

Supporting Medical Relation Extraction via Causality-Pruned Semantic Dependency Forest

Yifan Jin^{1,2,*}, Jiangmeng Li^{1,2,*}, Zheng Lian^{1,2}, Chengbo Jiao³, Xiaohui Hu^{1,2}

¹University of Chinese Academy of Sciences

²Institute of Software Chinese Academy of Sciences

³University of Electronic Science and Technology of China
{yifan2020,jiangmeng2019,lianzheng2017,hxh}@iscas.ac.cn
chengbojiao@hotmail.com

Abstract

Medical Relation Extraction (MRE) task aims to extract relations between entities in medical texts. Traditional relation extraction methods achieve impressive success by exploring the syntactic information, e.g., dependency tree. However, the quality of the 1-best dependency tree for medical texts produced by an out-of-domain parser is relatively limited so that the performance of medical relation extraction method may degenerate. To this end, we propose a method to jointly model semantic and syntactic information from medical texts based on causal explanation theory. We generate dependency forests consisting of the semantic-embedded 1-best dependency tree. Then, a task-specific causal explainer is adopted to prune the dependency forests, which are further fed into a designed graph convolutional network to learn the corresponding representation for downstream task. Empirically, the various comparisons on benchmark medical datasets demonstrate the effectiveness of our model.

1 Introduction

Medical relation extraction (MRE) refers to identifying relations among entities from medical literature and reports. It plays a very important role in downstream tasks such as medical knowledge graph construction (Li et al., 2020; Rotmensch et al., 2017) and biomedical knowledge discovery (Quirk and Poon, 2016). On the other hand, as the number of medical literature increases, it becomes increasingly important to automatically discover the relation among entities in the literature (Peng et al., 2017).

The addition of syntactic structure has been demonstrated to be beneficial for various natural language processing tasks (Zareemoodi and Hafari, 2017; Zhou et al., 2005; Le and Zuidema,

2015). As a type of syntactic structure, the dependency tree capturing long-distance connections between words can indeed improve benchmark relation extraction methods (Tian et al., 2021; Chen et al., 2021; Zhang et al., 2018; Sun et al., 2020). We demonstrate an example in Figure 1. Specifically, the 1-best dependency tree of the sentence “*Aminopropylindenes* derived from Grundmann’s ketone as a novel chemotype of *oxidosqualene cyclase* inhibitors” in the CPR dataset. *Aminopropylindenes* and *oxidosqualene cyclase* are the entities, and the relation between them is “down regulator”, denoted as “CPR:4”.

However, in the medical field, the quality of the 1-best dependency tree generated by the out-of-domain parsers, e.g., parsers for the news domain, is relatively deficient. Generally, the main verb in a sentence is treated as the root node in the dependency tree, while, as the example shown in Figure 1, the entity *Aminopropylindenes*, apparently a *noun*, is treated as the root node. To solve this problem, multiple methods with dependency forests have been proposed (Song et al., 2019; Jin et al., 2020; Guo et al., 2021). Such approaches focus on redesigning the parser or substituting the parser with a semantic encoder, but the semantic and dependency tree syntactic information is used in a biased manner. Furthermore, the causality between the edges in the dependency forests and the performance of the model is not explored by benchmark methods.

To this end, we propose a novel approach, namely Causality-Pruned semantic dependency forest Graph Convolutional Network (CP-GCN). To acquire the dependency forests enriched with semantic and syntactic information in an unbiased manner, we first obtain the 1-best dependency tree, as the sentence syntactic information, which is generated by the out-of-domain parser, and then fuse the syntactic information with the semantic information by using a switch gate network. The seman-

*Equal contribution.

†Corresponding author.

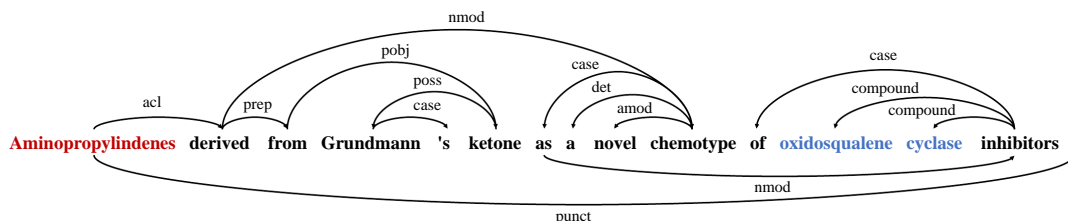


Figure 1: 1-best dependency tree for a biological sentence generated by the parser. *Aminopropylindenes* and *oxidosqualene cyclase* are the entities in the sentence.

