

GEGA: Graph Convolutional Networks and Evidence Retrieval Guided Attention for Enhanced Document-level Relation Extraction

Anonymous ACL submission

Abstract

Document-level relation extraction (DocRE) aims to extract relations between entities from unstructured document text. Compared to sentence-level relation extraction, it requires more complex semantic understanding from a broader text context. Currently, some studies are utilizing logical rules within evidence sentences to enhance the performance of DocRE. However, in the data without provided evidence sentences, researchers often obtain a list of evidence sentences for the entire document through evidence retrieval (ER). Therefore, DocRE suffers from two challenges: firstly, the relevance between evidence and entity pairs is weak; secondly, there is insufficient extraction of complex cross-relations between long-distance multi-entities. To overcome these challenges, we propose GEGA, a novel model for DocRE. The model leverages graph neural networks to construct multiple weight matrices, guiding attention allocation to evidence sentences. It also employs multi-scale representation aggregation to enhance ER. Subsequently, we integrate the most efficient evidence information to implement both fully supervised and weakly supervised training processes for the model. We evaluate the GEGA model on three widely used benchmark datasets: DocRED, Re-DocRED, and Revisit-DocRED. The experimental results indicate that our model has achieved comprehensive improvements compared to the existing SOTA model.

1 Introduction

Relation extraction (RE) is a crucial technology used to automatically identify and classify semantic relations between entities in natural language texts. Existing relation extraction tasks can be divided into two types: sentence-level relation extraction and document-level relation extraction (DocRE) (Peng et al., 2017; Verga et al., 2018). In sentence-level relation extraction datasets, each data entry contains only one sentence, and there is a single

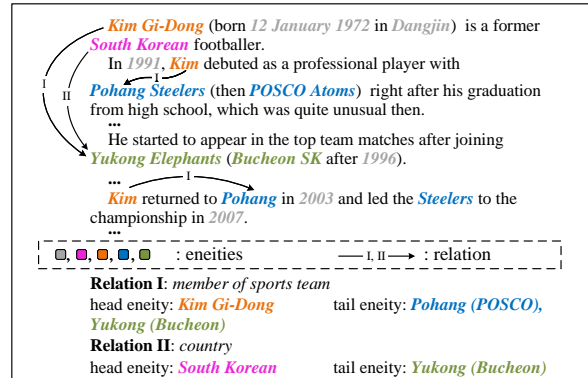


Figure 1: Examples of relations from DocRED, with entities marked in different colors, and curves indicating various relations between the entities.

entity pair within the sentence for which the relation needs to be predicted. In contrast, DocRE datasets contain multiple sentences per data entry, corresponding to multiple entity pairs whose relations need to be predicted. As depicted in Figure 1, each entity pair may appear multiple times within the document and may have different relation types (Yao et al., 2019), necessitating the analysis of a larger contextual scope to determine the relation for each entity pair.

Furthermore, assuming there are n entities in a data entry, predicting the relation for an entity pair involves pairing the anchor entity with the remaining $n-1$ entities one by one. This approach results in significant unnecessary memory overhead, as most pairs of entities do not have any relation between them.

Existing methods can be mainly divided into three categories (Zhou et al., 2021): sequence-based, graph-based, and transformer-based. Sequence-based models commonly employ pre-trained language models to produce word embeddings and character embeddings, transforming sequences of words or characters within texts into vector representations for processing (Ye et al.,

2020; Tang et al., 2020). Models based on dependency graphs utilize dependency information to construct document-level graph (Zeng et al., 2021; Li et al., 2021; Zhang et al., 2023b,a), which are then processed through graph neural networks for inference. Transformer-based models utilize the self-attention mechanism to refine the representation of each word by assessing its contextual relations with all other words in the text (Xiao et al., 2022; Ma et al., 2023).

The aforementioned three types of relation extraction methods suffer from two limitations. First, relation extraction between entity pairs generally requires only a set of sentences as supporting evidence, without the need to focus on redundant irrelevant information. However, these methods utilize all the information in the long text for relation extraction. Therefore, Ma et al. (2023) proposed an evidence retrieval-based relation extraction method. However, this method retrieves a list of evidence information for the entire document, resulting in poor relevance between this information and the relational entity pairs. Second, in DocRE, multiple entities are discretely distributed across different sentences or even paragraphs. Extracting their relations requires fully learning and understanding the semantics of the long text at the document-level. Existing methods rely on dependency parsing to construct multidimensional graph structures for semantic understanding and relation reasoning. However, they perform poorly when faced with complex intersecting relations due to the large number of entity pairs involved.

To address the aforementioned two issues and the insufficient annotation of evidence sentences in the dataset (Yao et al., 2019), we propose a novel model for DocRE: GEGA. This model is trained under both fully supervised and weakly supervised settings. First, we utilize a complex model (Teacher) trained with full supervision to infer over distant supervision data, extracting evidence sentences and assigning token weights. Then, we use this weight information as supervisory signals to guide the training of a simplified model (Student). Finally, we fine-tune the student model to adapt it to specific tasks and datasets, thereby improving performance. In summary, this article has two contributions:

- (1) We propose a novel DocRE model, GEGA¹

¹The implementation code for GEGA can be obtained from the GitHub link: [Links covered up in blind reviewing]

(Graph Convolutional Networks and Evidence Retrieval Guided Attention). This model combines graph structures and Transformers to retrieve evidence sentences highly relevant to the relational entity pairs from the document, guiding the attention to assign higher weights to this evidence information, thereby enhancing the performance of relation extraction.

(2) Experiments conducted on the three public datasets DocRED, Re-DocRED and Revisit-DocRED show that GEGA can achieve the new SOTA² results on document-level relation extraction compared to existing methods under the same experimental settings.

2 Preliminary

2.1 Task formulation for DocRE and DocER

Let's assume we have a document D containing n entities $e = \{e_1, e_2, \dots, e_n\}$. Each entity e_i in the document has a corresponding position p_i , and there may exist relations between entities. Our objective is to extract a set of relations $R = \{r_1, r_2, \dots, r_m\}$ from document D , where each relation r_i can be represented by a triple (e_i, e_j, r_{ij}) , with r_{ij} being the relation label between entities e_i and e_j . Therefore, the task of DocRE can be formulated as follows: $R = \{(e_i, e_j, r_{ij}) \mid e_i, e_j \in E, i \neq j\}$, r_{ij} is the relation label predicted by the relation classifier based on the contextual information of entity pairs (e_i, e_j) .

In addition, Document-level Evidence Retrieval (DocER) aims to retrieve a list of evidence sentences $evi_{[0,1,\dots,n]}$ from document D to enhance relation extraction. Nowadays, researchers have extended relation triples (e_i, e_j, r_{ij}) by adding evidence sentence lists, resulting in relation quadruplets $(e_i, e_j, r_{ij}, evi_{[0,1,\dots,n]})$. Relations between entity pairs can be predicted solely using the sentences from the evidence lists, without relying on the entire document.

3 Related Work

Our work is built upon a substantial body of recent work on document-level RE and ER.

