

Knowledge Triplets Derivation from Scientific Publications via Dual-Graph Resonance

Kai Zhang¹, Pengcheng Li⁴, Kaisong Song², Yangyang Kang², Xurui Li²,
Xuhong Zhang³, Xiaozhong Liu^{1*}

¹Worcester Polytechnic Institute, Worcester USA

²Alibaba Group, Hangzhou China

³Indiana University, Bloomington USA

⁴Wuhan University Wuhan China

{kzhang8, xliu14}@wpi.edu, {kaisong.sks, xurui, yangyang.kangyy}@alibaba-inc.com
fantasticlpc@whu.edu.cn, zhangxuh@indiana.edu

Abstract

Scientific Information Extraction (SciIE) is a vital task and is increasingly being adopted in domain-specific (e.g., Biomedical) data mining to conceptualize and epitomize knowledge triplets from scientific literature. Existing relation extraction methods aim to extract explicit triplet knowledge from documents, however, they can hardly perceive unobserved factual relations. Recent generative methods have more flexibility, but their generated relations will encounter trustworthiness problems. In this paper, we propose a novel **Extraction-Contextualization-Derivation (ECD)** strategy to generate a document-specific and entity-expanded dynamic graph from a shared static knowledge graph. Then, we propose a novel **Dual-Graph Resonance Network (DGRN)** which can generate richer explicit and implicit relations under the guidance of static and dynamic knowledge topologies. Experiments conducted on a public PubMed corpus validate the superiority of our method against several state-of-the-art baselines.

Keywords: Graph Neural Network, Biomedical Triple Extraction.

1. Introduction

As scientific literature grows at an expeditious pace, it becomes increasingly labor-intensive for scholars to curate the massive information and consume their interested knowledge. For example, PubMed, as one of the most commonly used searching biomedical publication databases, contains more than 34 million publications¹. Even with a focused research interest, such as oncology, it is still very laborious to locate useful information from noisy retrieval results. To address this challenge, SciIE approaches can be employed to extract structural information from scientific articles, which has drawn great attentions from Natural Language Processing (NLP) community (Gupta and Manning, 2011; Viswanathan et al., 2021).

Recently, great efforts have been made regarding SciIE tasks and obtained substantial achievements. Lee et al. (2020) and Beltagy et al. (2019) trained pre-trained language models with biomedical corpus for representation learning and downstream task fine-tuning. Based on these models, many extractive methods are proposed to extract triples explicitly present in the scientific documents (Ye et al., 2020; Nayak and Ng, 2020). Despite some success, extractive methods can only obtain limited biomedical knowledge due to the limited searching space. To enable implicit biomed-

ical triple extraction, generative extraction methods provide more flexibility, however they can not guarantee correctness of the generated knowledge and may cause trustworthiness concerns (Zhang et al., 2020). For example in Figure 1, the generative method produces an incorrect triplet $\langle \text{skin}, \text{cancerous}, \text{throat cancer} \rangle$ because there exists no direct relation between entities “skin” and “throat cancer”, while “throat cancer” is generated based on the trigger word “cancerous”. Recent graph-based methods provide the possibility to produce trusted and unobserved biomedical triplets via multi-hop path-reasoning on graphs, such as $\langle \text{melanoma}, \text{belongs to}, \text{skin cancer} \rangle$ can be derived from $\langle \text{skin}, \text{cancerous}, \text{skin cancer} \rangle$ and $\langle \text{skin}, \text{cancerous}, \text{melanoma} \rangle$. However, such methods only alleviate the problem and their extracted triplets are restricted by limited document entities.

High-quality and large-scale biomedical knowledge graphs (BKG) have been studied extensively and constructed (Wang et al., 2021), which provide possibilities to expand limited document entities and enable the generative methods to generate more trusted and unobserved biomedical triplets². However, one may hesitate to adopt BKG directly due to the potential of massive noise. In this study, we propose a novel *Extraction-Contextualization-Derivation (ECD)* strategy to address this prob-

* Corresponding Author

¹<https://pubmed.ncbi.nlm.nih.gov/about/>

²We resort to the biomedical domain since there exists rich data resources in both text and graph format.

Biomedical document: The skin is ... and causes cancerous process, ... that aggravates ... skin cancer	
Method: Extractive method Triples: < skin , cancerous , skin cancer >	✓
Method: Generative method Triples: < skin , cancerous , skin cancer > < skin , cancerous , throat cancer >	✓ ✗
Method: Knowledge Graph (KG) method Triples: < melanoma , belongs to , skin cancer >	✓
Method: Generative method + KG Triples: < skin , cancerous , skin cancer > < skin , cancerous , melanoma >	✓ ✓

Figure 1: Comparison among different methods. Entities and relations observed in the document are denoted in red color, and unobserved ones are denoted in blue. Correct/wrong triplets are labeled as check/cross marks.

lem - using large-scale BKG as the *Static Graph* to encapsulate the biomedical domain knowledge and the derived knowledge sub-BKG as *Dynamic Graph* to characterize document-related knowledge. The interactions between static and dynamic graphs ensure the comprehensiveness and trustworthiness of the knowledge generation.

Prior ScIE studies rarely equip generative methods with domain knowledge in an end-to-end fashion, and such investigation can be especially critical for biomedical knowledge mining which poses several challenges. Firstly, the input document and knowledge graph are information complementary and should be fully interactively modeled in the encoding process. Secondly, multi-hop path reasoning on graphs should be utilized for guiding copy mechanism to provide trusted biomedical triplets. To this end, we propose a novel Dual-Graph Resonance Network with a fourfold contribution:

- We propose an “Extraction-Contextualization-Derivation (ECD)” strategy to derive a document-related dynamic graph from a shared static graph which can be used for implicit entity expansion.
- We propose a “Dual-Graph Resonance Network (DGRN)” frame to generate both observed and unobserved knowledge triplets by jointly modeling the input document and the dual graphs.
- A new dataset (Bio-Sci), derived from open-source biomedical domain corpus, which contains 32,330 publications and each of them equipped with implicit and explicit triplets is released for biomedical NLP research community.
- Extensive experiments conducted on Bio-Sci show an average 4.91% improvement on F1 score against the best SOTA method which validate the effectiveness and superiority of our method.

2. Related Work

Existing ScIE works extract or generate knowledge triplets from different parts of scientific publications, such as content (Luan et al., 2019; Augenstein et al., 2017), abstract, introduction and citation sentences³ (Nakov et al., 2004). The mainstream approaches can be concluded as follow.

Extractive models have been extensively studied. Wei et al. (2020); Beltagy et al. (2019) integrated pre-trained language models into an encoder-decoder framework for performance improvements. Nayak and Ng (2020); Zhao et al. (2021) introduced a joint learning framework to model connections between relations and their corresponding entity pairs. Moreover, Takanobu et al. (2019); Bai and Zhao (2018) utilized a hierarchical structure that featured connections among different content layers to find facts within the content summarization for relation extraction (Lu et al., 2022). Further, with the emergence of powerful large language models, zero-shot fashion can be leveraged for triple extraction via chatting with such models (Wei et al., 2023). However, these methods ignore fine-grained entity-level information interaction and integration.

