

# Efficient Document-level Event Extraction via Pseudo-Trigger-aware Pruned Complete Graph

Tong Zhu<sup>1</sup>, Xiaoye Qu<sup>2</sup>, Wenliang Chen<sup>1\*</sup>, Zhefeng Wang<sup>2</sup>,  
 Baoxing Huai<sup>2</sup>, Nicholas Yuan<sup>2</sup> and Min Zhang<sup>1</sup>

<sup>1</sup>Institute of Artificial Intelligence, School of Computer Science and Technology,  
 Soochow University, China

<sup>2</sup>Huawei Cloud, China

tzhu7@stu.suda.edu.cn, {wlchen,minzhang}@suda.edu.cn

{quxiaoye,wangzhefeng,huaibaoxing}@huawei.com, nicholas.jing.yuan@gmail.com

## Abstract

Most previous studies of document-level event extraction mainly focus on building argument chains in an autoregressive way, which achieves a certain success but is inefficient in both training and inference. In contrast to the previous studies, we propose a fast and lightweight model named as PT-PCG. In our model, we design a novel strategy for event argument combination together with a non-autoregressive decoding algorithm via pruned complete graphs, which are constructed under the guidance of the automatically selected pseudo triggers. Compared to the previous systems, our system achieves competitive results with 19.8% of parameters and much lower resource consumption, taking only 3.8% GPU hours for training and up to 8.5 times faster for inference. Besides, our model shows superior compatibility for the datasets with (or without) triggers and the pseudo triggers can be the supplements for annotated triggers to make further improvements. Codes are available at <https://github.com/Spico197/DocEE>.

## 1 Introduction

Event Extraction (EE) aims at filling event tables with given texts. Different from sentence-level EE (SEE) which focuses on building trigger-centered trees [Chen *et al.*, 2015; Nguyen *et al.*, 2016; Liu *et al.*, 2018; Wadden *et al.*, 2019; Lin *et al.*, 2020], document-level EE (DEE) is to decode argument combinations from abundant entities across multiple sentences and fill these combinations into event record tables as shown in Figure 1, where annotated triggers are often not available [Yang *et al.*, 2018; Zheng *et al.*, 2019; Xu *et al.*, 2021].

One of the challenges for DEE is event argument combination without triggers. Triggers in DEE datasets are either absent or in a low annotation quality since the large-scale datasets are usually generated via distantly supervised (DS) alignment with existing knowledge bases (KB) [Mintz *et al.*,

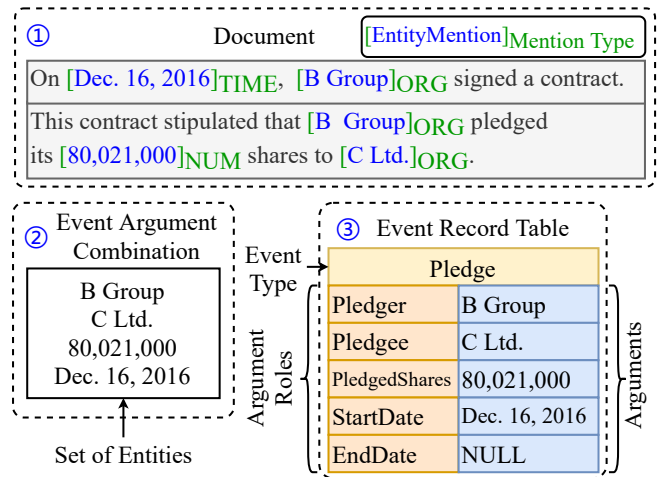


Figure 1: Examples of event extraction.

2009; Chen *et al.*, 2017]. Consequently, the absence of triggers drives the development of trigger-free argument combination methods in recent studies. Yang *et al.* [2018] first identify a key sentence and then fill the event record table by finding arguments near the sentence, while partial global features and arguments are still missing. Zheng *et al.* [2019] and Xu *et al.* [2021] fully utilize the global context and make significant improvements by building an directed acyclic graph (DAG). However, such DAG-based methods require massive computing resources as they rely on an autoregressive fashion to decode argument combinations, which is inefficient in both training and inference for long documents. Meanwhile, building DAGs consumes lots of memories to store previous paths. As a result, it takes **almost one week with a minimum of four 32GB GPUs to train such DAG-based DEE models, and the inference speed is also extremely slow.**

Considering the above challenge in DEE and the speed & memory consuming problems in the DAG-based methods, in this paper, we aim to 1) propose a universal event argument combination strategy that works on both trigger-aware and trigger-free DEE; and 2) provide a blazing fast and lightweight model for document-level event extraction.

\*Corresponding author

We propose a novel non-autoregressive approach named as **Pseudo-Trigger-aware Pruned Complete Graph (PTPCG)**. Specifically, we formulate each argument combination as a pruned complete graph, where the important arguments are identified and treated as a group of pseudo triggers with bidirectional connections to each other and other ordinary arguments are linked from these pseudo triggers in a directed manner. Based on the pruned complete graph with the pseudo triggers, we design an efficient algorithm with the non-autoregressive decoding strategy for event argument combination extraction.