DOC-Relation Extraction (RE). Previous studies can be divided into three major categories (Zhou et al., 2021):

²We outperform all other methods using BERT as the pretraining language model on the DocRED leaderboard. Please refer to GEGA's submission: <https://codalab.lisn.upscayl.fr/competitions/365#results>

Sequence-based methods. Zeng et al. (2014); Cai et al. (2016); Tang et al. (2020); Yao et al. (2019), and Sorokin and Gurevych (2017) use methods such as Conditional Random Fields (CRF) or Recurrent Neural Networks (RNN, (Cho et al., 2014)) to accurately identify and label entities in text and predict the relations between these entities by learning the contextual sequential semantic information of each word. For example, Tang et al. (2020) use different neural network architectures to perform sequence encoding of the entire document to learn the semantic representation of entity pairs and extract the relations between them. Yao et al. (2019) employ BiLSTM to simultaneously consider the forward and backward information in the text sequence, learning and understanding the different semantic paths from one entity to another, thus inferring the relations between entity pairs.

Graph-based methods. Veličković et al. (2018); Christopoulou et al. (2019); Sahu et al. (2019) model entities and relations in documents as graph structures, then use Graph Neural Networks (GNNs) to learn the relations between entities. This approach effectively leverages structural information and global context among entities. Additionally, researchers have proposed hard pruning and soft pruning strategies for dependency tree structures to optimize the model’s speed and storage efficiency. Zhang et al. (2018) and Mandya et al. (2020) use hard pruning strategies to retain words near the shortest path between two entities, maximizing the removal of irrelevant content while integrating relevant information. Guo et al. (2019) proposed a soft pruning method that directly takes the full dependency tree as input and automatically learns how to selectively focus on relevant substructures that are useful for the relation extraction task. Subsequently, Li et al. (2021); Nan et al. (2020) proposed refined strategies to enhance cross-sentence relation reasoning by automatically inducing latent document-level graphs. This strategy allowing the model to incrementally aggregate relevant information for both local and global reasoning.

Transformer-based methods. This approach does not use any graph structures but instead adapts to the document-level relation extraction task by fine-tuning pre-trained models (Wang et al., 2019). Ye et al. (2020) introduced a copy-based training objective into the basic pre-trained language model, enabling the model to better capture coreference information. Tang et al. (2020) employed a hierarchical aggregation method to obtain reason-

ing information at different granularities at the document-level. Zhou et al. (2021) addressed the multi-label and multi-entity issues in relation extraction datasets through adaptive thresholds and local context pooling.

DOC-Evidence Retrieval (ER). Currently, a few studies have investigated the importance of evidence information in document-level relation extraction tasks. Yao et al. (2019) directly incorporated evidence sentence instances supporting entity relations into a new dataset. However, in the absence of evidence sentences, evidence information needs to be generated through an Evidence Retrieval (ER) task. Huang et al. (2021) employed heuristic rules to select informative sets of paths from the entire document to discover evidence sentences and further optimized relation extraction by combining BiLSTM. Ma et al. (2023) integrated evidence information into a Transformer-based DocRE system by directly guiding attention, without introducing any additional trainable parameters for the ER task. Compared to our work, they did not incorporate graph neural networks for end-to-end learning to derive the advantages of attention weights. In contrast, GEGA provides a more reliable allocation of weights for evidence information by constructing a fully connected graph and its corresponding fully connected matrix to learn structured information.

4 Methodology

This section elucidates the main framework of the proposed method, illustrated in Figure 2, the model can be segmented into tripartite: Input Encoder Layer, GEGA Module and Classification Layer.

4.1 Input Encoder Layer

This work adheres to the methodology employed in prior research to incorporate specialized markers $[CLS]$ and $[SEP]$ at the begin and end of a designated document $doc=[sent_N]_{N=1}^L$ for the purpose of outlining the document’s boundaries, where $sent=[t_n]_{n=1}^l$. Subsequently, input the document into the pre-trained BERT model. When the input length exceeds 512, the document is divided into two overlapping segments³. The first segment has a length of 512 and the second segment comprises the difference between the total input length and 512. Ultimately, the contextual embedding repre-

³The length of each data in the publicly available datasets is less than 1024.

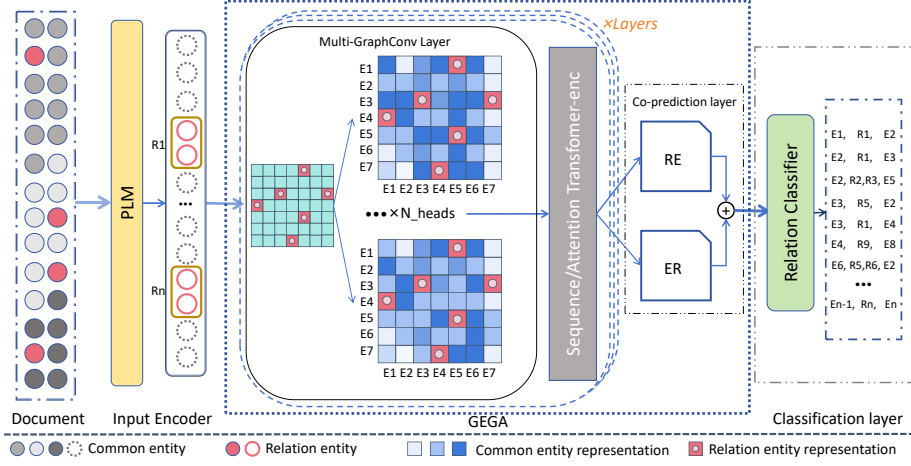


Figure 2: The overall architecture of our method. The gray circles with different depths belong to different sentences, and the color depth of the square is the basis to distinguish the attention weight score.

sentation H and attention matrix A of the token are derived:

$$(H, A) = [(h_1, a_1), (h_2, a_2), \dots, (h_{l \times L}, a_{l \times L})] \quad (1)$$

= BERT([doc])

where l is the length of the sentence (i.e., the number of tokens), and L is the length of the input document (i.e., the number of sentences).

For each entity, the efficacy of the max pooling function is pronounced when the inter-entity relations are explicitly articulated. Nevertheless, in the context of this study, the relations among entities remain ambiguous. It is understood that an entity may be referenced by one or several mentions, and a mention might uniquely identify an entity or fail to ascertain a definite corresponding entity. This necessitates the calculation of an entity’s embedding based on the embeddings of each associated mention. Following methodology of Jia et al. (2019), a soft version of the max function *LogSumExp* (*LSE*), is utilized to compute the embeddings of entities while concurrently capturing the attention distribution:

$$e_{\text{emb}} = LSE(x_1, \dots, x_n) = \log \sum_{i=1}^{|\mathcal{M}_e|} \exp(h_{i \in e}) \quad (2)$$

$$e_{\text{att}} = \frac{1}{|\mathcal{M}_e|} \sum_{i \in e} a_i \quad (3)$$

where e is an entity comprising multiple mentions, while a is the distribution of attention for mention i within entity e . $|\mathcal{M}_e|$ is the number of mentions for entity e .

4.2 GEGA Module

The GEGA module comprehends four parts: the Attention Concentration Layer, the Multi-GraphConv Layer, the Transformer-enc Layer, and the Collaborative Prediction Layer.