Graph-based models can provide both dependencies among entities and path reasoning potential for inference. Zeng et al. (2020b) proposed a mention-to-entity graph aggregation model which can capture the relation of entities across sentences. Instead of integrating graph structure into neural network models, Peng et al. (2017); Guo et al. (2019); Xu et al. (2021); Huang et al. (2021b) enhanced the mention-to-entity graph paradigm by introducing multi-hop path reasoning and reconstructing the graph based on the obtained path information. Another fashion is to transform sentences into a graph and perform multi-view GCN to obtain the relation (Chen et al., 2022; Rabbani et al., 2023; Zhang et al., 2023). Unfortunately, existing graph-based models can not synthesize new domain knowledge.

Generative models are recently proposed to generate triplets flexibly from input documents. Zeng et al. (2018) proposed a CopyRE model to select entities or relations via copy mechanism. Zeng et al. (2020a); Ye et al. (2020) further improved this paradigm by introducing multi-task learning and contrastive learning frameworks. Other generative models utilized additional information. For example, Zhang et al. (2021b) proposed a Knowledge-Graph (KG)-enriched Abstract Meaning Representation (AMR) framework which uses external information to enrich the AMR graph extracted from scientific papers. Garg et al. (2021) leveraged transformers to refine semantic embed-

³called citance in the following sections

ding of a given text for better generation(Zhang et al., 2022; Giorgi et al., 2022). However, Zhang et al. (2020) claimed the factual correctness and trustworthiness problems of these methods which ignore prior knowledge to ensure that generated knowledge triplets are more reliable.

While these studies have accomplished notable advancements, our DGRN steps further by leveraging the synergistic power of generative and graph-based methods within a cohesive framework.

3. Dual-Graph Construction

Generative extraction methods may encounter trustworthiness issues. To address this, we propose enhancing these methods with guidance from biomedical knowledge. We first introduce a static background knowledge graph (BKG), which is built on public resources. We then generate a document-specific dynamic BKG using a dynamic graph generator, which enables the generation of unobserved biomedical triplets. The detail construction process is depicted in Figure 2.

3.1. Static Biomedical Knowledge Graph

Large-scale biomedical knowledge graphs are recently constructed, and some have been made publicly available. Wang et al. (2021) constructed a public knowledge graph with 1.47 million triplets and 96,397 entities from multiple sources, such as PubMed, DrugCentral etc. Thus, the knowledge graph owns massive biomedical knowledge, which will expand document entities greatly. Given any BKG as an external public resource, we formulate it as a shared static graph $G_S = \{(e_i, r_{i,j}, e_j) | e_i, e_j \in \mathcal{E}, r_{i,j} \in \mathcal{R}\}$, where \mathcal{E} and \mathcal{R} represent an entity set and a relation set respectively. However, direct application of G_S may introduce noisy information into generative models and pollute the SciIE results. Thus, we propose a dynamic biomedical knowledge graph construction strategy.

3.2. Dynamic Sub-Graph

Inspired by Viswanathan et al. (2021); Neumann et al. (2019), we let $D = \{S_{abs}, S_{int}, S_{cit}\}$ be a biomedical document, which consists of three important sections, i.e., *abstract*, *introduction* and *citance*. Each section $S_* = \{w_i\}_{i=1}^{N_*}$ refers to a sequence of words of length N_* . We aim to derive a dynamic graph automatically by enriching consisted document relations under the prior guidance of the static graph G_S . Note that we consider the three sections together when constructing the dynamic graph. Particularly, we propose an **Extraction-Contextualization-Derivation (ECD)** framework which consists of three key steps:

Extraction. We first resort to the tool of SciSpacy⁴(Neumann et al., 2019) (We detailed the usage in Appendix 9), which extracts all the biomedical entities from the document D and obtain an entity (i.e., node) set \mathcal{E}° . Then, we retrieve predefined relations from the static graph G_S and produce an edge set \mathcal{R}° . Finally, we construct an initial dynamic graph $G_D^\circ = \{(e_i, r_{i,j}, e_j) | e_i, e_j \in \mathcal{E}^\circ, r_{i,j} \in \mathcal{R}^\circ\}$, where e_i denotes a head entity, e_j is a tail entity, and $r_{i,j}$ denotes the relation. G_D° contains limited document entities and it should be expanded to cover richer biomedical knowledge.

Contextualization. Intuitively, we can choose k-hop expansion (we use one-hop in our experiments) for contextualization because the directly connected entities are always similar. However, such a simple expansion will bring ungovernable noisy entities and relations, which can pollute the dynamic graph generation quality. Instead, we propose a Dynamic Graph Generator (DGG) which produces triplets and expands G_D° as $G_D = \{(e_i, r_{i,j}, e_j) | e_i, e_j \in \mathcal{E}_D, r_{i,j} \in \mathcal{R}_D\}$, where \mathcal{E}_D and \mathcal{R}_D denote the expanded entity set and relation set respectively, $\mathcal{E}^\circ \subseteq \mathcal{E}_D$ and $\mathcal{R}^\circ \subseteq \mathcal{R}_D$.

Inspired by the masked language model (Devlin et al., 2018), we can randomly mask entities in the graph G_D° and train a DGG to recover the original topology. The trained generator, then, can detect the unobserved triplets while avoiding introducing noise. The training procedure is detailed below:

We mask 10% entities randomly from G_D° and label the masked entities \mathcal{E}_+° as positive instances, and label the remaining unmasked entities \mathcal{E}_-° as negative instances, and $\mathcal{E}^\circ = \mathcal{E}_+^\circ \cup \mathcal{E}_-^\circ$.

Then, we let e_i be the representation of any entity $e_i \in \mathcal{E}_-^\circ$ and e_j be the representation of any entity $e_j \in \mathcal{E}_+^\circ$. The initial entity representations are obtained based on pre-trained language models, like BERT and its variants (Devlin et al., 2018; Lee et al., 2020) as can be seen in Section 4.2. For each entity e_j to be recovered, we can calculate the selected probability based on an attention mechanism as below:

$$p(e_j | \mathcal{E}_-^\circ) = \delta \left(\text{Dense} \left(\sum_{e_i \in \mathcal{E}_-^\circ} \alpha_i e_i \right) \right) \quad (1)$$

$$\alpha_i = \sigma([e_i; e_j] \mathbf{W}_0)$$

where $\delta(\cdot)$ is a sigmoid function, $\text{Dense}(\cdot)$ is a fully-connected layer, $[\ ; \]$ denotes a vector concatenation operation, $\sigma(\cdot)$ denotes the ReLU function and \mathbf{W}_0 represents a trainable weight matrix.