The experiments results show that our PTPCG can reach competitive results with only 19.8% parameters of DAG-based SOTA models, just taking 3.8% GPU hours to train and up to 8.5 times faster in terms of inference. Besides, our PTPCG is highly flexible and scalable that can be used as a general architecture for non-autoregressive trigger-based event extraction. If only one pseudo trigger is selected for each combination, the pruned complete graph becomes a trigger-centered tree like SEE. Furthermore, pseudo triggers can be adopted as supplements to enhance the annotated-trigger-based methods.

In summary, our contributions include:

- We propose a novel non-autoregressive event argument combination paradigm based on pruned complete graph with pseudo triggers, which is compatible in document-level event extraction with (or without) triggers.
- Our model is fast and lightweight for end-to-end document-level event extraction and we conduct extensive experiments to show the efficiency and efficacy.
- To our best knowledge, our present approach is the first work that explores the effects of using some arguments as pseudo triggers in DEE, and we design a metric to help select a group of pseudo triggers automatically. Furthermore, such metric can also be used for measuring the quality of annotated triggers in DEE.

## 2 Methodology

As shown in Figure 2, our model can be divided into four components: 1) **Event detection** performs multi-label classification to identify all possible event types. 2) **Entity extraction** extracts all entities from the documents and encodes these entities into dense vectors. 3) **Combination extraction** builds pruned complete graphs and decodes argument combinations from such graphs. 4) **Event record generation** combines the results of event types and extracted argument combinations to generate the final event records.

### 2.1 Event Detection

For a document  $\mathcal{D}$ , a bidirectional long short-term memory (BiLSTM) network [Hochreiter and Schmidhuber, 1997] is used to encode each sentence  $s_i$  into token-wise hidden states  $(h_i^{(1)}, h_i^{(2)}, \dots, h_i^{(|s_i|)})$ , where  $h_i^{(j)} \in \mathbb{R}^{d_h}$  is the concatenation of two direction representations  $\overrightarrow{h_i^{(j)}} \parallel \overleftarrow{h_i^{(j)}}$  and  $|s_i|$  is the length of  $i$ -th sentence. The last hidden states in each direction are concatenated to get the sentence representation  $g_i \in \mathcal{G}$  as  $g_i = \overrightarrow{h_i^{(|s_i|)}} \parallel \overleftarrow{h_i^{(0)}}$ .

Subsequently, we follow Doc2EDAG [Zheng *et al.*, 2019] and use randomly initialized event queries with multi-head attention on  $\mathcal{G}$  to make a binary classification for each event type. The loss function of event detection component  $\mathcal{L}_{\text{det}}$  is defined as a binary cross entropy loss.

### 2.2 Entity Extraction

The entity extraction task can be formulated as a sequence tagging task in BIO scheme [Zheng *et al.*, 2019]. To enhance the entity recognition, we add more *money*, *date*, *percentage ratio* and *shares* entities into the dataset by simple regular expression matching. To deal with the document-level entity extraction, we first split the whole document into sentences and perform a sentence-level mention extraction. Then, we use BiLSTM with CRF [Lample *et al.*, 2016] to extract entity mentions with the same BiLSTM used in event detection. The training objective of entity extraction is to minimize the negative log-likelihood loss  $\mathcal{L}_{\text{ent}}$  of CRF for each sentence.

For all tokens of a mention, a max-pooling operation is applied on token-level representations to get the mention representation  $\tilde{m}_j$ . As mention types have been proved effective for downstream sub-modules [Zheng *et al.*, 2019; Xu *et al.*, 2021], we convert the predicted discrete mention types into vectors by looking up an embedding table. After concatenating  $\tilde{m}_j$  and type embeddings  $l_j$ , we get the final mention representation  $m_j = \tilde{m}_j \parallel l_j \in \mathbb{R}^{d_a}$ , where  $d_a = d_h + d_l$  and  $d_l$  denotes the dimension of  $l_j$ . Finally, all the mentions for an entity are aggregated to get the entity representation  $\hat{e}_i$  via another max-pooling. For better modeling entity semantics and the latent connections for combination extraction, an additional BiLSTM layer is applied to get the set of entity representations  $\mathcal{E} = \{\tilde{e}_i\}_{i=1}^{|\mathcal{E}|}$ .

### 2.3 Combination Extraction

In this section, we introduce the details of selecting pseudo triggers, building pruned complete graphs, and decoding combinations from the graphs.

#### Pseudo Trigger Selection

It is hard to annotate conventional triggers manually in documents due to the document length and dataset scale. We instead select a group of pseudo triggers for each event type in an automatic way.

Empirically, the triggers are keywords that play two roles: 1) triggers can be used to identify combinations; 2) triggers are fingerprints that can distinguish different combinations. Combinations are made of arguments and we can extract a specific combination by finding all the corresponding arguments. To this end, we design an importance metric that evaluates the possibility of whether a group of arguments can serve as the pseudo triggers. In general, we first select a group of argument roles as candidates for each event type according to the importance scores, and take the corresponding arguments as the pseudo triggers.