4.2.1 Attention Concentration Layer

We employ Attention Concentration Layer to transform the initial dependency tree into a fully connected weighted graph based on the dependency relations within the sentence. This approach can be construed as a soft pruning strategy (Xu et al., 2015) juxtaposed with the conventional hard pruning strategy (Guo et al., 2019). By assigning weights to the sequence data, as opposed to outright deletion, a greater amount of contextual information can be preserved, thereby fostering enhancements in module efficacy. Subsequently, we utilize the multi-head attention mechanism, wherein the input vector is mapped to several heads using a linear transformation layer to produce an adjacency matrix with varied weight distributions, denoted as $\tilde{A}^{(\text{head}_i)} = \text{Attention}(QW_i^Q, KW_i^K)$. We employ parallel computing to expedite the computational process.

$$\tilde{A}^{(\text{head}_i)} = \text{softmax} \left(\frac{QW_i^Q \times (KW_i^K)^T}{\sqrt{d}} \right) \quad (4)$$

where $Q, K \in \mathbb{R}^{N \times d_{\text{model}}}$, N is the length of the sequence, d_{model} is the dimensionality of the input feature, and $W_i^Q, W_i^K \in \mathbb{R}^{d_{\text{model}} \times d_k}$ is the weight parameter associated with the linear transformation.

4.2.2 Multi-GraphConv Layer

The Graph Convolutional Networks (GCNs (Kipf and Welling, 2016)) is a deep learning framework tailored for the processing of graph-structured data. It is a semi-supervised learning method based on graph structure that aggregates and propagates

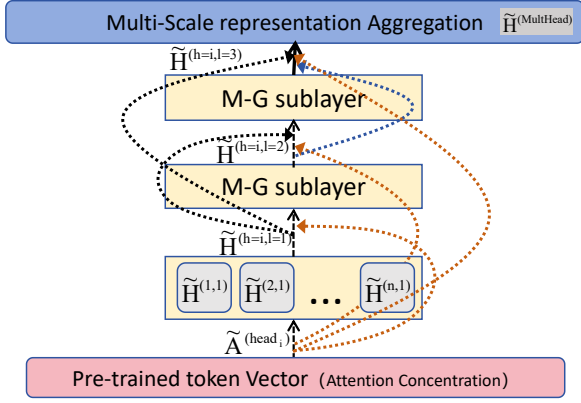


Figure 3: The overall architecture diagram of Multi-GraphConv (M-G) Layer includes three sub layers, each containing n heads.

node features in the graph structure, thus deriving node representations. The Multi-GraphConv (M-G) Layer is a densely connected graph structure data processing module that is constructed based on GCNs. Illustrated in Figure 3, the n -th sublayer’s output within the Multi-GraphConv (M-G) Layer serves as the subsequent $l - n$ sublayers’ input, with the N -th layer receiving an aggregation of all output features from the initial $n - 1$ layers.

Initially, the linear transformation of the adjacency matrix $\tilde{A}^{(\text{head}_i)}$ on the input features is computed for each head i . Subsequently, the impact of neighboring nodes’ features on the present node is determined for each layer l , and the current node’s features are consolidated with the previous layer’s output:

$$\tilde{H}^{(i,l)} = \text{ReLU} \left((\tilde{A}^{(\text{head}_i)})^{(i)} W^{(i,l)} \right) + \tilde{H}^{(i,l-1)} \quad (5)$$

where, L is the quantity of layers in the graph convolution layer, $W^{(i,l)}$ is the weight parameter of the i -th head in the l -th layer.

The feature representation resulting from the output of each head is combined to form the final output of this layer:

$$\tilde{H}^{(\text{MultiHead})} = \text{Concat} \left(\tilde{H}^{(1,l)}, \dots, \tilde{H}^{(i,l)} \right) W^O \quad (6)$$

where $W^O \in \mathbb{R}^{hd \times d_{\text{model}}}$ is the weight parameter of the linear transformation applied to the ultimate output.

4.2.3 Transformer-enc Layer

The Transformer-enc layer is composed of multiple encoder layers stacked together. These encoder

layers bear resemblance to the encoder layers delineated in the transformer model introduced by Vaswani et al. (2017). However, distinctively, our approach involves solely utilizing the output generated by the final three layers of the Encoder for the purpose of averaging. Each encoder layer incorporates self-attention mechanism and Feedforward Neural Network (FFN). This module facilitates the derivation of hidden representations of entities and an attention distribution matrix that are used as input for subsequent layers. The calculations can be outlined as follows:

$$\text{self-Att}(\tilde{H}_Q, \tilde{H}_K, \tilde{H}_V) = \text{softmax} \left(\frac{\tilde{H}_Q \tilde{H}_K^T}{\sqrt{d_k}} \right) \tilde{H}_V \quad (7)$$

$$L_a = \text{LayerNorm}(\tilde{H} + \text{self-Att}(\tilde{H})) \quad (8)$$

$$(\tilde{H}, \tilde{A}) = L_a + \text{FFN}(L_a) \quad (9)$$

Where $\tilde{H}_Q, \tilde{H}_K, \tilde{H}_V$ is the query, key, and value representations obtained from the linear transformation of \tilde{H} , d_k is the dimension of the attention head. LayerNorm is the layer normalization operation.

4.2.4 Collaborative Prediction Layer

The local context extraction methodology, as described by Zhou et al. (2021), is employed to ascertain the importance of individual tokens in relation to the entity pair (Es, Eo) , which is interpreted as the sentence-level importance. Erecting on this base, document-level importance was deduced by apportioning diverse attention weights in accordance to the contribution of each sentence within the document to the prediction of entity relations, and by establishing a fixed threshold. Sentences that exceed this threshold are selected as evidence sentences. The sentence-level importance $q_i^{(Es, Eo)}$ and document-level importance $p_j^{(Es, Eo)}$ can be computed as follows:

$$q_i^{(Es, Eo)} = \sum_{i=1}^{\tilde{H}} \tilde{A}_{Es} \cdot \tilde{A}_{Eo} \quad (10)$$

$$p_j^{(Es, Eo)} = \sum_{j=1}^l q_i^{(Es, Eo)} \quad (11)$$

We apportion more attention to evidence sentences and less to non-evidence sentences through evidence supervision to further coordinate the prediction results of document-level relation extraction. As depicted in Figure 4, we train a teacher model on the Human-Annotated Data (which encompasses relation labels and evidence sentences) of DocRED (Step 1). We utilize the trained teacher model to predict the entity relations and evidence

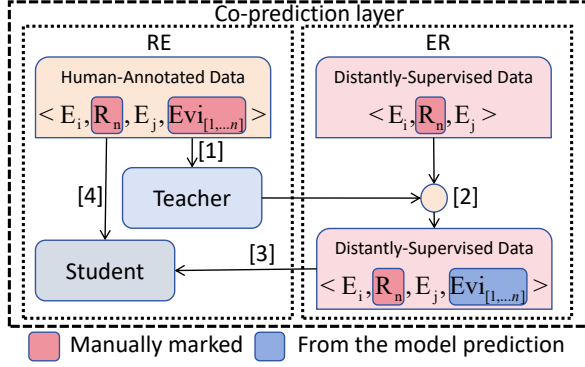


Figure 4: Step diagram of Co-prediction for RE and ER.

sentence distribution in Distantly-Supervised Data (which encompasses relation labels but lacks evidence sentences) (Step 2). Subsequently, we train a student model on the Distantly-Supervised Data that incorporates evidence sentences (Step 3), and retrain the student model using Human-Annotated Data (Step 4). Additionally, we define a row vector $\mathbf{z}^{(Es, Eo)}$ consisting only of 0s and 1s, generated based on Human-Annotated Data. This vector indicates whether each sentence is an evidence sentence for the relation triples: 1 if it is, and 0 if it is not.

$$\mathbf{z}^{(Es, Eo)} = \sum_{sent=1}^L \mathbf{z}^{(Es, Eo)} / \mathbf{1}^\top \sum_{sent=1}^L \mathbf{z}^{(Es, Eo)} \quad (12)$$

where $\mathbf{1}$ is the row vector composed of all 1, L is the total count of sentences contained within the document.