To train the DGG, we use the cross-entropy function to minimize the loss between the masked enti-

⁴<https://allenai.github.io/scispacy/>

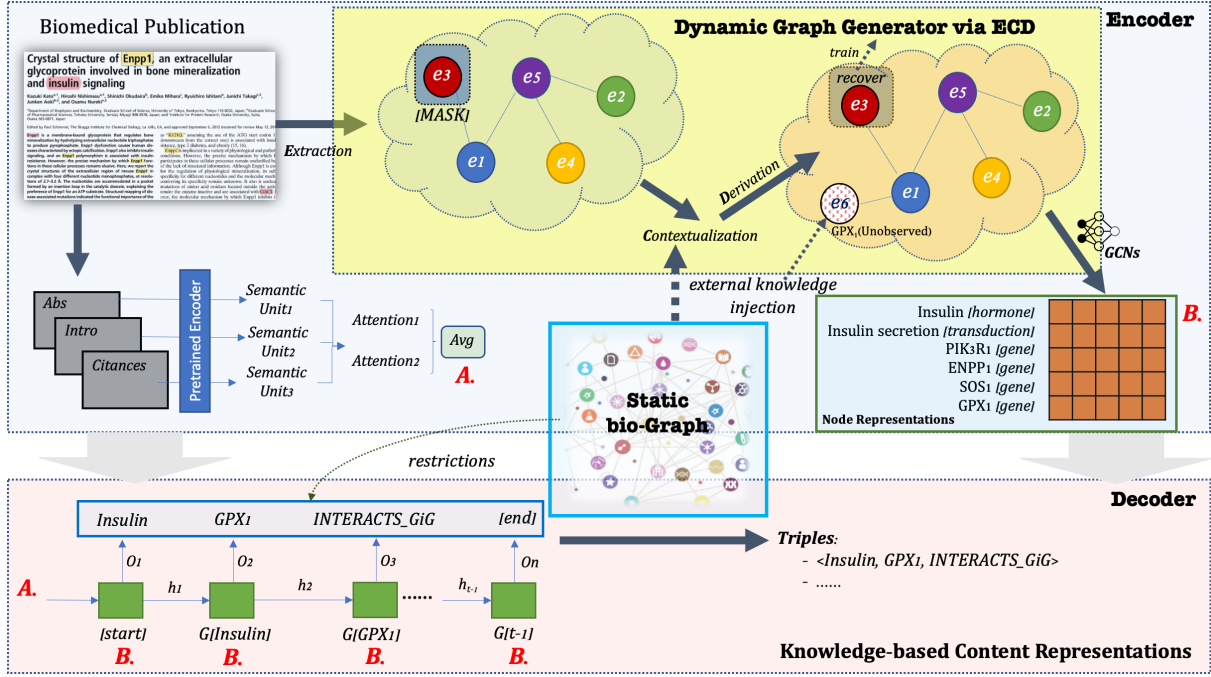


Figure 2: The DGRN architecture and ECD strategy. Given any biomedical publication, a pretrained language model and the ECD strategy can be used for content information modeling and dynamic graph generation. Then, both content and expanded entity-level information (denoted in red alphabets) can be used for decoder to generate observed and unobserved under the guidance of dual-graph.

ties and their corresponding recovered ones:

$$\mathcal{L}_{DGG} = \frac{1}{|\mathcal{E}_+^o|} \sum_{e_j \in \mathcal{E}_+^o} -\log(p(e_j | \mathcal{E}_-^o)) \quad (2)$$

Derivation In the derivation step, we feed the initial dynamic graph G_D^o to the trained DGG which considers one-hop entities as positive instances and the initial entities \mathcal{E}^o as negative instances. Then, we choose those entities with high probabilities and produce the expanded entity set $\mathcal{E}_D = \mathcal{E}^o \cup \{e_j\}_{\geq \text{Threshold}}$ and the expanded dynamic graph G_D .

4. Dual-Graph Resonance Network

In this section, we propose an innovative Dual-Graph Resonance Network (DGRN) which combines the dual graphs and generative method into a unified framework. The method is depicted in Figure 2, which consists of three modules: *Text Encoder*, *Graph Encoder* and *Triplet Decoder*.

4.1. Text Encoder Module

Given any document D , its constituent section S_* is fed into a BioBERT (Lee et al., 2020) to produce token-level representations ($\mathbf{H}^{abs} = \{\mathbf{h}_i^{abs}\}$, $\mathbf{H}^{int} = \{\mathbf{h}_i^{int}\}$ and $\mathbf{H}^{cit} = \{\mathbf{h}_i^{cit}\}$) and section-level representations (\mathbf{h}_{CLS}^{abs} , \mathbf{h}_{CLS}^{int} and

\mathbf{h}_{CLS}^{cit}). Viswanathan et al. (2021) found that citation sentences are more relevant to the document topics in comparison with abstract and introduction. Based on this, we opt for abstract-aware attention to measure the importance of each token representation \mathbf{h}_i^{cit} through a scoring function using a feed-forward neural network:

$$\alpha_i^{abs} = \text{softmax}((\mathbf{h}_{CLS}^{abs})^\top \sigma(\mathbf{W}_1 \mathbf{h}_i^{cit} + \mathbf{b}_1))$$

$$\mathbf{p}^{abs-cit} = \sum_{i \in [1, N_{cit}]} \alpha_i^{abs} \times \mathbf{h}_i^{cit} \quad (3)$$

Similarly, we also use introduction-aware attention to measure the importance of each token representation \mathbf{h}_i^{cit} through a scoring function as below:

$$\alpha_i^{int} = \text{softmax}((\mathbf{h}_{CLS}^{int})^\top \sigma(\mathbf{W}_2 \mathbf{h}_i^{cit} + \mathbf{b}_2))$$

$$\mathbf{p}^{int-cit} = \sum_{i \in [1, N_{cit}]} \alpha_i^{int} \times \mathbf{h}_i^{cit} \quad (4)$$

where \mathbf{W}_* and \mathbf{b}_* are trainable model parameters, $\mathbf{p}^{abs-cit}$ and $\mathbf{p}^{int-cit}$ represent the abstract-aware citance representation and introduction-aware citance representation, respectively. Finally, we take the average of the sum of the two vectors and obtain the document representation \mathbf{p}^{doc} as below:

$$\mathbf{p}^{doc} = \frac{1}{2} (\mathbf{p}^{abs-cit} + \mathbf{p}^{int-cit}) \quad (5)$$

4.2. Graph Encoder Module

We apply multi-layer Graph Convolutional Network (GCN) (Kipf and Welling, 2017) on the dynamic graph G_D to aggregate the features from neighbors to obtain node representation. For any entity e_i at the l -th layer, the graph convolutional operation can be applied by the formula as below:

$$\mathbf{e}_i^{(l+1)} = \sigma\left(\sum_{e_j \in \mathcal{N}_{e_i}^{(l)}} \mathbf{W}_3^{(l)} \mathbf{e}_j^{(l)} + \mathbf{b}_3^{(l)}\right) \quad (6)$$

where \mathcal{N}_{e_i} denotes the neighbors for the node e_i , $\mathbf{W}_3^{(l)}$ and $\mathbf{b}_3^{(l)}$ are trainable model parameters. In the initial stage, the $\mathbf{e}_i^{(0)} = \frac{1}{t-s+1} \sum_{i \in [s,t]} \mathbf{h}_i^*$ indicates entity i ranges from t -th token to s -th token in any section S_* . Intuitively, the GCN layers can encapsulate rich topological information.

4.3. Triplet Decoder Module

A decoder is adopted to generate knowledge triplets. Given training data, the decoder can copy an entity from the graph G_D as the head entity of the triplet, and then generate a relation for the triplet. Then, it can copy the tail entity from G_D . Repeating this process, the decoder could generate multiple triplets. In time step t ($1 \leq t$), we can calculate the decoder output \mathbf{o}_t and hidden state \mathbf{h}'_t as follows:

$$\mathbf{o}_t, \mathbf{h}'_t = f(\mathbf{x}_t, \mathbf{h}'_{t-1}) \quad (7)$$

where $f(\cdot)$ represents the RNN-based decoder function, \mathbf{h}'_{t-1} indicates the hidden state of time step $t-1$, and \mathbf{x}_t is the input representation of time step t and defined as below:

$$\mathbf{x}_t = [\mathbf{o}_{t-1}; \mathbf{c}_t] \mathbf{W}_4 \quad (8)$$

where \mathbf{o}_{t-1} denotes the entity or relation representation copied from dynamic graph in time step $t-1$, \mathbf{c}_t is the attention vector (Bahdanau et al., 2014) and \mathbf{W}_4 is a trainable weight matrix. In the initial step, input representation $\mathbf{x}_0 = \mathbf{p}^{doc}$ (see Eq. 5).