tic information is captured in different representation subspaces using multi-head attention (Vaswani et al., 2017). To extract dependency forests’ edges that are causally related to the MRE performance, we construct a causal explanation dataset based on Granger causality (Granger, 1969, 1980) and train a task-specific causal explainer. We then obtain task-specific explanations of the dependency forests generated by the trained explainer and prune the dependency forests by following the corresponding explanations, which aim to eliminate the task-irrelevant information from the dependency forests. The pruned dependency forests are encoded by DCGCNs (Guo et al., 2019b) for MRE task. Empirically, the comparisons demonstrate that CP-GCN achieves state-of-the-art on benchmark relation extraction tasks, e.g., for the sentence-level relation extraction task, our model obtains 67.3 and 92.9 scores on CPR and PGR, respectively. The **contributions** are summarized as follows:

- We propose an approach to generate dependency forests enriched with semantic and syntactic information in an unbiased manner.
- We propose a causal pruning approach to remove task-irrelevant information from the dependency forests, which is achieved by using a task-specific explainer trained on a causal explanation dataset for the target MRE task.
- CP-GCN achieves state-of-the-art on benchmark MRE datasets, and the ablation comparisons further support the effectiveness of each part of our model.

2 Related Work

2.1 Medical Relation Extraction

Previous work performs the MRE task by constructing the 1-best dependency tree of sentences (Peng et al., 2017; Song et al., 2018). However, the accuracy of the 1-best dependency tree generated by the out-of-domain parser is relatively low, resulting

in a fall in MRE performance. Therefore, (Song et al., 2019) proposes to use dependency forests to solve this problem, which uses EDGEWISE and KBESTEISNER algorithm to pick edges to construct dependency forests. (Jin et al., 2020) encodes all effective dependency trees generated by a parser into dependency forests. (Guo et al., 2021) utilizes multi-head attention and Kirchhoff’s Matrix-Tree Theorem (MMT) (Koo et al., 2007) to automatically generate latent dependency forests without the usage of any parser. In general, (Song et al., 2019) and (Jin et al., 2020) focus more on the syntactic information in the 1-best dependency tree generated by the out-of-domain parser, while (Guo et al., 2021) directly discards the syntactic information and focuses only on the semantic information.

2.2 Causal Explanation

Causal explanation is designed to explain the importance of each module in a machine learning model on the prediction, which receives increasing attention recently (Datta et al., 2016; Schwab and Karlen, 2019; Lin et al., 2021). There are several viable forms of causality, including Granger causality (Granger, 1969), causal Bayesian networks (Pearl, 1985), and structural causal models (Pearl, 2009). (Chattopadhyay et al., 2019) proposes an attribution method based on the first principles of causality. (Schwab and Karlen, 2019) models the explanation task of image deep learning models as a causal learning task and proposes a causal explanation model based on Granger causality. (Lin et al., 2021) proposes a framework for explaining graph neural networks using the first principles of Granger causality.

3 Preliminaries

3.1 Task Definition

Our task is to extract relation between entities in a sentence, focusing on both binary relation extraction and ternary relation extraction. Formally, the input to our task is a sentence $S =$

$\{w_1, w_2, \dots, w_n\}$ with n words and w_i denotes the i -th word in the sentence. \mathcal{S} is annotated with entity mentions E_1 and E_2 ¹. The output is the relation between entities from a predefined relation set $\mathcal{R} = \{r_1, r_2, \dots, r_m\}$, where m denotes the number of relations.

3.2 Densely-Connected Graph Convolutional Networks

Graph Neural Network is a set of models that can effectively encode the information of graph structure, the classical models including Graph Attention Networks (GATs) (Velickovic et al., 2017), Graph Convolutional Networks (GCNs) (Kipf and Welling, 2016), etc. Densely-Connected Graph Convolutional Networks (DCGCNs) (Guo et al., 2019b) is a variant of GCNs, which introduces dense connections to GCNs. Thus being able to build multi-layer GCNs models with a large depth and learn richer information than the shallower GCNs models. More specifically, DCGCNs differs from GCNs in that the embedding of node v in the l -th layer receives information from all the preceding layers, which can be formulated as follows:

$$\mathbf{h}_v^{(l)} = \rho \left(\sum_{u \in \mathcal{N}(v)} \mathbf{W}^{(l)} \times \mathbf{g}_u^{(l)} + \mathbf{b}^{(l)} \right) \quad (1)$$

where \times denotes matrix multiplication, $\mathbf{h}_v^{(l)}$ is the embedding of node v in the l -th layer, ρ is an activation function, $\mathcal{N}(v)$ denotes the neighbours of node v , $\mathbf{W}^{(l)}$ and $\mathbf{b}^{(l)}$ are the weight matrix and bias vector of the l -th layer respectively, and $\mathbf{g}_u^{(l)}$ indicates the information about node u from all the preceding layers. Mathematically, $\mathbf{g}_u^{(l)}$ can be calculated by concatenating the initial embedding \mathbf{x}_u and the node embedding $\mathbf{h}_u^{(1)}; \dots; \mathbf{h}_u^{(l-1)}$ produced in layer $1, \dots, l-1$, respectively.