4.3 Classification Layer

We begin with the computation of a weighted average of entity importance at the sentence-level $\mathbf{q}_i^{(Es, Eo)}$, and subsequently cascading it with the previous entity representation. Following this, we apply the *tanh* activation function to normalize the input to range between $(-1, +1)$, resulting in the contextual representation of the two associated entities. The computation is elaborated as:

$$\begin{aligned} \mathbf{c}_{Es} &= \tanh \left(\mathbf{W}_{Es} \left[\mathbf{h}_{Es}; \tilde{\mathbf{H}}^\top \mathbf{q}_i^{(Es, Eo)} \right] + \mathbf{b}_{Es} \right) \\ \mathbf{c}_{Eo} &= \tanh \left(\mathbf{W}_{Eo} \left[\mathbf{h}_{Eo}; \tilde{\mathbf{H}}^\top \mathbf{q}_i^{(Es, Eo)} \right] + \mathbf{b}_{Eo} \right), \end{aligned} \quad (13)$$

where $\mathbf{W}_{Es}, \mathbf{W}_{Eo} \in \mathbb{R}^{d \times 2d}$, $\mathbf{b}_{Es}, \mathbf{b}_{Eo} \in \mathbb{R}^d$.

Finally, apply the grouped bilinear classifier proposed by Zheng et al. (2019) to calculate the relation category scores. $\text{Score}^{(Es, Eo)} = \mathbf{c}_{Es}^\top \mathbf{W}_{Rn} \mathbf{c}_{Eo} + \mathbf{b}_{Rn}$. The possibility of the entity pair (Es, Eo) possessing a relation Rn is computed thusly: $\mathbf{P}(Rn | Es, Eo) = \text{Sigmoid}(\text{Score}^{(Es, Eo)})$.

5 Experiments

5.1 Dataset and Evaluation

DocRED⁴ (Yao et al., 2019) is a benchmark dataset for document-level relation extraction tasks, released by Tsinghua University. DocRED comprises numerous documents from Wikipedia and Wikidata, each annotated with entities, relations between entities, and evidence sentences that support relation triples. It serves as the predominant benchmark for DocRE model training and evaluation.

Re-DocRED⁵ (Tan et al., 2022b) and Revisit-DocRED⁶ (Huang et al., 2022) are modified datasets of DocRED. They supplement a large number of relation triples to solve the problems of incomplete annotations, coreferential errors, and inconsistent logic in docred. Annotation quality has high accuracy and consistency, which provides a more reliable benchmark for DocRE-model training and evaluation.

We assess GEGA using an Nvidia Tesla V100 16GB GPU and evaluate it with F1, Ign-F1, and Evi-F1 metrics. Ign-F1 represents the calculated F1 score attained by excluding relational facts present in both the training and development/testing datasets. Evi-F1 serves as a significant measure for assessing the performance of ER and constitutes a new benchmark for assessing the quality of relation extraction models.

5.2 Single and Fusion

In the task of RE, the most ideal scenario is that the evidence sentence set of the dataset already contains all contextual information necessary to predict entity relations, thereby enabling accurate relation prediction results based solely on the evidence sentence set. However, manually labeled data and distant supervision data often fall short in this regard. Therefore, it is necessary to extract contextual information from the entire document to predict entity relations. We divide the above problem into two evaluation methods: (1) Single: extract entity relations from the entire document and obtain the corresponding prediction scores; (2) Fusion: predict entity relations based on a collection of evidence sentences and combine the prediction results with those from the Single method. This is similar to the Fusion of Evidence approach in Xie et al. (2022).

⁴<https://github.com/thunlp/DocRED>

⁵<https://github.com/tonytan48/Re-DocRED>

⁶<https://github.com/AndrewZhe/Revisit-DocRED>

Category	Model (With BERT _{base})	Dev			Test		
		Ign-F1	F1	Evi-F1	Ign-F1	F1	Evi-F1
• without Distant Supervision							
Sequence	CNN (Yao et al., 2019)	41.58	43.45	-	40.33	42.26	-
	BiLSTM (Yao et al., 2019)	48.87	50.94	-	48.78	51.06	-
	GAIN (Zeng et al., 2020)	59.14	61.22	-	59.00	61.24	-
Graph	MRN (Li et al., 2021)	59.74	61.61	-	59.52	61.74	-
	DocuNet (Zhang et al., 2021)	59.86	61.83	-	59.93	61.86	-
	GTN (Zhang et al., 2023a)	60.86	62.73	-	60.77	62.75	-
	SD-DocRE (Zhang et al., 2023b)	60.85	62.81	-	60.91	62.85	-
	ATLOP (Zhou et al., 2021)	59.22	61.09	-	59.31	61.30	-
Transformer	EIDER (Xie et al., 2022)	60.51	62.48	50.71	60.42	62.47	51.27
	SAIS (Xiao et al., 2022)	59.98	62.96	53.70	60.96	62.77	52.88
	DREEAM (Ma et al., 2023)	60.51	62.55	52.08	60.03	62.49	51.71
Teacher	GEGA-single (Ours)	59.98 \pm 0.12	61.95 \pm 0.12	52.19 \pm 0.15	59.31	61.52	51.90
	GEGA-fusion (Ours)	60.55 \pm 0.08	62.65 \pm 0.08	-	60.11	62.53	-
• with Distant Supervision							
Graph	AA (Lu et al., 2023)	61.31	63.38	-	60.84	63.10	-
Transformer	KD-DocRE (Tan et al., 2022a)	62.62	64.81	-	62.56	64.76	-
	DREEAM (Ma et al., 2023)	63.92	65.83	55.68	63.73	65.87	55.43
Student	GEGA-single (Ours)	64.02 \pm 0.15	65.83 \pm 0.15	56.09 \pm 0.18	63.82	65.85	55.89
	GEGA-fusion (Ours)	64.26 \pm 0.13	66.38 \pm 0.13	-	63.90	66.31	-

Table 1: Experimental results (%) for the dev and test set of DocRED. Using BERT-base as a pre-trained language model. The best score has been displayed in bold. The scores of other models refer to their respective papers.

5.3 Compared Methods

To ensure a fair comparison of the performance of DocRE baselines, we compare our model with three state-of-the-art methods, all using BERT-base as the pre-trained language model (PLM), which are: (1) Sequence-based methods: CNN (Yao et al., 2019), LSTM (Yao et al., 2019), BiLSTM (Yao et al., 2019). (2) Graph-based methods: GAIN (Zeng et al., 2020), MRN (Li et al., 2021), DocuNet (Zhang et al., 2021), GTN (Zhang et al., 2023a), SD-DocRE (Zhang et al., 2023b), AA (Lu et al., 2023). (3) Transformer-based methods: ATLOP (Zhou et al., 2021), EIDER (Xie et al., 2022), SAIS (Xiao et al., 2022), PRiSM (Choi et al., 2023), DREEAM (Ma et al., 2023). On this foundation, we categorize the above methods into two major classes: without Distant Supervision and with Distant Supervision.

6 Results and Analyses

We test the trained student and teacher models in both the Single and Fusion stages, and report the results on DocRED, Re-DocRED and Revisit-DocRED, where the results on Revisit-DocRED are moved to appendix.