Attention Vector Entities and relations are generated and treated differently based on their different positions. In the time step t ($t\%3 = 0, 1$) (generating head or tail), the attention vector \mathbf{c}_t is calculated by copying entities from the entity set \mathcal{E}_D by the following formula:

$$\mathbf{c}_t = \sum_{e_i \in \mathcal{E}_D} \beta_i \times \mathbf{e}_i \quad (9)$$

$$\beta_i = \text{softmax}\left(\sigma([\mathbf{h}'_{t-1}; \mathbf{e}_i] \mathbf{W}_5)\right)$$

where \mathbf{h}'_{t-1} is the hidden state of the decoder in the $t-1$ time step, and \mathbf{W}_5 are trainable parameters.

In the time step t ($t\%3 = 2$) (generating relation), \mathbf{c}_t can be calculated by copying relations from relation set \mathcal{R}_D by the following formula:

$$\mathbf{c}_t = \sum_{r_{i,j} \in \mathcal{R}_D} \gamma_{i,j} \times \mathbf{r}_{i,j} \quad (10)$$

$$\gamma_{i,j} = \text{softmax}\left(\sigma([\mathbf{h}'_{t-1}; \mathbf{r}_{i,j}; \mathbf{p}_{i,j}] \mathbf{W}_6)\right)$$

where \mathbf{W}_6 are trainable parameters, $\mathbf{r}_{i,j}$ is the relation representation, $\mathbf{p}_{i,j}$ is the representation of edge between entity e_i and e_j via path reasoning.

Path Reasoning In the dynamic graph G_D , the head entity and tail entity in a triplet are not always directly connected. Thus, we introduce a path reasoning method that can model dependency among entities with multi-hop distances in the graph. Similar to (Zeng et al., 2020b), given the head entity and tail entity, we can define the representation of directed edge from entity e_i to entity e_j as below:

$$\mathbf{e}_{i,j} = \sigma(\mathbf{W}_7[\mathbf{e}_i; \mathbf{e}_j] + \mathbf{b}_7) \quad (11)$$

where \mathbf{W}_7 and \mathbf{b}_7 are trainable parameters, \mathbf{e}_i and \mathbf{e}_j denote representations of entity e_i and entity e_j respectively (see Eq. 6).

Based on the vectorized edge representation, the path between head entity e_i and tail entity e_j passing through entity e_o is represented as follow:

$$\mathbf{p}_{i,j} = [\mathbf{e}_{i,o}; \mathbf{e}_{o,j}; \mathbf{e}_{j,o}; \mathbf{e}_{o,i}] \quad (12)$$

For computation efficiency, we choose one-hop path, while it can be extended to multi-hop paths.

Entity Prediction To copy a head/tail entity, we calculate the confidence vector $\mathbf{q} = [q_1, \dots, q_{|\mathcal{E}_D|}]$ for all the entities in \mathcal{E}_D . We also apply a softmax on \mathbf{q} to obtain the probability distribution $\mathbf{p}^{entity} = [p_1^{entity}, \dots, p_{|\mathcal{E}_D|}^{entity}]$ by the formulas as below:

$$q_t = \sigma(\mathbf{o}_t \mathbf{W}_8 + \mathbf{b}_8) \quad (13)$$

$$\mathbf{p}^{entity} = \text{softmax}(\mathbf{q})$$

where \mathbf{W}_8 and \mathbf{b}_8 are trainable parameters. We select the entity with the highest probability as the predicted entity and use its embedding to produce the next time step input \mathbf{c}_{t+1} . *Note that the tail entity can not be the same as the head entity.*

Relation Prediction We now use a fully connected layer to calculate a confidence vector $\mathbf{q}' = [q'_1, \dots, q'_{|\mathcal{R}_D|}]$ of all the valid relations. Specifically, we apply a softmax on \mathbf{q}' to obtain the probability distribution $\mathbf{p}^{relation} = [p_1^{relation}, \dots, p_{|\mathcal{R}_D|}^{relation}]$ by the formulas as below:

$$\mathbf{q}' = \sigma(\mathbf{o}_t \mathbf{W}_9 + \mathbf{b}_9) \quad (14)$$

$$\mathbf{p}^{relation} = \text{softmax}(\mathbf{q}')$$

where \mathbf{W}_9 and \mathbf{b}_9 are trainable parameters. We select the relation with the highest probability as

the prediction relation and use its embedding to produce input c_{t+1} in the next time step.

Objective Function Our DGRN is trained with the negative log-likelihood loss function. Suppose $Y = \{y_1, y_2, \dots, y_T\}$ is the target result for triplets generation, the loss function is defined as:

$$\mathcal{L}_{SEQ} = \frac{1}{T} \sum_{t=1}^T -\log(p(y_t|y_{<t})) \quad (15)$$

where T is the maximum steps of the decoder, $p(y_t|y_{<t})$ denotes the conditional probability of target y_t given previous output sequence $y_{<t}$.

There are two negative log-likelihood loss functions used for training our DGRN; one is DGG loss (see Eq. 2) and one is the seq2seq loss (see Eq. 15). We optimize \mathcal{L}_{DGG} and \mathcal{L}_{SEQ} iteratively. We use backpropagation to calculate the gradients of all the trainable parameters and update them with Adam optimizer (Loshchilov and Hutter, 2019).

5. Experiments

5.1. Datasets

Motivation One of our contributions in this work is that we construct a new dataset, Bio-Sci, that derived from open-source PubMed corpus. The motivation behind Bio-Sci is that existing datasets rarely contain implicit relation triples and most relations are derived from sentence level (Luo et al., 2022). Our assumption is that some applications such as paper recommendation reason generation requires not only knowledge from query words but also from cross-sentence/document relations and thus we propose to inject both explicit and implicit triples for PubMed publications for a further investigation towards real-world applications such as paper recommendation.

Construction Following Wang et al. (2021), we use their public triplet dataset built from multiple public datasets⁵ as the bases to construct our static biomedical knowledge graph. This dataset consists of 1,426,025 triplets, 41,078 entities, and 27 relation types. Besides, we also construct a dataset with 32,330 biomedical publications collected from PubMed Central⁶, and split the dataset into training, development and testing sets with a split of 22,330/5,000/5,000. "To generate knowledge triples for every training/development document, we draw inspiration from the work of Xing et al. (2020) and Luo et al. (2022) and employ the pubmed_parser (Achakulvisut et al., 2020) to extract biomedical concepts for establishing connections between the entities in the Static Graph and each document. Subsequently, we annotate

the respective triples for each document. In the case of the testing set, we deliberately omit 10% of the entities and their corresponding sentences from each document, allowing us to assess the model's capability in entity recovery.

Human Agreement Test To evaluate the quality of our dataset, we conduct human agreement tests on a randomly selected small-scale and human-annotated testing set. Specifically, we first randomly selected 50 documents and provide 8 triplets as label candidates for each document. Two annotators with professional biomedical knowledge participated in the annotation task and chose the most appropriate triplets. The Kappa value of the manual annotation is 0.88. Besides, we conduct a paired t-test between the best performance baseline and our DGRN, the p-value is less than 0.05.