Formally, the importance score is obtained by the existence and the distinguishability. For a subset of predefined argument roles in type  $t_i$ ,  $\mathcal{R} = \{r_j\}_{j=1}^{|\mathcal{R}|}$  are selected as the pseudo trigger candidates.  $|\mathcal{R}|$  is a hyper-parameter that denotes the

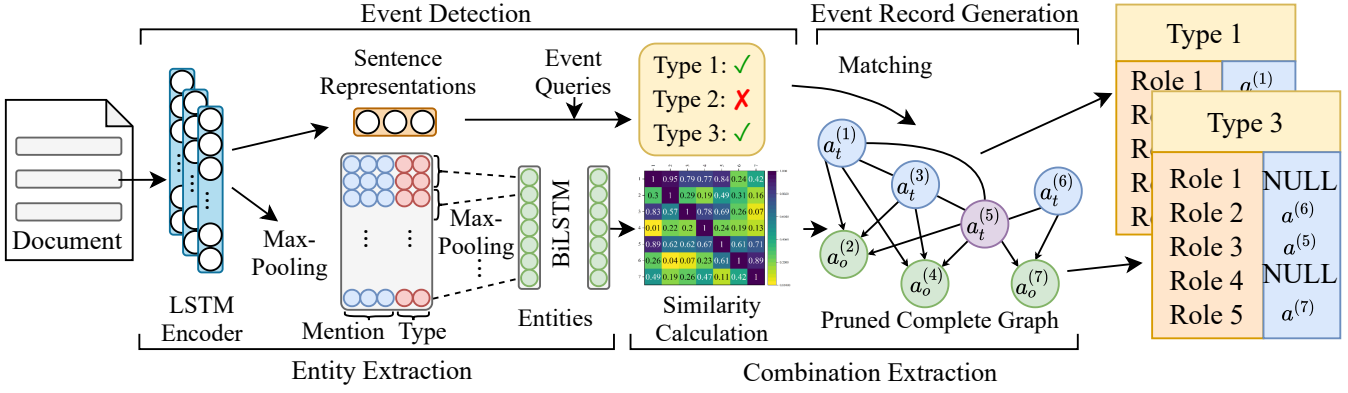


Figure 2: Overview of our PTPCG. Event types and entity mentions are first extracted with a shared LSTM encoder, then similarities between entity pairs are calculated to help recover the adjacent matrix of pruned complete graph. After combinations are decoded from the adjacent matrix, the extracted types are paired with combinations to generate final event records.

number of selected pseudo triggers for each argument combination. The existence measures whether the arguments of  $\mathcal{R}$  can identify combinations.  $N_e^{(\mathcal{R})}$  is the number of event records that at least one corresponding argument of  $\mathcal{R}$  is not NULL, and  $N^{(i)}$  is the number of total records of  $t_i$ . The distinguishability is defined to satisfy that the triggers can distinguish different combinations, where  $N_u^{(\mathcal{R})}$  is the number of records that the arguments of  $\mathcal{R}$  do not appear in other records in the same document. With the multiplication of existence and distinguishability, pseudo triggers are selected by picking the candidate with the highest importance score.

$$\text{Existence}(\mathcal{R}) = \frac{N_e^{(\mathcal{R})}}{N^{(i)}}, \quad \text{Distinguish}(\mathcal{R}) = \frac{N_u^{(\mathcal{R})}}{N^{(i)}} \quad (1)$$

$$\text{Importance}(\mathcal{R}) = \text{Existence}(\mathcal{R}) \times \text{Distinguish}(\mathcal{R})$$

### Pruned Complete Graph Construction

Based on the DEE task setting and data analysis, we propose an assumption that *arguments in the same combination are close to each other in the semantic space*. Following this assumption, we take the pseudo triggers as the core of argument combinations and formulate each combination as a pruned complete graph. As shown in the pruned complete graph of Figure 2, for any two arguments as pseudo triggers  $a_t^{(i)}$  and  $a_t^{(j)}$  in the same combination, they are bidirectionally connected, where the adjacent matrix  $y_A^{(i,j)} = y_A^{(j,i)} = 1$ . For a pseudo trigger  $a_t^{(i)}$  and an ordinary argument  $a_o^{(j)}$  in the same combination, they are connected with a directional link and  $y_A^{(i,j)} = 1$ . Besides, each argument  $a^{(i)}$  has a self-loop connection where  $y_A^{(i,i)} = 1$ . Other entries in  $y_A$  are zeros, where entities that do not participate in any combination are isolated nodes in the graph.