6.1 Results on DocRED

Table 1 indicates that GEGA achieves superior Ign-F1 and F1 metrics compared to the established DocRE Baselines on both the development set and the test set. The single stage of the student model

has achieved performance levels comparable to the leading DREEAM (Ma et al., 2023). Notably, the fusion stage of the student model achieved the highest recorded scores, surpassing DREEAM by 0.34% (Ign F1) and 0.55% (F1) on the development set, as well as by 0.17% (Ign F1) and 0.44% (F1) on the test set.

GEGA also performs well in the test of the new benchmark Evi-F1, surpassing the previous most advanced DREEAM by 0.41% (Evi-F1) and 0.46% (Evi-F1) on the development set and test set respectively. In the table, it is evident that the Transformer-based and Graph-based models outperform the Sequence-based ones, validating the rationality behind integrating infographics with the Transformer.

6.2 Results on Re-DocRED

Table 2 presents the feedback outcomes of GEGA on the development and test sets of RE-DocRED, demonstrating that our GEGA achieves state-of-the-art results compared to other methods utilizing BERT-base as a pre-trained language model. Notably, GEGA has outperformed all other methods in the table during the fusion stage without Distant Supervision (Teacher). GEGA secured the highest Ign F1 and F1 scores in the fusion stage with Distant Supervision (Student), with improvements of 2.08% (Ign F1) and 2.52% (F1) respectively on the development set over the second-place GTN-BERT (Zhang et al., 2023a), and by 1.58% (Ign F1) and 1.97% (F1) respectively on the test set.

Category	Model (With BERT _{base})	Dev		Test	
		Ign-F1	F1	Ign-F1	F1
Graph	GAIN-BERT (Zeng et al., 2020)	71.99	73.49	71.88	73.44
	DocuNet-BERT (Zhang et al., 2021)	73.68	74.65	73.60	74.49
	GTN-BERT (Zhang et al., 2023a)	75.03	75.85	74.85	75.77
Transformer	ATLOP-BERT (Zhou et al., 2021)	73.35	74.22	73.22	74.02
	KMGRE-BERT (Jiang et al., 2022)	73.33	74.44	73.39	74.46
	KD-DocRE-BERT (Tan et al., 2022a)	73.76	74.69	73.67	74.55
	PRiSM-BERT (Choi et al., 2023)	72.92	74.25	72.35	73.69
	GEGA-single (Ours)	73.69 \pm 0.06	74.53 \pm 0.06	73.42 \pm 0.05	74.21 \pm 0.03
Teacher	GEGA-fusion (Ours)	76.06 \pm 0.07	77.41 \pm 0.07	75.28 \pm 0.03	76.61 \pm 0.03
Student	GEGA-single (Ours)	75.72 \pm 0.06	76.64 \pm 0.06	75.25 \pm 0.05	76.21 \pm 0.05
	GEGA-fusion (Ours)	77.11 \pm 0.08	78.37 \pm 0.08	76.43 \pm 0.06	77.74 \pm 0.07

Table 2: Performance (%) on the dev/test set of Re-DocRED. We use the same presentation method as Table 1. Other model results are replicated from the academic paper (Zhang et al., 2023a).

Model (With BERT _{base})	DocRED-Dev		
	Ign-F1	F1	Evi-F1
• GNNs layer ablation			
GEGA-single	64.02	65.83	56.09
— Attention Concentration Layer	63.37	65.34	55.69
— Multi-GraphConv Layer	62.94	64.31	55.19
— Transformer-enc Layer	63.87	65.66	55.93
• Training phase ablation			
GEGA-single	64.02	65.83	56.09
— self-training	62.19	63.45	53.98
— fine-tuning	63.91	65.80	55.63
— Distant Supervision-training	59.98	61.95	52.19

Table 3: Ablation analysis on DocRED-Dev.

6.3 Effect Analysis of GCNs and ER

Based on the test scores on DocRED and Re-DocRED, we observe that graph-based models such as DocuNET (Zhang et al., 2021) and GTN-BERT (Zhang et al., 2023a) have achieved superior performance in the field. Graphs have an advantage in conveying document-level contextual information, so we used grid search to select a 2-layer GNNs to guide multiple attention maps. Additionally, DREEAM (Ma et al., 2023) is the first method to enhance relation extraction performance purely through evidence-guided attention, it has already achieved excellent scores on the DocRED set. By integrating GCNs with evidence retrieval, we further improved its scores by 0.41% (Evi-F1) and 0.46% (Evi-F1) on the dev and test sets, respectively. Therefore, we conclude that GCNs and ER significantly enhance performance in the relation extraction task.

6.4 Ablation Studies

We conducted ablation experiments were conducted on the development set to analyze the GNNs layer and Training phase of GEGA. The single phase of GEGA (Student) was utilized as the test

benchmark. The results of the score post-ablation of each part are presented in Table 3. Initially, we removed the Attention Concentration Layer, which resulted in a minor decline in performance. Subsequently, upon removing the Multi-GraphConv Layer, a significant performance deterioration was observed, implying the importance of constructing multiple attention distribution graphs for relation extraction. Upon removal of the Transformer-enc Layer, we noted a relatively minor decline in performance. We speculate that this may be related to using Transformer based BERT as PLM.

Additionally, during the training process, we performed ablation on the self-training, fine-tuning, and Distant Supervision-training stages, in order to further analyze their impact. The results indicate that when the ER self-training phase is omitted, performance declines, whereas the absence of the fine-tuning stage did not lead to a noticeable decline in performance. Further more, omitting the Distant Supervision-training stage caused severe performance degradation. These findings highlight the effectiveness of our ER method in enhancing relation extraction.

7 Conclusion

We propose GEGA, the first model to employ GCNs and ER jointly guided attention to enhance DocRE. We validate the superiority of our model on three widely used datasets: DocRED, Re-DocRED and Revisit-DocRED. GEGA is trained using parallel computing in both fully supervised and semi-supervised settings, without incurring additional overhead, making it convenient for use in the era of Large Language Models (LLMs). In the future, we aim to leverage the scalability of GEGA and apply it to a broader range of scenarios, including entity recognition, event extraction, and more.

8 Limitations

The model GEGA is subject to two limitations. Firstly, when utilizing Multi-GraphConv Layers to induce multiple fully connected attention distribution matrices, there is a possibility of generating one matrix that differs significantly from others in terms of weight distribution. This could lead to significant deviations in prediction results. We hypothesize that guiding the construction of multiple fully connected attention matrices using evidence information may reduce the occurrence of such undesirable situations, a conjecture that will be verified in future work. Secondly, it is acknowledged that the relations between most entity pairs can be predicted based on the local context of the entities. However, our model utilizes evidence sentences retrieved from the entire document corpus, which are strongly correlated with the entity pairs of interest, rather than evidence sentences obtained specifically for individual relation triples. This approach may result in the model carrying more global contextual information while reducing the utilization of local context information.

9 Ethics Statement

Our proposed GEGA demonstrates outstanding scalability and applicability, serving as an excellent solution for both DocRE and DocER tasks. This method is evaluated solely on publicly available datasets, ensuring no compromise on individual privacy. Furthermore, we provide the source code implementation of GEGA to enable researchers to reproduce its performance authentically, fostering academic exchange in the field of DocRE.