5.2. Experimental Settings

Since our experiments are conducted on biomedical data, we choose BioBert as our text encoder and implement DGRN with Pytorch and DGL (Wang et al., 2019). The hyper-parameter settings are detailed in Table 1. We follow the commonly used *precision*, *recall* and *F1-score* as the evaluation metrics⁷ (Yao et al., 2019). All the methods run on a server configured with 4 Tesla P100 GPUs, 16 CPUs, and 512G memory.

Baselines In Table 2, we compare our DGRN with several state-of-the-art baselines:

HRL is a hierarchical extraction paradigm which approaches relation extraction via hierarchical reinforcement learning (Takanobu et al., 2019).

CASREL is a cascade binary tagging framework, which models relations as functions that map subjects to objects in a sentence (Wei et al., 2020).

CopyRE is a Seq2Seq model which leverages copy mechanism through textual similarity for relation copy. (Zeng et al., 2018).

CopyMTL is a multi-task learning framework equipped with copy mechanism to allow the model to predict multi-token entities (Zeng et al., 2020a).

GAIN is a double graph reasoning network that aggregates mentions and their paths for better triplets extraction. (Zeng et al., 2020b).

AGGCN is a soft-pruning approach which automatically selects the relevant sub-structures of text for the relation extraction. (Guo et al., 2019).

For fair comparisons, we also utilize some state-of-the-art works that use pre-trained models:

KECI is an end-to-end model which utilizes external domain knowledge graph for joint entity and relation extraction (Lai et al., 2021).

TEMPGEN is a cross-attention guided model for triplet template generation (Huang et al., 2021a).

⁵<https://synlethdb.sist.shanghaiitech.edu.cn/#/download>

⁶https://ftp.ncbi.nlm.nih.gov/pub/pmc/oa_bulk/oa_comm

⁷Explanations are detailed in 9

GCN for Dual-Graph		DGRN Training	
Parameter	Value	Parameter	Value
Number of layers	1, 2, 3	learning rate	10^{-3} , 10^{-5}
emb size	100	dropout	0.2, 0.5 , 0.8
hidden size	808	batch_size	10, 20 , 50
		weight decay	10^{-3}

Table 1: The experimental settings of our method. The best parameter settings are highlighted.

Model	Model Type	P	R	F1
HRL	Extractive	61.17	21.81	32.16
CASREL	Extractive	71.12	32.94	45.03
CopyRE	Generative	54.73	25.72	34.99
CopyMTL	Generative	56.91	29.64	38.98
GAIN	Graph	56.01	21.43	31.00
AGGCN	Graph	61.43	33.91	43.70
KeBioLM	KG	61.18	32.88	42.77
UmlsBert	KG	59.61	29.17	39.17
KECI-BioBert	Extractive	61.93	40.81	49.20
TEMPGEN-BioBert	Generative	63.13	33.71	43.95
GAIN-BioBert	Graph	56.61	21.87	31.55
DGRN	Generative+Graph	73.77*	39.69*	51.61*

Table 2: Comparison among different models. Models above the double line do not use pre-trained model. Superscript * indicates statistical significance at $p < 0.05$ level compared to the best performance of baselines.

KeBioLM is a biomedical pretrained language model which uses external knowledge graph for relation extraction (Yuan et al., 2021).

UmlsBert integrates domain knowledge during the pre-training process for knowledge augmentation (Michalopoulos et al., 2021).

Note that for fair comparisons, we retain their best performance settings reported in their paper and use BioBert as the base for all the baselines that use pre-trained model⁸.

5.3. Comparative Study

Table 2 provides the main experiment results. One can witness the DGRN superiority compared with the best-performed baseline (KECI) with 4.90% improvement in F1. Interestingly, we find that generative methods can not compete with extractive approaches since generative methods require more searching space in the decoding stage than extractive models which only need to select target tokens among candidates. Meanwhile, the graph-based models achieve convincing performance, which indicates that graphs can provide essential topological information for knowledge extraction. We further investigate the performance of other state-of-the-art works using pre-trained model (BioBert). As can be observed in Table 2, DGRN achieves a 3.95% F1 improvement than TEMPGEN-BioBert

⁸The dataset and implementation can be found at <https://github.com/MatthewKKai/DGRN>

Model	ECD Step	P	R	F1
Extract	1	47.07	23.02	30.92
Extract+Text	1	59.03	38.21	46.39
Extract+Context+Text	1,2	41.17	20.81	27.65
DGRN (Full)	1,2,3	73.77	39.69	51.61

Table 3: Ablation study on DGRN components.

and a 14.42% improvement in precision. However, we notice that KECI-BioBert gets the highest recall among all models. This can be explained by the fact that KECI-BioBert first constructs a span graph to filter the noisy tokens before encoding the input while our method first encodes the sentence and can only process 1024 tokens in each iteration. However, our DGRN still performs the best which can be attributed to the aggregation of interactions between two graph. Moreover, we also compared our DGRN with knowledge-based methods of which KeBioLM achieves a 42.77 score on F1. We attributed the incomparable results to the fact that KeBioLM is designed for sentence-level relation extraction. Overall, DGRN outperforms all baselines because it utilizes topology information with dual-graph resonance to guide knowledge triplet generation while avoids disturbance from noisy information, which assures its effectiveness.

5.4. Ablation Study

We also explored how different configurations can impact our model’s performance. **Extract** only uses the initial graph. **Extract+Text** considers text modeling additionally. **Extract+Context+Text** considers one-hop path expansion by adding up to 20 entities. **DGRN** is our fully configured model. The experimental results are displayed in Table 3.

As can be seen from Table 3, **Extract+Text** outperforms **Extract** due to the semantic richness provided by text content. However, interestingly, **Extract+Context+Text** performs the worst which we attribute to the fact that a simple expansion strategy will introduce noisy entities/relations from the static knowledge graph. Beyond the encoder part, we also conduct experiments on by removing the Path Reasoning within the decoder. The results are 67.83, 38.87, 49.61 for precision, recall and f1, respectively. Our fully configured **DGRN** considers all the necessary components and achieves the best result, which again validates the effectiveness of our *Extraction-Contextualization-Derivation* strategy and the complete DGRN frame.