After obtaining entity representations, a dot scaled similarity function (Eqn 2) is applied to estimate their semantic

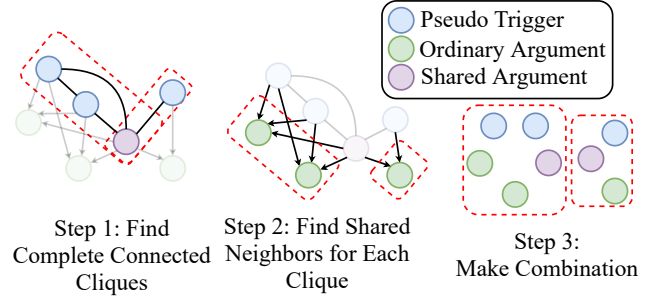


Figure 3: Combination decoding (best viewed in color).

distances:

$$\tilde{e}_i = e_i \times W_s^\top + b_s, \quad \tilde{e}_j = e_j \times W_e^\top + b_e \quad (2)$$

$$\tilde{A}_{i,j} = \text{sigmoid} \left( \tilde{e}_i^\top \tilde{e}_j / \sqrt{d_h} \right)$$

where  $\tilde{A}$  denotes the similarity matrix,  $W_s, W_e \in \mathbb{R}^{d_a \times d_a}$  and  $b_s, b_e \in \mathbb{R}^{d_a}$  are trainable parameters for semantic space linear projection.

In training, we use binary cross entropy function to formulate the combination loss  $\mathcal{L}_{\text{comb}}$ .

To predict the binary adjacent matrix  $A$  of the pruned complete graph for further decoding, the threshold  $\gamma$  is used here (Eqn 3).

$$A_{i,j} = \begin{cases} 1 & \tilde{A}_{i,j} \geq \gamma \\ 0 & \text{otherwise} \end{cases} \quad (3)$$

### Non-Autoregressive Combination Decoding

Event argument combinations are extracted based on the predicted adjacent matrix  $A$  with a non-autoregressive decoding algorithm.

First, all the pseudo triggers are identified based on nodes' out-degree and each pseudo trigger group is recognized as a clique. If the out-degree of an entity is greater than 0 except for self-loop, then the entity is treated as a pseudo trigger. For  $|\mathcal{R}| = 1$ , all the combinations are pseudo-trigger-centered

trees, where each combination is made of a pseudo trigger with its neighbors. Otherwise, Bron-Kerbosch (BK) algorithm [Bron and Kerbosch, 1973] is applied first to find all possible cliques (step 1 in Figure 3). To apply BK algorithm, the links between arguments must be undirected, thus we first extract all bidirectional links as undirected input.

As shown in Figure 3, the next step is to find the ordinary (non-pseudo-trigger) arguments in the combination. We further exploit all the neighbors of each pseudo trigger in the clique. After that, an intersection operation is performed to find commonly shared ordinary arguments. The combination is consist of a pseudo trigger clique and their commonly shared ordinary arguments. For those extreme records that have only one argument in the combinations, all predicted entities are aggregated together as a default combination.

The document length does not affect the event record generation speed, and the time complexity of our event records generation is polynomial ( $\mathcal{O}(N_t \times N_c)$ ), while DAG-based model is  $\mathcal{O}(N_t \times N_s \times N_r)$ , where  $N_t$ ,  $N_c$ ,  $N_s$ ,  $N_r$  are the number of predicted types, combinations, predicted spans and the average number of roles per type. In most cases,  $N_s \times N_r \gg N_c$ , so our PTPCG consistently holds the speed advantage.

## 2.4 Event Records Generation

After the set of combinations  $\mathcal{C}$  are obtained from the pruned complete graphs, the next step is to fill these combinations into event tables. First, all the combinations should match with event types. Since event detection is a multi-label classification task, there may be more than one type prediction. For all type predictions  $\mathcal{T}_p = \{t_j\}_{j=1}^{|\mathcal{T}_p|}$  and combinations  $\mathcal{C}$ , we perform a Cartesian product and get all type-combination pairs  $\{ \langle t_j, c_k \rangle \mid 1 \leq j \leq |\mathcal{T}_p|, 1 \leq k \leq |\mathcal{C}| \}$ .

For each pair  $\langle t_j, c_k \rangle$ , we use an event-relevant feed-forward network (FFN) as classifier to get possible role results for all arguments  $\mathcal{E}_k$  in  $c_k$ , and use sigmoid function to get the probabilities  $p_{role}^{(j)}(r_j | c_k)$ . The loss  $\mathcal{L}_{role}$  is calculated by a binary cross entropy function.

In the role classification task, an entity can act as more than one role in a record, while a role in a table can only be filled with one entity. Following this setting, we take the entity  $e_{i^*}^{(k)}$  as the argument of role  $r_j^{(q)}$  if and only if it satisfies the constraint in Eqn 4. If  $\forall j \in [1, |\mathcal{T}_p|], p_{role}^{(j)}(r_j | c_k) < 0.5$ , then the pair is recognized as an invalid candidate and will be dropped.

$$i^* = \operatorname{argmax}_q p_{role}^{(j)}(r_j^{(q)} | c_k) \quad (4)$$

## 2.5 Optimization

Our PTPCG is an end-to-end model with joint training and scheduled sampling [Bengio *et al.*, 2015] strategy. The overall loss is a weighted sum of all losses as below:

$$\mathcal{L} = \alpha_1 \mathcal{L}_{\text{det}} + \alpha_2 \mathcal{L}_{\text{ent}} + \alpha_3 \mathcal{L}_{\text{comb}} + \alpha_4 \mathcal{L}_{\text{role}} \quad (5)$$

where  $\alpha_1, \alpha_2, \alpha_3, \alpha_4$  are hyper-parameters to reduce the unbalanced loss effect.