References

Rui Cai, Xiaodong Zhang, and Houfeng Wang. 2016. [Bidirectional recurrent convolutional neural network for relation classification](#). In *In Proceedings of the ACL*, pages 756–765.

Kyunghyun Cho, Bart van Merriënboer, Çağlar Gulçehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. [Learning phrase representations using rnn encoder–decoder for statistical machine translation](#). In *In Proceedings of the EMNLP*, pages 1724–1734.

Minseok Choi, Hyesu Lim, and Jaegul Choo. 2023. [Prism: Enhancing low-resource document-level relation extraction with relation-aware score calibration](#). In *In Findings of the ACL-IJCNLP*, pages 39–47.

Fenia Christopoulou, Makoto Miwa, and Sophia Ananiadou. 2019. [Connecting the dots: Document-level neural relation extraction with edge-oriented graphs](#). In *In Proceedings of the EMNLP-IJCNLP*, pages 4925–4936.

Zhijiang Guo, Yan Zhang, and Wei Lu. 2019. [Attention guided graph convolutional networks for relation extraction](#). In *In Proceedings of the ACL*, pages 241–251.

Quzhe Huang, Shibo Hao, Yuan Ye, Shengqi Zhu, Yansong Feng, and Dongyan Zhao. 2022. [Does recommend-revise produce reliable annotations? an analysis on missing instances in docred](#). In *In Proceedings of ACL*, pages 6241–6252.

Quzhe Huang, Shengqi Zhu, Yansong Feng, Yuan Ye, Yuxuan Lai, and Dongyan Zhao. 2021. [Three sentences are all you need: Local path enhanced document relation extraction](#). In *In Proceedings of the ACL-IJCNLP*, pages 998–1004.

Robin Jia, Cliff Wong, and Hoifung Poon. 2019. [Document-level n-ary relation extraction with multiscale representation learning](#). In *In Proceedings of the NAACL-HLT*, pages 3693–3704.

Feng Jiang, Jianwei Niu, Shasha Mo, and Shengda Fan. 2022. [Key mention pairs guided document-level relation extraction](#). In *In Proceedings of the COLING*, pages 1904–1914.

Thomas N Kipf and Max Welling. 2016. [Semi-supervised classification with graph convolutional networks](#). In *In Proceedings of the ICLR*.

Jingye Li, Kang Xu, Fei Li, Hao Fei, Yafeng Ren, and Donghong Ji. 2021. [Mrn: A locally and globally mention-based reasoning network for document-level relation extraction](#). In *In Findings of the ACL-IJCNLP*, pages 1359–1370.

Chonggang Lu, Richong Zhang, Kai Sun, Jaemin Kim, Cunwang Zhang, and Yongyi Mao. 2023. [Anaphor assisted document-level relation extraction](#). In *In Proceedings of the EMNLP*, pages 15453–15464.

Youmi Ma, An Wang, and Naoaki Okazaki. 2023. [Dreeam: Guiding attention with evidence for improving document-level relation extraction](#). *arXiv preprint arXiv:2302.08675*.

Angrosh Mandya, Danushka Bollegala, and Frans Coenen. 2020. [Graph convolution over multiple dependency sub-graphs for relation extraction](#). In *In Proceedings of the COLING*, pages 6424–6435.

Guoshun Nan, Zhijiang Guo, Ivan Sekulić, and Wei Lu. 2020. [Reasoning with latent structure refinement for document-level relation extraction](#). In *In Proceedings of the ACL*, pages 1546–1557.

Nanyun Peng, Hoifung Poon, Chris Quirk, Kristina Toutanova, and Wen-tau Yih. 2017. [Cross-sentence n-ary relation extraction with graph lstms](#). In *Proceedings of the TACL*, 5:101–115.

708	Sunil Kumar Sahu, Fenia Christopoulou, Makoto Miwa, and Sophia Ananiadou. 2019. Inter-sentence relation extraction with document-level graph convolutional neural network . In <i>In Proceedings of the ACL</i> , pages 4309–4316.	763
709		764
710		765
711		766
712		767
713	Daniil Sorokin and Iryna Gurevych. 2017. Context-aware representations for knowledge base relation extraction . In <i>In Proceedings of the EMNLP</i> , pages 1784–1789.	768
714		769
715		770
716		771
717	Qingyu Tan, Ruidan He, Lidong Bing, and Hwee Tou Ng. 2022a. Document-level relation extraction with adaptive focal loss and knowledge distillation . In <i>In Findings of the ACL</i> , pages 1672–1681.	772
718		773
719		774
720		775
721	Qingyu Tan, Lu Xu, Lidong Bing, Hwee Tou Ng, and Sharifah Mahani Aljunied. 2022b. Revisiting docred-addressing the false negative problem in relation extraction . In <i>In Proceedings of the EMNLP</i> , pages 8472–8487.	776
722		777
723		778
724		779
725		
726	Hengzhu Tang, Yanan Cao, Zhenyu Zhang, Jiangxia Cao, Fang Fang, Shi Wang, and Pengfei Yin. 2020. Hin: Hierarchical inference network for document-level relation extraction . In <i>In Proceedings of the PAKDD</i> , pages 197–209. Springer.	780
727		781
728		782
729		783
730		
731	Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need . In <i>Proceedings of the NeurIPS</i> , 30.	784
732		785
733		786
734		787
735	Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. Graph attention networks . In <i>In Proceedings of the ICLR</i> .	788
736		
737		
738		
739	Patrick Verga, Emma Strubell, and Andrew McCallum. 2018. Simultaneously self-attending to all mentions for full-abstract biological relation extraction . In <i>In Proceedings of the ACL-HLT</i> , pages 872–884.	789
740		790
741		791
742		792
743	Hong Wang, Christfried Focke, Rob Sylvester, Nilesh Mishra, and William Wang. 2019. Fine-tune bert for docred with two-step process . <i>arXiv preprint arXiv:1909.11898</i> .	793
744		794
745		
746		
747	Yuxin Xiao, Zecheng Zhang, Yuning Mao, Carl Yang, and Jiawei Han. 2022. Sais: Supervising and augmenting intermediate steps for document-level relation extraction . In <i>In Proceedings of the NAACL</i> , pages 2395–2409.	795
748		796
749		797
750		798
751		799
752	Yiqing Xie, Jiaming Shen, Sha Li, Yuning Mao, and Jiawei Han. 2022. Eider: Empowering document-level relation extraction with efficient evidence extraction and inference-stage fusion . In <i>In Findings of the ACL</i> , pages 257–268.	800
753		801
754		802
755		803
756		
757	Kelvin Xu, Jimmy Lei Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhutdinov, Richard S Zemel, and Yoshua Bengio. 2015. Show, attend and tell: neural image caption generation with visual attention . In <i>In Proceedings of the ICML</i> , pages 2048–2057.	804
758		805
759		806
760		807
761		
762		
	Yuan Yao, Deming Ye, Peng Li, Xu Han, Yankai Lin, Zhenghao Liu, Zhiyuan Liu, Lixin Huang, Jie Zhou, and Maosong Sun. 2019. Docred: A large-scale document-level relation extraction dataset . In <i>In Proceedings of the ACL</i> , pages 764–777.	808
		809
		810
		811
		812
	Deming Ye, Yankai Lin, Jiaju Du, Zhenghao Liu, Peng Li, Maosong Sun, and Zhiyuan Liu. 2020. Coreferential reasoning learning for language representation . In <i>In Proceedings of the EMNLP</i> , pages 7170–7186.	813
	Daojian Zeng, Kang Liu, Siwei Lai, Guangyou Zhou, and Jun Zhao. 2014. Relation classification via convolutional deep neural network . In <i>In Proceedings of the COLING</i> , pages 2335–2344.	814
		815
	Shuang Zeng, Yuting Wu, and Baobao Chang. 2021. Sire: Separate intra- and inter-sentential reasoning for document-level relation extraction . In <i>In Findings of the ACL-IJCNLP</i> , pages 524–534.	
	Shuang Zeng, Runxin Xu, Baobao Chang, and Lei Li. 2020. Double graph based reasoning for document-level relation extraction . In <i>In Proceedings of the EMNLP</i> , pages 1630–1640.	
	Liang Zhang, Zijun Min, Jinsong Su, Pei Yu, Ante Wang, and Yidong Chen. 2023a. Exploring effective inter-encoder semantic interaction for document-level relation extraction . In <i>In Proceedings of the IJCAI</i> , pages 5278–5286.	
	Liang Zhang, Jinsong Su, Zijun Min, Zhongjian Miao, Qingguo Hu, Biao Fu, Xiaodong Shi, and Yidong Chen. 2023b. Exploring self-distillation based relational reasoning training for document-level relation extraction . In <i>In Proceedings of the AAAI</i> , pages 13967–13975.	
	Ningyu Zhang, Xiang Chen, Xin Xie, Shumin Deng, Chuanqi Tan, Moshua Chen, Fei Huang, Luo Si, and Huajun Chen. 2021. Document-level relation extraction as semantic segmentation . In <i>In Proceedings of the IJCAI</i> , page 3999–4006.	
	Yuhao Zhang, Peng Qi, and Christopher D Manning. 2018. Graph convolution over pruned dependency trees improves relation extraction . In <i>In Proceedings of the EMNLP</i> , pages 2205–2215.	
	Heliang Zheng, Jianlong Fu, Zheng-Jun Zha, and Jiebo Luo. 2019. Learning deep bilinear transformation for fine-grained image representation . In <i>Proceedings of the NeurIPS</i> , 32.	
	Wenxuan Zhou, Kevin Huang, Tengyu Ma, and Jing Huang. 2021. Document-level relation extraction with adaptive thresholding and localized context pooling . In <i>In Proceedings of the AAAI</i> , pages 14612–14620.	
	A Loss	
	To accommodate the requirements of the RE and ER tasks within both the teacher model and the	