5.5. Robustness Study

In contrast to shorter documents, intricate and lengthy documents present a greater challenge for SciIE. Furthermore, the number of triplets within each document significantly influences the model’s

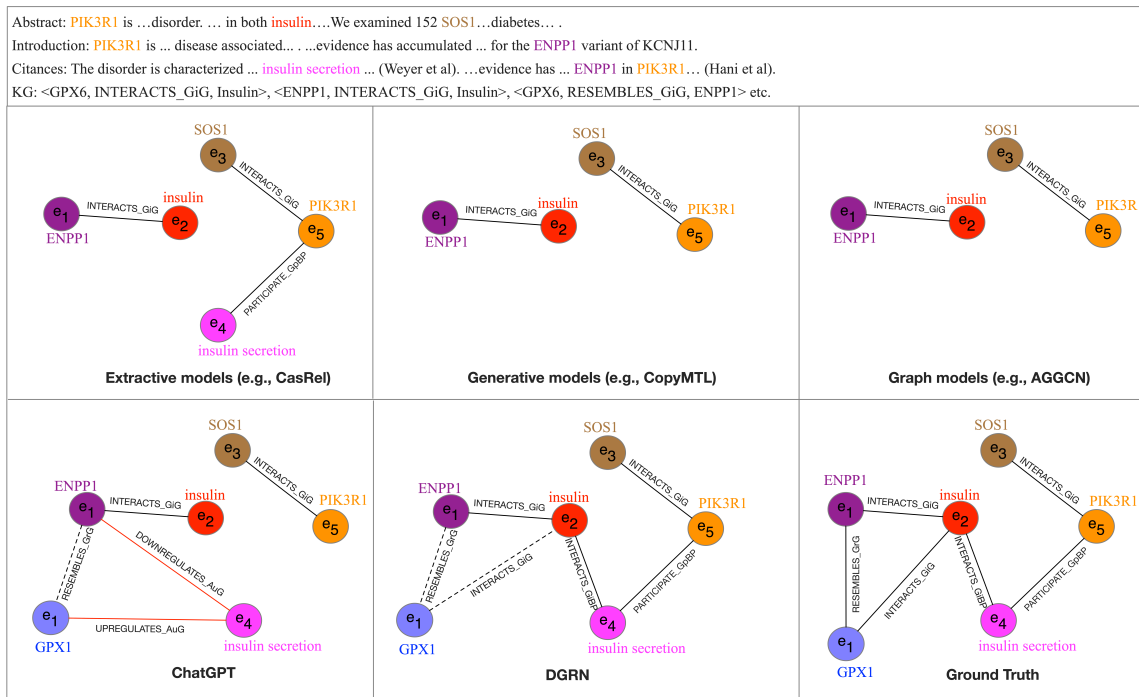


Figure 3: The exemplar generation results of baseline models (CASREL, CopyMTL, AGGCN) and DGRN.

Model	ECD Step	P	R	F1
Extract	1	45.61	19.97	27.78
Extract+Text	1	51.53	37.62	43.49
Extract+Context+Text	1,2	39.01	19.22	25.75
DGRN (Full)	1,2,3	73.77	39.69	51.61

Table 4: Robustness study on dynamic graph generation with training instances containing more than 10 triplets.

Model	Mask	Precision	Recall	F1-score
DGRN (Full)	5%	64.03	33.91	44.34
	10%	73.77	39.69	51.61
	15%	70.91	38.02	49.50
	20%	69.21	37.17	48.37
	25%	60.71	32.03	41.94

Table 5: Results on different percent of masked entities.

performance. Consequently, we have established an additional training set comprising instances with more than 10 triplets. The results in Table 4 clearly illustrate that complex documents have a detrimental impact on the performance of less sophisticated models. Nevertheless, the fully configured DGRN remains largely unaffected due to the synergistic relationship between the document content and the dynamically generated graph. This observation underscores the robustness of our model.

5.6. Masking Study

To validate the masking efficiency, we mask different percentages of nodes for DGG training. The

experimental results are presented in Table 5.

Table 5 reveals a notable trend: as the percentage of masked entities in the dynamic graph increases, the model's performance experiences a decline. This phenomenon is rooted in the dynamic graph's initiation process, which starts from a single document. As a result, the graphs constructed are typically rather small in scale, consisting of merely 10 to 20 nodes. Consequently, a higher masking percentage introduces more unseen entities, thereby amplifying the challenge of successful recovery. Further, the results of 5% masking setting are not comparable with the results of 10% setting. We attribute this to the fact that by 5% masking, only 0-1 nodes will be masked which leads to a performance degradation. In summary, masking approximately 10% of nodes proves beneficial in augmenting the DGG's ability to acquire valuable ontological information.

5.7. Case Study

As displayed in Figure 3, while other baselines successfully produce correct triplets, they can not compete with DGRN since they can't learn well from both rich BKG and the latent content. DGRN also shows its inference ability by generating "<GPX1, RESEMBLES_GiG, ENPP1>" "<GPX1, INTERACTS_GiG, Insulin>" based on the evidence derived from BKG. We further leverage in-context learning to let ChatGPT, one of the current most powerful AI tools, tackle this task. Despite the wrong relations it predicts for "insulin

secretion", it still surprises us by knowing the "RE-SEMABLE_GrG" relation between "GPX1" and "ENPP1" since this requires inference from external BKG. However, the better results of DGRN again support the superiority of our model.

6. Conclusions & Future Work

In this paper, we discuss the challenges of existing extractive and generative methods and make two efforts to deal with them. First, we propose a novel Extraction-Contextualization-Derivation (ECD) strategy to generate a document-specific dynamic graph. Then, we propose a novel Dual-Graph Resonance Network (DGRN) to generate richer triplets under the guidance of dual-graph. Extensive experiments validate the effectiveness of our proposed method. We also release a new dataset comes from Biomedical domain to encourage the community further explore this task.

In the future, we will integrate heterogeneous graph simplification (Wu et al., 2019) and sub-graph mining (Jiang et al., 2020) into our frame for better performance. Also, along with the representation learned from each paper and its citation network, we could use the triple extracted by our DGRN as backbone for paper recommendation reasons generation.

7. Ethics Statement

After carefully reviewing the COLING Ethics Policy, we sincerely show our respect and obey to consent all. To reproduce the experiment outcomes, we will make our dataset and code publicly available.

8. Acknowledgements

We gratefully acknowledge support from NSF, Award # 2122232 - SCISIPBIO: Constructing Heterogeneous Scholarly Graphs to Examine Social Capital During Mentored K Awardees Transition to Research Independence: Explicating a Matthew Mechanism.

9. Bibliographical References

Bilal Abu-Salih. 2021. Domain-specific knowledge graphs: A survey. *Journal of Network and Computer Applications*, 185:103076.

Titipat Achakulvisut, Daniel Acuna, and Konrad Kornding. 2020. *Pubmed parser: A python parser for pubmed open-access xml subset and medline xml dataset xml dataset*. *Journal of Open Source Software*, 5(46):1979.

Rie Kubota Ando and Tong Zhang. 2005. *A framework for learning predictive structures from multiple tasks and unlabeled data*. *Journal of Machine Learning Research*, 6:1817–1853.

Galen Andrew and Jianfeng Gao. 2007. *Scalable training of L_1 -regularized log-linear models*. In *Proceedings of the 24th International Conference on Machine Learning*, pages 33–40.

Isabelle Augenstein, Mrinal Das, Sebastian Riedel, Lakshmi Vikraman, and Andrew McCallum. 2017. *SemEval 2017 task 10: ScienceIE - extracting keyphrases and relations from scientific publications*. In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, pages 546–555, Vancouver, Canada. Association for Computational Linguistics.

Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. *Neural machine translation by jointly learning to align and translate*. *arXiv preprint arXiv:1409.0473*.

Hongxiao Bai and Hai Zhao. 2018. *Deep enhanced representation for implicit discourse relation recognition*. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 571–583, Santa Fe, New Mexico, USA. Association for Computational Linguistics.

Iz Beltagy, Kyle Lo, and Arman Cohan. 2019. *Scibert: Pretrained language model for scientific text*. In *EMNLP*.

BSI. 1973a. *Natural Fibre Twines*, 3rd edition. British Standards Institution, London. BS 2570.

BSI. 1973b. *Natural fibre twines*. BS 2570, British Standards Institution, London. 3rd. edn.