## 3 Experiments

### 3.1 Datasets

We use ChFinAnn [Zheng *et al.*, 2019] and DuEE-fin [Li, 2021] datasets to make fair comparisons across all methods.

(1) **ChFinAnn** is by far the largest DEE dataset constructed by distantly supervised alignment **without trigger annotations**, which is widely used in the previous studies. This dataset contains 32k financial announcements and 48k event records, in which 29% documents have more than one record and 98% of records have arguments scattered across different sentences. On average a document contains 20 sentences and the longest document contains 6.2k Chinese characters. (2) **DuEE-fin** is another dataset for DEE **with trigger annotations**. It contains 13 event types and 11.7k documents, where the test set is evaluated online. Each record in DuEE-fin has a labeled trigger word without specific positions, and 36% of records share the same trigger in a document.

### 3.2 Experiment Settings

We choose the hyper-parameters of our system according to the performance on the development set in ChFinAnn. In PTPCG, we use 2 layers of shared BiLSTM for event detection and entity extraction, and another 2 layers of BiLSTM for entity encoding. We use the same vocabulary as [Zheng *et al.*, 2019] and randomly initialize all the embeddings where  $d_h=768$  and  $d_l=32$ . Adam [Kingma and Ba, 2015] optimizer is used with a learning rate of 5e-4 and the mini-batch size is 64. The weights in Equation 5 are 0.05, 1.0, 1.0, and 1.0, and  $\gamma$  in Equation 3 is 0.5. Following the setting in Zheng *et al.* [2019], we train our models for 100 epochs and select the checkpoint with the best F1 score on the development set to evaluate on the test set.

### 3.3 Baselines and Metrics

**Baselines.** (1) **DCFEE** [Yang *et al.*, 2018] has two variants here: DCFEE-O extracts only one record from one document while DCFEE-M extracts records as much as possible. (2) **Doc2EDAG** [Zheng *et al.*, 2019] constructs records as argument chains (DAG) and uses an auto-regressive way to extract final results. (3) **GreedyDec** is a baseline from Zheng *et al.* [2019] which fills one event table greedily. (4) **GIT** [Xu *et al.*, 2021] is another variant of Doc2EDAG and utilizes a graph neural network to further encode entities and add more features in DAG generation.

**Metrics.** We use the same evaluation setting as in the previous studies [Zheng *et al.*, 2019; Xu *et al.*, 2021]. For each predicted record, a golden record is selected without replacement by matching the record that has the same event type and the most shared arguments, and F1 scores are calculated by comparing arguments. DuEE-fin uses an online evaluation fashion and the platform only provides micro-averaged scores.

### 3.4 Main Results

As shown in Table 1, PTPCG achieves better or competitive results than the previous systems. On the ChFinAnn dataset, our PTPCG achieves the best scores evaluated on

Model	#Params (w/o Emb)	GPU Hours	ChFinAnn-Single			ChFinAnn-All			DuEE-fin w/o Tgg			DuEE-fin w/ Tgg		
			P	R	F1	P	R	F1	P	R	F1	P	R	F1
DCFEE-O*	32M (16M)	192.0	73.2	71.6	72.4	69.7	57.8	63.2	56.2	48.2	51.9	51.9	49.6	50.7
DCFEE-M*	32M (16M)	192.0	64.9	71.7	68.1	60.1	61.3	60.7	38.7	52.3	44.5	37.3	48.6	42.2
GreedyDec*	64M (48M)	604.8	83.9	77.3	80.4	81.9	51.2	63.0	59.6	41.8	49.1	59.0	42.1	49.2
Doc2EDAG*	64M (48M)	604.8	83.2	89.3	86.2	81.1	77.0	79.0	66.7	50.0	57.2	67.1	51.3	<b>58.1</b>
GIT*	97M (81M)	633.6	85.0	88.7	86.8	82.4	<b>77.6</b>	<b>79.9</b>	<b>68.2</b>	43.4	53.1	<b>70.3</b>	46.0	55.6
PTPCG $_{ \mathcal{R} =1}$	32M (16M)	<b>24.0</b>	<b>86.3</b>	<b>90.1</b>	<b>88.2</b>	<b>83.7</b>	75.4	79.4	64.5	<b>56.6</b>	<b>60.3</b>	63.6	<b>53.4</b>	<b>58.1</b>

Table 1: Main results. #Params is estimated on the ChFinAnn dataset. w/o Emb means the number of parameters without vocabulary embeddings. All models are trained with 100 epochs. -Single denotes the evaluation results on documents with only one event record. Tgg denotes the manually annotated triggers. \*We reproduce the results using their open-source codes.