Figure 5: Loss value variation of the GEGA model trained on the DocRED dataset

student model, distinct forms of loss function computation have been devised for the relation classification approach. The loss variation is shown in Figure 5.

A.1 RE loss:

For the RE tasks of the two models previously described, we implement the Adaptive Thresholding Loss (ATL) as proposed by ATLOP. During the training phase, we use a threshold class (TH) to learn a threshold such that the logits of the positive class $\mathcal{R}\mathcal{P}$ exceed it, and the logits of the negative class $\mathcal{R}\mathcal{N}$ fall below it.

$$\mathcal{L}_{\text{RE}} = - \sum_{Rn \in \mathcal{R}\mathcal{P}} \frac{\exp(\text{Score}_{Rn}^{(Es, Eo)})}{\sum_{Rn' \in \mathcal{R}\mathcal{P} \cup \{\text{TH}\}} \exp(\text{Score}_{Rn'}^{(Es, Eo)})} - \frac{\exp(\text{Score}_{\text{TH}}^{(Es, Eo)})}{\sum_{Rn' \in \mathcal{R}\mathcal{N} \cup \{\text{TH}\}} \exp(\text{Score}_{Rn'}^{(Es, Eo)})} \quad (14)$$

A.2 ER loss:

The tasks targeted by the teacher model and the student model differ in detail. The teacher model is trained on Human-Annotated Data, which includes reliable manually annotated evidence sentences, while the student model is trained on Distantly-Supervised Data that incorporates evidence sentences identified through the teacher model’s ER. Considering these distinctions, there is a requirement for specialized loss computation methods. Consequently, we propose both document-level and sentence-level loss calculations.

Document-level Loss: By integrating the document-level importance distribution $p_j^{(Es, Eo)}$

with the original manually annotated evidence sentences $z^{(Es, Eo)}$ from the dataset, we induce a localized z context representation that contributes significantly to RE. We use the Kullback-Leibler Divergence (KL divergence), a method for measuring the difference between two probability distributions within the same event space. In this paper, it is used to assess the degree of divergence between $p_j^{(Es, Eo)}$ and $z^{(Es, Eo)}$:

$$\begin{aligned} \mathcal{L}_{\text{ER}}^{\text{doc}} &= \text{KL}(z^{(Es, Eo)} \| p_j^{(Es, Eo)}) \\ &= \sum_{sent=1}^L z^{(Es, Eo)} \log \frac{z^{(Es, Eo)}}{p_j^{(Es, Eo)}} \end{aligned} \quad (15)$$

Sentence-level Loss: We use the teacher model trained on Human-Annotated Data to perform ER testing on Distantly-Supervised Data, predicting its sentence-level evidence distribution $\tilde{q}_i^{(Es, Eo)}$. Thereafter, utilize Kullback-Leibler (KL) divergence to compute the difference in sentence-level probability distributions between the teacher model and the student model.

$$\begin{aligned} \mathcal{L}_{\text{ER}}^{\text{sent}} &= \text{KL}(q_i^{(Es, Eo)} \| \tilde{q}_i^{(Es, Eo)}) \\ &= \sum_{t=1}^l q_i^{(Es, Eo)} \log \frac{q_i^{(Es, Eo)}}{\tilde{q}_i^{(Es, Eo)}} \end{aligned} \quad (16)$$

Finally, we apply the prevalent weighted summation technique for document-level and sentence-level losses to equilibrate the losses of RE and ER, where λ serving as a hyperparameter.

$$\mathcal{L} = (1 - \lambda)\mathcal{L}_{\text{RE}} + \lambda\mathcal{L}_{\text{ER}} \quad (17)$$

B Hyperparameter

We ran tests on the above three datasets using different random seeds five times and reported the average test accuracy. In the manuscript, a comprehensive Table 4 delineating the hyperparameters and configuration has been furnished, encapsulating pivotal settings utilized across the experimental trials. These hyperparameters, meticulously curated and fine-tuned, exert control over multifarious facets of the model’s dynamics throughout the training and evaluation phases.

C Results on Revisit-DocRED

Table 5 shows the test results of GEGA (Student) on Revisit-DocRED. According to the test method of Huang et al. (2022), we trained on the DocRED

Hyperparameter/Configuration	teacher	student	self-train	finetune	evaluation
train-file	annotated	-	distant	annotated	-
test-file	-	distant	-	-	test
dev-file	dev	-	dev	dev	-
num-class	97	97	97	97	97
gradient-accumulation-steps	1	-	2	1	-
train-batch-size	4	-	4	4	-
test-batch-size	8	4	8	8	8
num-labels	4	4	4	4	4
evi- λ	0.1	0.1	0.1	0.1	-
lr-transformer	5e-5	-	3e-5	1e-6	-
lr-added	-	-	-	3e-6	-
max-grad-norm	1.0	-	5.0	2.0	-
evi-thresh	0.2	0.2	0.2	0.2	0.2
warmup-ratio	0.06	-	0.06	0.06	-
num-train-epochs	30.0	-	2.0	10.0	-
eval-mode	-	-	-	-	single/fusion

Table 4: Hyperparameter/Configuration Settings for Training and Evaluation of GEGA.