$$\mathbf{g}_u^{(l)} = [\mathbf{x}_u; \mathbf{h}_u^{(1)}; \dots; \mathbf{h}_u^{(l-1)}] \quad (2)$$

3.3 Dependency Tree Generation

To construct the 1-best dependency tree, we use Standard CoreNLP Toolkits (SCT) (Manning et al., 2014) to obtain the dependency tree \mathcal{T} for each input sentence \mathcal{S} and represent \mathcal{T} by a adjacency matrix $\mathbf{T} = (t_{i,j})_{n \times n}$ ², where $t_{i,j}$ is the dependency type (e.g., *dobj*) between w_i and w_j , e.g.,

¹ E_1, E_2 and E_3 for ternary relation extraction.

²The adjacency matrix \mathbf{T} adds the self-loop of each word to the dependency tree \mathcal{T} with the ‘‘self’’ dependency type and regards the dependencies between words as unoriented.

$t_{i,j} = 0$ if the connection between w_i and w_j do not exist. Then, we encode $t_{i,j}$ to the corresponding embedding $c_{i,j}^t$ with a learnable matrix, and use $\mathbf{C} = (c_{i,j}^t)_{n \times n}$ to denote the syntactic matrix.

4 Methodology

In this section, we introduce our proposed CP-GCN model shown in Figure 2.

4.1 Causality-Pruned Semantic Dependency Forest Generator

In the medical domain, the quality of the 1-best dependency tree generated by the out-of-domain parsers is relatively deficient. Thus, we propose the Causality-pruned Semantic dependency Forest Generator (CSFG) to generate dependency forests enriched with syntactic and semantic information and derive task-relevant information from them.

4.1.1 Semantic Embedding Module

In order to construct dependency forests that combine both semantic and syntactic information in an unbiased manner, we propose a semantic embedding module to incorporate the semantic information of the sentence into the 1-best dependency tree.

Specifically, we model semantic information using the multi-head attention mechanism (Vaswani et al., 2017) with N heads, which captures the semantic relevance between words in a sentence. For the p -th head, we compute the semantic matrix \mathbf{A}^p by using the query vector \mathbf{Q} and the key vector \mathbf{K} :

$$\mathbf{A}^p = \frac{(\mathbf{Q} \times \mathbf{W}^{\mathbf{Q}}) \times (\mathbf{K} \times \mathbf{W}^{\mathbf{K}})^{\top}}{\sqrt{d}} \quad (3)$$

where $\mathbf{W}^{\mathbf{Q}}$ and $\mathbf{W}^{\mathbf{K}}$ are learnable transformation matrices for \mathbf{Q} and \mathbf{K} , respectively, and d is the dimension of \mathbf{K} .

Finally, the p -th dependency forest can be obtained by summing the syntactic matrix \mathbf{C} and the semantic matrix \mathbf{A}^p with a switch gate network and a softmax function:

$$\mathbf{F}^p = \text{softmax}((1 - \alpha) \mathbf{A}^p + \alpha \mathbf{C}) \quad (4)$$

where $\alpha \in [0, 1]$ is a hyper-parameter to balance the syntactic matrix \mathbf{C} and the semantic matrix \mathbf{A}^p , and \mathbf{F}^p is the adjacency matrix of the p -th dependency forest.