ChFinAnn-Single. As for ChFinAnn-All, our model outperforms Doc2EDAG and is very close to GIT. On the DuEE-fin dataset, PTPCG achieves the best F1 scores whether or not trigger words are provided, demonstrating good compatibility and universality. Our PTPCG outperforms GIT and improves the overall F1 with 7.2% w/o manually annotated triggers. For DuEE-fin w/ Tgg, we find our approach declines from 60.3% to 58.1%. After analyzing the dataset, we find that the manually annotated triggers have less importance scores (62.9) than our pseudo triggers (83.8,  $|\mathcal{R}|=1$ ). Although adding annotated triggers when  $|\mathcal{R}|=1$  can boost the importance score to 93.7, it is a trade-off between the number of triggers (annotated & pseudo) and the adjacent matrix prediction. We will discuss more in Section 3.7.

### 3.5 Comparison on Model Parameters and Speed

**Comparison on the Number of Parameters.** As shown in Table 1, among all the models, PTPCG is the most lightweight one and is in the same scale with DCFEE, while PTPCG outperforms DCFEE-O with 16.2% absolute gain in F1 scores on ChFinAnn-All. Without considering 16M vocabulary embedding parameters, PTPCG takes only 19.8% parameters of GIT and reaches competitive results.

**Comparison on Speed.** Benefiting from the lightweight architecture design and non-autoregressive decoding style, our PTPCG is fast in both training and inference.

Comparing the training time as shown in Table 1, it takes 633.6 GPU hours to train GIT, while PTPCG is 26.4 times faster and takes only 24.0 GPU hours. These facts indicate that PTPCG is more efficient in training and requires much lower computation resource cost than other methods, and the whole training process has been reduced from almost a week with 4 NVIDIA V100 GPUs to just one day with 1 GPU.

According to the inference speed test in Figure 4, PTPCG is more scalable than other models. With the growth of batch size, PTPCG becomes faster and finally stabilized near 125 documents per second, while Doc2EDAG and GIT are barely growing with batch size, peaking at 19 and 15 docs/s respectively. PTPCG is up to 7.0 and 8.5 times faster compared to Doc2EDAG and GIT. As an enhanced version of Doc2EDAG, GIT is 21.2% slower than Doc2EDAG on average, and raises OOM error on a 32GB memory GPU when batch size is 128. To further check the effect of different encoders, we substitute all BiLSTM layers into transformer encoders and create

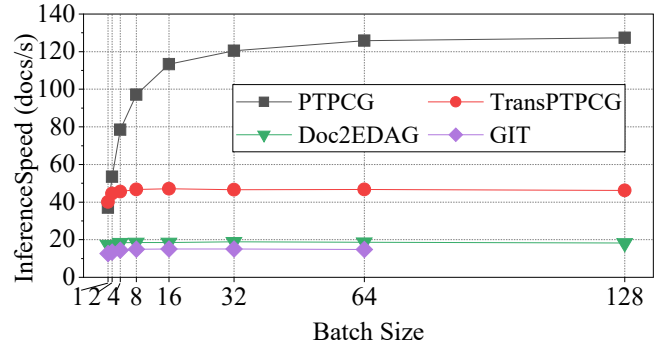


Figure 4: Inference speed comparison with baselines (left) and with different  $|\mathcal{R}|$  (right) with 1 NVIDIA V100 GPU for all models.

the TransPTPCG model with the same scale to Doc2EDAG parameters. The results in Figure 4 show that TransPTPCG is up to 2.5 times faster than Doc2EDAG, validating the advantage of non-autoregressive combination decoding.

### 3.6 Pseudo Triggers as Supplements

PTPCG is capable of handling with or without manually annotated triggers, and the automatically selected pseudo triggers can be supplements to enhance the performance.

As shown in Table 2, we find the pseudo triggers could assist the trigger words to boost the importance score from 62.9 to 93.7, and the results show that this helps identify combinations and bring 0.7% improvement in offline dev evaluation and 0.8% improvement in the online test set.

To further validate the effectiveness of importance scores for pseudo trigger selection, we select groups of pseudo triggers with the middle and the lowest importance scores instead of the highest ones and analyze the results in Table 3. The results show that there is a positive correlation between the

Pseudo Trigger	Impt.	Dev			Online Test		
		P	R	F1	P	R	F1
×	62.9	<b>76.5</b>	56.8	65.2	<b>69.8</b>	48.5	57.3
✓	<b>93.7</b>	70.5	<b>62.3</b>	<b>66.2</b>	63.6	<b>53.4</b>	<b>58.1</b>

Table 2: PTPCG results on DuEE-fin w/ annotated triggers. Impt denotes the importance score.

Impt.	ChFinAnn			Impt.	DuEE-fin w/o Tgg		
	P	R	F1		P	R	F1
88.3	83.7	<b>75.4</b>	<b>79.4</b>	83.8	66.7	<b>54.6</b>	<b>60.0</b>
62.5	88.7	60.9	72.2	37.0	65.5	49.6	56.4
22.4	<b>90.5</b>	59.5	71.8	15.0	<b>74.2</b>	45.3	56.3

Table 3: PTPCG results with different importance scores.

importance and overall scores. Nevertheless, the highest importance score (88.3%) is not equal to 100.0%, which may limit the upper bound of decoding. We will explain more about error analysis in Section 3.7.