A. Castor and L. E. Pollux. 1992. The use of user modelling to guide inference and learning. *Applied Intelligence*, 2(1):37–53.

Hao Chen, Zepeng Zhai, Fangxiang Feng, Ruifan Li, and Xiaojie Wang. 2022. *Enhanced multi-channel graph convolutional network for aspect sentiment triplet extraction*. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2974–2985.

J.L. Chercœur. 1994. *Case-Based Reasoning*, 2nd edition. Morgan Kaufman Publishers, San Mateo, CA.

N. Chomsky. 1973. *Conditions on transformations*. In *A festschrift for Morris Halle*, New York. Holt, Rinehart & Winston.

- James W. Cooley and John W. Tukey. 1965. [An algorithm for the machine calculation of complex Fourier series](#). *Mathematics of Computation*, 19(90):297–301.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. [arXiv preprint arXiv:1810.04805](#).
- Umberto Eco. 1990. [The Limits of Interpretation](#). Indian University Press.
- Krishna Garg, Jishnu Ray Chowdhury, and Cornelia Caragea. 2021. Keyphrase generation beyond the boundaries of title and abstract. [arXiv preprint arXiv:2112.06776](#).
- John Giorgi, Gary D Bader, and Bo Wang. 2022. A sequence-to-sequence approach for document-level relation extraction. [arXiv preprint arXiv:2204.01098](#).
- Zhijiang Guo, Yan Zhang, and Wei Lu. 2019. Attention guided graph convolutional networks for relation extraction. [arXiv preprint arXiv:1906.07510](#).
- Sonal Gupta and Christopher Manning. 2011. [Analyzing the dynamics of research by extracting key aspects of scientific papers](#). In [Proceedings of 5th International Joint Conference on Natural Language Processing](#), pages 1–9, Chiang Mai, Thailand. Asian Federation of Natural Language Processing.
- Dan Gusfield. 1997. [Algorithms on Strings, Trees and Sequences](#). Cambridge University Press, Cambridge, UK.
- Paul Gerhard Hoel. 1971a. [Elementary Statistics](#), 3rd edition. Wiley series in probability and mathematical statistics. Wiley, New York, Chichester. ISBN 0 471 40300.
- Paul Gerhard Hoel. 1971b. [Elementary Statistics](#), 3rd edition, Wiley series in probability and mathematical statistics, pages 19–33. Wiley, New York, Chichester. ISBN 0 471 40300.
- Kung-Hsiang Huang, Sam Tang, and Nanyun Peng. 2021a. [Document-level entity-based extraction as template generation](#). In [Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing](#), pages 5257–5269, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Quzhe Huang, Shengqi Zhu, Yansong Feng, Yuan Ye, Yuxuan Lai, and Dongyan Zhao. 2021b. Three sentences are all you need: Local path enhanced document relation extraction. [arXiv preprint arXiv:2106.01793](#).
- Zhuoren Jiang, Zheng Gao, Jinjong Lan, Hongxia Yang, Yao Lu, and Xiaozhong Liu. 2020. Task-oriented genetic activation for large-scale complex heterogeneous graph embedding. In [Proceedings of The Web Conference 2020](#), pages 1581–1591.
- Vani Kanjirangat and Fabio Rinaldi. 2021. Enhancing biomedical relation extraction with transformer models using shortest dependency path features and triplet information. [Journal of Biomedical Informatics](#), 122:103893.
- Thomas N Kipf and Max Welling. 2016. Semi-supervised classification with graph convolutional networks. [arXiv preprint arXiv:1609.02907](#).
- Thomas N. Kipf and Max Welling. 2017. Semi-supervised classification with graph convolutional networks. In [International Conference on Learning Representations \(ICLR\)](#).
- Tuan Lai, Heng Ji, ChengXiang Zhai, and Quan Hung Tran. 2021. [Joint biomedical entity and relation extraction with knowledge-enhanced collective inference](#). In [Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing \(Volume 1: Long Papers\)](#), pages 6248–6260, Online. Association for Computational Linguistics.
- Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang. 2020. Biobert: a pre-trained biomedical language representation model for biomedical text mining. [Bioinformatics](#), 36(4):1234–1240.
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). In [International Conference on Learning Representations](#).
- Keming Lu, I Hsu, Wenxuan Zhou, Mingyu Derek Ma, Muhao Chen, et al. 2022. Summarization as indirect supervision for relation extraction. [arXiv preprint arXiv:2205.09837](#).
- Yi Luan, Luheng He, Mari Ostendorf, and Hannaneh Hajishirzi. 2018. [Multi-task identification of entities, relations, and coreference for scientific knowledge graph construction](#). In [Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing](#), pages 3219–3232, Brussels, Belgium. Association for Computational Linguistics.
- Yi Luan, Dave Wadden, Luheng He, Amy Shah, Mari Ostendorf, and Hannaneh Hajishirzi. 2019. [A general framework for information extraction](#)

- using dynamic span graphs. 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 3036–3046, Minneapolis, Minnesota. Association for Computational Linguistics.
- Ling Luo, Po-Ting Lai, Chih-Hsuan Wei, Cecilia N Arighi, and Zhiyong Lu. 2022. Biored: a rich biomedical relation extraction dataset. Briefings in Bioinformatics, 23(5):bbac282.
- George Michalopoulos, Yuanxin Wang, Hussam Kaka, Helen Chen, and Alexander Wong. 2021. UmlsBERT: Clinical domain knowledge augmentation of contextual embeddings using the Unified Medical Language System Metathesaurus. In Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, pages 1744–1753, Online. Association for Computational Linguistics.
- Preslav I Nakov, Ariel S Schwartz, Marti Hearst, et al. 2004. Citances: Citation sentences for semantic analysis of bioscience text. In Proceedings of the SIGIR, volume 4, pages 81–88. Citeseer.
- Guoshun Nan, Zhijiang Guo, Ivan Sekulić, and Wei Lu. 2020. Reasoning with latent structure refinement for document-level relation extraction. arXiv preprint arXiv:2005.06312.
- Tapas Nayak, Navonil Majumder, Pawan Goyal, and Soujanya Poria. 2021. Deep neural approaches to relation triplets extraction: A comprehensive survey. Cognitive Computation, 13(5):1215–1232.
- Tapas Nayak and Hwee Tou Ng. 2020. Effective modeling of encoder-decoder architecture for joint entity and relation extraction. In Proceedings of the AAAI conference on artificial intelligence, volume 34, pages 8528–8535.
- Mark Neumann, Daniel King, Iz Beltagy, and Waleed Ammar. 2019. ScispaCy: Fast and robust models for biomedical natural language processing. In Proceedings of the 18th BioNLP Workshop and Shared Task, pages 319–327, Florence, Italy. Association for Computational Linguistics.
- Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. 2017. Automatic differentiation in pytorch.
- Nanyun Peng, Hoifung Poon, Chris Quirk, Kristina Toutanova, and Wen-tau Yih. 2017. Cross-sentence n-ary relation extraction with graph lstms. Transactions of the Association for Computational Linguistics, 5:101–115.
- Kashif Rabbani, Matteo Lissandrini, and Katja Hose. 2023. Extraction of validating shapes from very large knowledge graphs. Proceedings of the VLDB Endowment, 16(5):1023–1032.
- Mohammad Sadegh Rasooli and Joel R. Tetreault. 2015. Yara parser: A fast and accurate dependency parser. Computing Research Repository, arXiv:1503.06733. Version 2.
- Charles Joseph Singer, E. J. Holmyard, and A. R. Hall, editors. 1954–58. A history of technology. Oxford University Press, London. 5 vol.
- Jannik Strötgen and Michael Gertz. 2012. Temporal tagging on different domains: Challenges, strategies, and gold standards. In Proceedings of the Eight International Conference on Language Resources and Evaluation (LREC’12), pages 3746–3753, Istanbul, Turkey. European Language Resource Association (ELRA).
- S. Superman, B. Batman, C. Catwoman, and S. Spiderman. 2000. Superheroes experiences with books, 20th edition. The Phantom Editors Associates, Gotham City.
- Ryuichi Takanobu, Tianyang Zhang, Jiexi Liu, and Minlie Huang. 2019. A hierarchical framework for relation extraction with reinforcement learning. In AAAI, pages 7072–7079.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. Advances in neural information processing systems, 30.
- Vijay Viswanathan, Graham Neubig, and Pengfei Liu. 2021. Citationie: Leveraging the citation graph for scientific information extraction. arXiv preprint arXiv:2106.01560.
- Minjie Wang, Da Zheng, Zihao Ye, Quan Gan, Mufei Li, Xiang Song, Jinjing Zhou, Chao Ma, Lingfan Yu, Yu Gai, Tianjun Xiao, Tong He, George Karypis, Jinyang Li, and Zheng Zhang. 2019. Deep graph library: A graph-centric, highly-performant package for graph neural networks. arXiv preprint arXiv:1909.01315.
- Shike Wang, Fan Xu, Yunyang Li, Jie Wang, Ke Zhang, Yong Liu, Min Wu, and Jie Zheng. 2021. Kg4sl: knowledge graph neural network