Model (With BERT _{base})	Test	
	Ign-F1	F1
CNN-BERT* (Yao et al., 2019)	29.70	30.04
LSTM-BERT* (Yao et al., 2019)	31.32	31.77
BiLSTM-BERT* (Yao et al., 2019)	32.50	32.91
GAIN-BERT (Zeng et al., 2020)	41.27	41.64
ATLOP-BERT (Zhou et al., 2021)	41.62	41.90
KMGRE-BERT (Jiang et al., 2022)	42.78	43.16
KD-DocRE-BERT (Tan et al., 2022a)	43.22	43.68
GTN-BERT (Zhang et al., 2023a)	44.84	45.33
DREEAM-BERT* (Ma et al., 2023)	55.32	56.48
GEGA-single (Ours)	45.34 \pm 0.12	45.58 \pm 0.10
GEGA-fusion (Ours)	55.89 \pm 0.07	56.97 \pm 0.05

Table 5: Performance (%) on the test set of Revisit-DocRED. Results marked with * are obtained by our code reproduction. Other model results are replicated from the academic paper (Zhang et al., 2023a).

training set, and then tested on the test set provided by Revisit-DocRED. It can be seen that GEGA has a great performance improvement compared with other methods.

D Supplementary Model Analysis

Here, we conducted experimental analyses on the number of GNNs layers and the number of heads in the Multi-GraphConv Layer of GEGA. We selected the number of GNNs layers from $\{1, 2, 3, 4, 5\}$ and the number of heads from $\{1, 2, 3, 4\}$. The experimental results are shown in Figure 6. Ultimately, the analysis concluded that the performance is optimal when $layers=2$ and $heads=2$.

E Case Study

To further demonstrate the superior performance of GEGA, we extracted a case from the DocRED

dataset and used it to compare the performance of GEGA with two other advanced methods (SAIS and DREEAM).

E.1 RE

From Figure 7, we observe five types of relations among five entities. SAIS correctly identified four types of relations but overlooked the "country of citizenship" relation between "Robert F./Mark R." and "United States". Additionally, it incorrectly identified a "country of citizenship" relation between "Robert F./Mark R." and "American" as well as a "located in the administrative territorial entity" relation with "Terry McAuliffe". While DREEAM correctly identified all the existing relations, it excessively identified a "country" relation between "American" and "United States". In contrast, GEGA perfectly extracted the correct relational network in this case.

E.2 ER

During evidence retrieval, both GEGA and DREEAM labeled the evidence source for the relation between *United States* and *Virginia* as [S1, S2], missing [S0]. Additionally, DREEAM incorrectly labeled the evidence information for the relation between *Terry McAuliffe* and *United States* as [S0, S6]. SAIS incorrectly labeled the evidence information for the relation between *Virginia* and *Terry McAuliffe* as [S0, S1, S2, S6]. We believe that GEGA's superior performance is also attributed to its better ER performance.

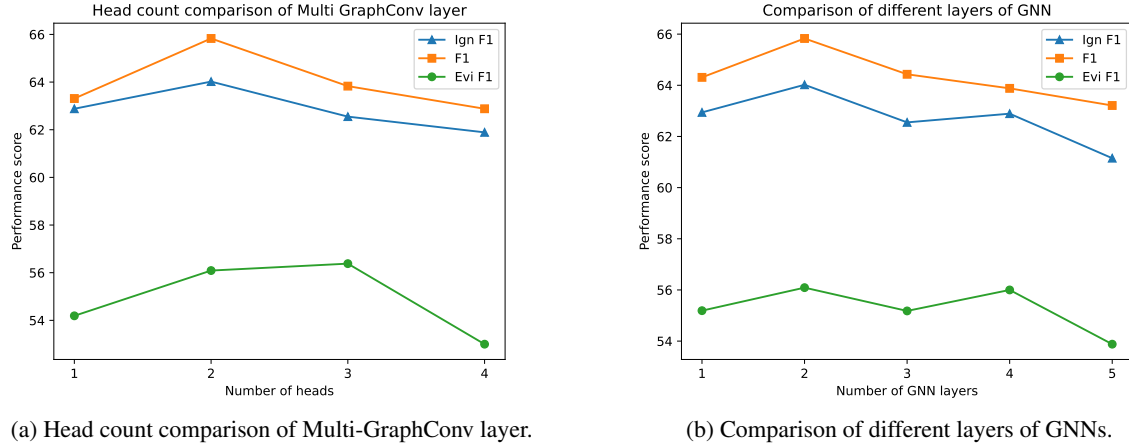


Figure 6: Comparison of experimental results with different parameter settings of GEGA.

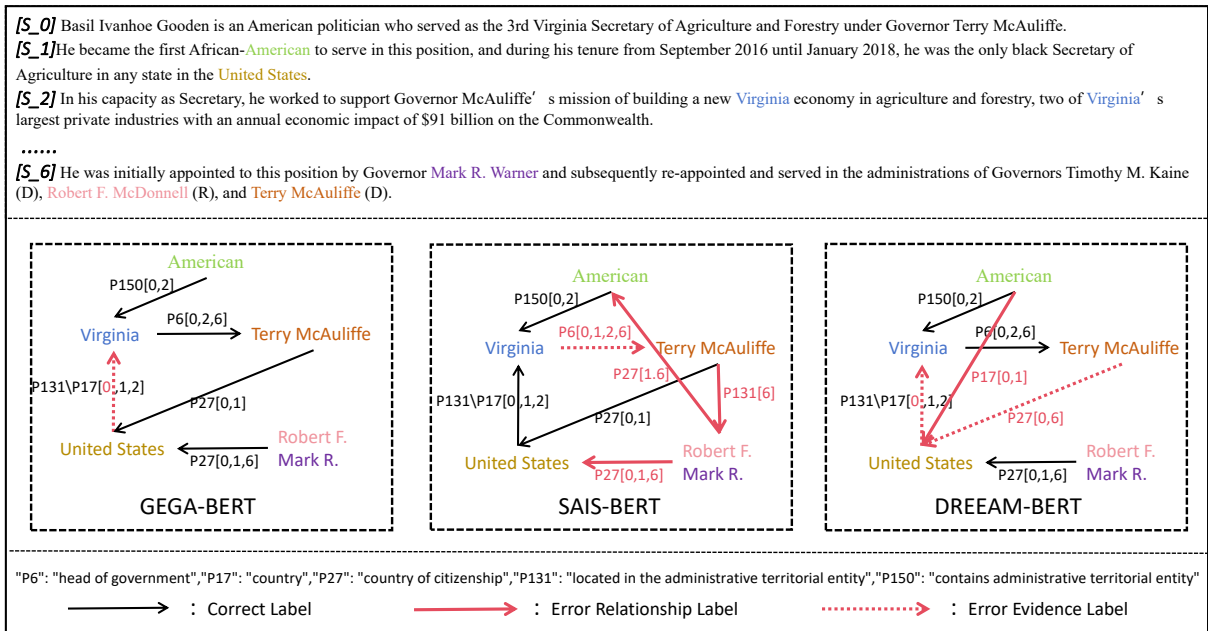


Figure 7: The comparison results of a case on three advanced models show that the entity is marked with special color, P_{number} is the relation label, and the red arrow is the prediction error.