### 3.7 Error Analysis

Although pruned complete graph structure is efficient for training and inference, it is not perfect in similarity calculation (adjacent matrix prediction) and combination decoding. Here we analyze the upper bounds of the combination decoding algorithm and discuss the future directions.

Results in Table 4 show that models with the lower number of pseudo triggers are better. However, models with more pseudo triggers have higher upper bounds for decoding and have higher importance scores. Why models with more pseudo triggers have greater importance but still result in lower performance? Results in Table 3 and 4 may answer this question. Models with the same  $|\mathcal{R}|$  has a strong correlation that higher importance brings higher metric scores which validates the effectiveness of the pseudo trigger selection strategy based on importance scores. However, more pseudo triggers bring more links to connect, and the model may be not robust enough for predicting each connection correctly, leading to an adjacent accuracy decline.

Overall, it is a trade-off that more pseudo triggers improves the upper bounds and reduces combination error rates, but also brings new challenges to recover connections between entities. We believe it is the future direction to improve the similarity calculation and adjacent matrix prediction.

$ \mathcal{R} $	Impt.	SE	ME	TotE	#links	Adj Acc.	F1
1	88.3	5.0	37.5	14.6	10,502	65.8	79.4
2	95.7	1.0	20.4	6.7	23,847	59.1	77.7
3	97.2	0.9	18.0	5.9	55,961	56.7	74.9
4	97.6	0.5	16.9	5.3	75,334	58.2	74.0
5	97.8	0.4	13.9	4.4	88,752	59.5	73.1
all	97.8	0.2	13.4	4.1	140,989	60.1	69.5

Table 4: Analysis of combination decoding errors on ChFinAnn. SE, ME, TotE are error upper bounds of single-record documents, multiple-record documents, and overall error respectively. #links is the number of connected links among all graphs in the test set. Adj Acc is the accuracy of adjacent matrix prediction given golden entities. F1 is micro-averaged.

## 4 Related Work

Fully annotated datasets are usually small in scale [Ebner *et al.*, 2020; Li *et al.*, 2021]. DS-constructed datasets are large

in scale, but they can hardly match triggers to records, so triggers are likely to be absent. Unlike trigger-based methods [Pouran Ben Veyseh *et al.*, 2021; Du and Cardie, 2020], Liu *et al.* [2019] argue that event types can be detected without triggers. Yang *et al.* [2018] extract event records using entities in a window of key sentences. It is efficient but misses lots of information from other sentences. To fully utilize all entities in the whole document, Doc2EDAG [Zheng *et al.*, 2019] formulates the argument combination to be a directed acyclic graph (DAG), making great progress in DEE. GIT [Xu *et al.*, 2021] is a variant of Doc2EDAG where a graph neural network is added to help entity encoding and further exploit the global memory mechanism during decoding. DAG extracts combinations in an auto-regressive way, which is very time consuming and needs huge space in global memory module to store all previous paths. Besides, Doc2EDAG and GIT are both large models, where Doc2EDAG and GIT use 12 and 16 layers of transformer encoders. To train such models, a minimum of four 32GB GPUs would have to be run for almost a week. Huang and Jia [2021] utilize BERT [Devlin *et al.*, 2019] as sentence representations and exploit relations between sentences via graph neural networks [Veličković *et al.*, 2018] to help identify fixed number of combinations. To speed up DEE, we propose a novel non-autoregressive decoding strategy and boost the training & inference speed with competitive results compared with the previous works.

## 5 Conclusion

Pursuing fast and general document-level event extraction (DEE), we propose a non-autoregressive model named as PTPCG. For DEE without triggers, we first select a group of pseudo triggers to build pruned complete graphs, and then train a lightweight model to extract all possible combinations, which enables non-autoregressive decoding and is up to 8.5x faster compared to SOTA models. For DEE with annotated triggers, pseudo triggers also show the power to make improvement and are even better than annotated trigger-only method. In summary, our model costs less resources yet achieves better or comparable results, compared with the previous systems.

Although PTPCG is not perfect in adjacent matrix prediction, we believe it has the ability to combine different DEE tasks and needs more explorations.

## Acknowledgments

This work was supported by the National Natural Science Foundation of China (Grant No. 61936010 and 61876115), the Priority Academic Program Development of Jiangsu Higher Education Institutions, and the joint research project of Huawei Cloud and Soochow University. We would also like to thank the anonymous reviewers for their valuable comments.

## References

[Bengio *et al.*, 2015] Samy Bengio, Oriol Vinyals, Navdeep Jaitly, and Noam Shazeer. Scheduled sampling for sequence prediction with recurrent neural networks. In *Proc. of NeurIPS*, pages 1171–1179, 2015.