- for synthetic lethality prediction in human cancers. *Bioinformatics*, 37(Supplement_1):i418–i425.
- Xiang Wei, Xingyu Cui, Ning Cheng, Xiaobin Wang, Xin Zhang, Shen Huang, Pengjun Xie, Jinan Xu, Yufeng Chen, Meishan Zhang, et al. 2023. Zero-shot information extraction via chatting with chatgpt. [arXiv preprint arXiv:2302.10205](#).
- Zhepei Wei, Jianlin Su, Yue Wang, Yuan Tian, and Yi Chang. 2020. A novel cascade binary tagging framework for relational triple extraction. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1476–1488.
- Felix Wu, Amauri Souza, Tianyi Zhang, Christopher Fifty, Tao Yu, and Kilian Weinberger. 2019. Simplifying graph convolutional networks. In *International conference on machine learning*, pages 6861–6871. PMLR.
- Rui Xing, Jie Luo, and Tengwei Song. 2020. Biorel: towards large-scale biomedical relation extraction. *BMC bioinformatics*, 21(16):1–13.
- Wang Xu, Kehai Chen, and Tiejun Zhao. 2021. Document-level relation extraction with reconstruction. In *The 35th AAAI Conference on Artificial Intelligence (AAAI-21)*.
- 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 *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 764–777, Florence, Italy. Association for Computational Linguistics.
- Hongbin Ye, Ningyu Zhang, Shumin Deng, Mosha Chen, Chuanqi Tan, Fei Huang, and Hua-jun Chen. 2020. Contrastive triple extraction with generative transformer. [arXiv preprint arXiv:2009.06207](#).
- Hongbin Ye, Ningyu Zhang, Shumin Deng, Mosha Chen, Chuanqi Tan, Fei Huang, and Hua-jun Chen. 2021. Contrastive triple extraction with generative transformer. In *Proceedings of the Thirty-Fifth AAAI Conference on Artificial Intelligence*, pages 14257–14265.
- Zheng Yuan, Yijia Liu, Chuanqi Tan, Songfang Huang, and Fei Huang. 2021. [Improving biomedical pretrained language models with knowledge](#). In *Proceedings of the 20th Workshop on Biomedical Language Processing*, pages 180–190, Online. Association for Computational Linguistics.
- Daojian Zeng, Haoran Zhang, and Qianying Liu. 2020a. Copymtl: Copy mechanism for joint extraction of entities and relations with multi-task learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 9507–9514.
- Shuang Zeng, Runxin Xu, Baobao Chang, and Lei Li. 2020b. [Double graph based reasoning for document-level relation extraction](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1630–1640, Online. Association for Computational Linguistics.
- Xiangrong Zeng, Daojian Zeng, Shizhu He, Kang Liu, and Jun Zhao. 2018. [Extracting relational facts by an end-to-end neural model with copy mechanism](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 506–514, Melbourne, Australia. Association for Computational Linguistics.
- Kai Zhang, Kaisong Song, Yangyang Kang, and Xiaozhong Liu. 2023. [Content- and topology-aware representation learning for scientific multi-literature](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 7490–7502, Singapore. Association for Computational Linguistics.
- Ningyu Zhang, Xiang Chen, Xin Xie, Shumin Deng, Chuanqi Tan, Mosha Chen, Fei Huang, Luo Si, and Huajun Chen. 2021a. Document-level relation extraction as semantic segmentation. [arXiv preprint arXiv:2106.03618](#).
- Sheng Zhang, Patrick Ng, Zhiguo Wang, and Bing Xiang. 2022. Reknow: Enhanced knowledge for joint entity and relation extraction. [arXiv preprint arXiv:2206.05123](#).
- Yuhao Zhang, Derek Merck, Emily Tsai, Christopher D. Manning, and Curtis Langlotz. 2020. [Optimizing the factual correctness of a summary: A study of summarizing radiology reports](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5108–5120, Online. Association for Computational Linguistics.
- Zixuan Zhang, Nikolaus Parulian, Heng Ji, Ahmed Elsayed, Skatje Myers, and Martha Palmer. 2021b. [Fine-grained information extraction from biomedical literature based on knowledge-enriched Abstract Meaning Representation](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on*

Natural Language Processing (Volume 1: Long Papers), pages 6261–6270, Online. Association for Computational Linguistics.

Tianyang Zhao, Zhao Yan, Yunbo Cao, and Zhoujun Li. 2021. Asking effective and diverse questions: a machine reading comprehension based framework for joint entity-relation extraction. In Proceedings of the Twenty-Ninth International Conference on International Joint Conferences on Artificial Intelligence, pages 3948–3954.

Hao Zhu, Yankai Lin, Zhiyuan Liu, Jie Fu, Tat-Seng Chua, and Maosong Sun. 2019. Graph neural networks with generated parameters for relation extraction. arXiv preprint arXiv:1902.00756.

Appendix A. SciSpacy Details

For SciSpacy usage, we use them to obtain the entities with “en_core_sci_sm” and “en_ner_bionlp13cg_md” corpus because these two settings are trained in large-scale biomedical domain corpus. The f1 score is 70.87% on the entity extraction task for “en_core_sci_sm”, and 86.75% for “en_ner_bionlp13cg_md”. We mainly leverage “en_ner_bionlp13cg_md” in our experiments. What we do for linking entities is we first extract the entities and retrieve the BKG to see if there are matches for them.

Appendix B. Metrics

Following the partial match strategy, a generated triple is regarded as correct if the predicted relation and both the subject and object entity are correct. In this context, TP is the number of correctly generated/extracted triples, FN is the number of triples that the model failed to generate/extract, and FP is the number of incorrectly generated/extracted triples.