- [Bron and Kerbosch, 1973] Coen Bron and Joep Kerbosch. Algorithm 457: Finding all cliques of an undirected graph. *Commun. ACM*, 16(9):575–577, September 1973.
- [Chen *et al.*, 2015] Yubo Chen, Liheng Xu, Kang Liu, Daojian Zeng, and Jun Zhao. Event extraction via dynamic multi-pooling convolutional neural networks. In *Proc. of ACL-IJCNLP*, pages 167–176, Beijing, China, July 2015.
- [Chen *et al.*, 2017] Yubo Chen, Shulin Liu, Xiang Zhang, Kang Liu, and Jun Zhao. Automatically labeled data generation for large scale event extraction. In *Proc. of ACL*, pages 409–419, Vancouver, Canada, July 2017.
- [Devlin *et al.*, 2019] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proc. of NAACL-HLT*, 2019.
- [Du and Cardie, 2020] Xinya Du and Claire Cardie. Event extraction by answering (almost) natural questions. In *Proc. of EMNLP*, pages 671–683, Online, November 2020. Association for Computational Linguistics.
- [Ebner *et al.*, 2020] Seth Ebner, Patrick Xia, Ryan Culkin, Kyle Rawlins, and Benjamin Van Durme. Multi-sentence argument linking. In *Proc. of ACL*, pages 8057–8077, 2020.
- [Hochreiter and Schmidhuber, 1997] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [Huang and Jia, 2021] Yusheng Huang and Weijia Jia. Exploring sentence community for document-level event extraction. In *Findings of EMNLP*, pages 340–351, Punta Cana, Dominican Republic, November 2021.
- [Kingma and Ba, 2015] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *Proc. of ICLR*, 2015.
- [Lample *et al.*, 2016] Guillaume Lample, Miguel Ballesteros, Sandeep Subramanian, Kazuya Kawakami, and Chris Dyer. Neural architectures for named entity recognition. In *Proc. of NAACL*, pages 260–270, San Diego, California, June 2016.
- [Li *et al.*, 2021] Sha Li, Heng Ji, and Jiawei Han. Document-level event argument extraction by conditional generation. In *Proc. of NAACL*, pages 894–908, Online, June 2021.
- [Li, 2021] X. et al Li. Ducee-fin: A document-level event extraction dataset in the financial domain released by baidu. <https://aistudio.baidu.com/aistudio/competition/detail/46>, 2021.
- [Lin *et al.*, 2020] Ying Lin, Heng Ji, Fei Huang, and Lingfei Wu. A joint neural model for information extraction with global features. In *Proc. of ACL*, pages 7999–8009, Online, July 2020.
- [Liu *et al.*, 2018] Xiao Liu, Zhunchen Luo, and Heyan Huang. Jointly multiple events extraction via attention-based graph information aggregation. In *Proc. of EMNLP*, pages 1247–1256, Brussels, Belgium, October-November 2018.
- [Liu *et al.*, 2019] Shulin Liu, Yang Li, Feng Zhang, Tao Yang, and Xinpeng Zhou. Event detection without triggers. In *Proc. of NAACL*, pages 735–744, Minneapolis, Minnesota, June 2019.
- [Mintz *et al.*, 2009] Mike Mintz, Steven Bills, Rion Snow, and Daniel Jurafsky. Distant supervision for relation extraction without labeled data. In *Proc. of ACL-IJCNLP*, pages 1003–1011, 2009.
- [Nguyen *et al.*, 2016] Thien Huu Nguyen, Kyunghyun Cho, and Ralph Grishman. Joint event extraction via recurrent neural networks. In *Proc. of NAACL*, pages 300–309, San Diego, California, June 2016.
- [Pouran Ben Veyseh *et al.*, 2021] Amir Pouran Ben Veyseh, Minh Van Nguyen, Nghia Ngo Trung, Bonan Min, and Thien Huu Nguyen. Modeling document-level context for event detection via important context selection. In *Proc. of EMNLP*, pages 5403–5413, Online and Punta Cana, Dominican Republic, November 2021. Association for Computational Linguistics.
- [Veličković *et al.*, 2018] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. Graph attention networks. In *Proc. of ICLR*, 2018.
- [Wadden *et al.*, 2019] David Wadden, Ulme Wennberg, Yi Luan, and Hannaneh Hajishirzi. Entity, relation, and event extraction with contextualized span representations. In *Proc. of EMNLP-IJCNLP*, pages 5784–5789, Hong Kong, China, November 2019.
- [Xu *et al.*, 2021] Runxin Xu, Tianyu Liu, Lei Li, and Baobao Chang. Document-level event extraction via heterogeneous graph-based interaction model with a tracker. In *Proc. of ACL-IJCNLP*, pages 3533–3546, Online, August 2021.
- [Yang *et al.*, 2018] Hang Yang, Yubo Chen, Kang Liu, Yang Xiao, and Jun Zhao. DCFEE: A document-level Chinese financial event extraction system based on automatically labeled training data. In *Proc. of ACL*, pages 50–55, Melbourne, Australia, July 2018.
- [Zheng *et al.*, 2019] Shun Zheng, Wei Cao, Wei Xu, and Jiang Bian. Doc2EDAG: An end-to-end document-level framework for Chinese financial event extraction. In *Proc. of EMNLP-IJCNLP*, pages 337–346, Hong Kong, China, November 2019.