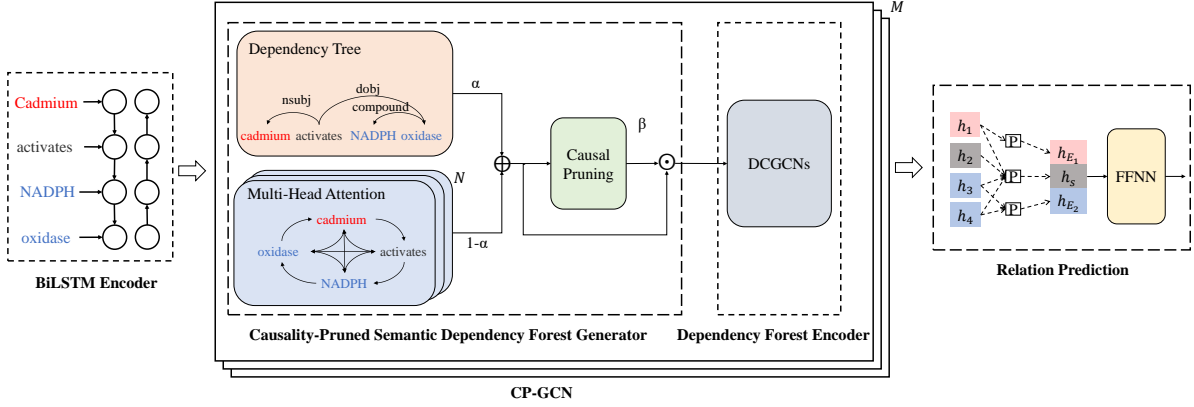


Figure 2: The overall architecture of CP-GCN with an example input sentence (“*Cadmium*” in red and “*NADPH oxidase*” in blue are two entities of the sentence). The model consists of three components: 1) BiLSTM Encoder obtains the sentence representation with a BiLSTM model. 2) CP-GCN is the main component of the model which contains M identical blocks, and each block contains two modules. Causality-Pruned Semantic Dependency Forest Generator combines the dependency tree and the multi-head attention with N heads to generate dependency forests and then prunes them using a task-specific causal explainer. Dependency Forest Encoder uses DCGCNs to encode the pruned dependency forests. 3) Relation Prediction module predicts relations using global and local max pooling and feedforward neural networks (FFNN).

4.1.2 Task-Specific Causal Pruning Module

In this part, our major objective is to extract dependency forests’ edges that are causally related to the MRE performance. Inspired by (Lin et al., 2021), we propose a method consisting of three processes for pruning dependency forests based on Granger causality. The first two processes aim to train a task-specific causal explainer, which are illustrated in Figure 3. The causal pruning process prunes the dependency forests with the trained causal explainer.

Causal explanation generation process. This process is designed to construct a causal explanation dataset for a specific MRE task. Given a pre-trained MRE model denoted by $f_{\text{MRE}}(\cdot)$ and the gold-standard relation r of the sentence \mathcal{S} . We start by using the semantic embedding module of the pre-trained MRE model to generate $N * M$ dependency forests of the sentence \mathcal{S} , denoted by $\mathcal{G} = \{\mathbf{G}^1, \mathbf{G}^2, \dots, \mathbf{G}^{N*M}\}$. For any dependency forest \mathbf{G}^i , it can be represented as $\mathbf{G}^i = (\mathbf{F}^i, \mathbf{H}_0)$, where \mathbf{F}^i is the fully-connected adjacency matrix indicating the weights of the edges, and \mathbf{H}_0 is the matrix of node features, which is the same for each dependency forest. Then, we need to extract the top \mathcal{K} edges from the dependency forests that are most relevant for predicting relation r . We implement this based on Granger causality.³

³Granger causality describes the causal relationships between two (or more) variables. Specifically, if we are better able to predict variable \tilde{y} using all information \mathbf{U} than excluding information about variable x , which means that

Specifically, we use $\mathcal{L}_{\mathcal{G}}$ to denote the model error of $f_{\text{MRE}}(\cdot)$ when taking the $N * M$ dependency forests into account, and $\mathcal{L}_{\mathcal{G} \setminus \{e_k\}}$ represents the model error excluding the edge e_k from each dependency forest. According to Granger causality, we can quantify the causal contribution of edge e_k to our MRE task by the change in model error after removing edge e_k :

$$\Delta_{e_k} = \mathcal{L}_{\mathcal{G} \setminus \{e_k\}} - \mathcal{L}_{\mathcal{G}} \quad (5)$$

where Δ_{e_k} represent the causal contribution of edge e_k .

To calculate $\mathcal{L}_{\mathcal{G}}$ and $\mathcal{L}_{\mathcal{G} \setminus \{e_k\}}$, we first take the $N * M$ dependency forests \mathcal{G} and $\mathcal{G} \setminus \{e_k\}$ as the input to the pre-trained model, respectively, and obtain their corresponding outputs $r_{\mathcal{G}}$ and $r_{\mathcal{G} \setminus \{e_k\}}$:

$$r_{\mathcal{G}} = f_{\text{MRE}}(\mathbf{F}^1, \dots, \mathbf{F}^{N*M}) \quad (6)$$

$$r_{\mathcal{G} \setminus \{e_k\}} = f_{\text{MRE}}(\mathbf{F}^1 \setminus \{e_k\}, \dots, \mathbf{F}^{N*M} \setminus \{e_k\}) \quad (7)$$

We then use the cross-entropy loss function to measure the model error, denoted as \mathbf{CE} .

$$\mathcal{L}_{\mathcal{G}} = \mathbf{CE}(r, r_{\mathcal{G}}) \quad (8)$$

$$\mathcal{L}_{\mathcal{G} \setminus \{e_k\}} = \mathbf{CE}(r, r_{\mathcal{G} \setminus \{e_k\}}) \quad (9)$$

Finally, we filter out the edges with the top \mathcal{K} causal contributions to form the causal explanation. In summary, our causal explanation dataset

the variable x helps predict variable \tilde{y} . Then we say that x Granger-causes \tilde{y} (Granger, 1980), denoted by $x \rightarrow \tilde{y}$.

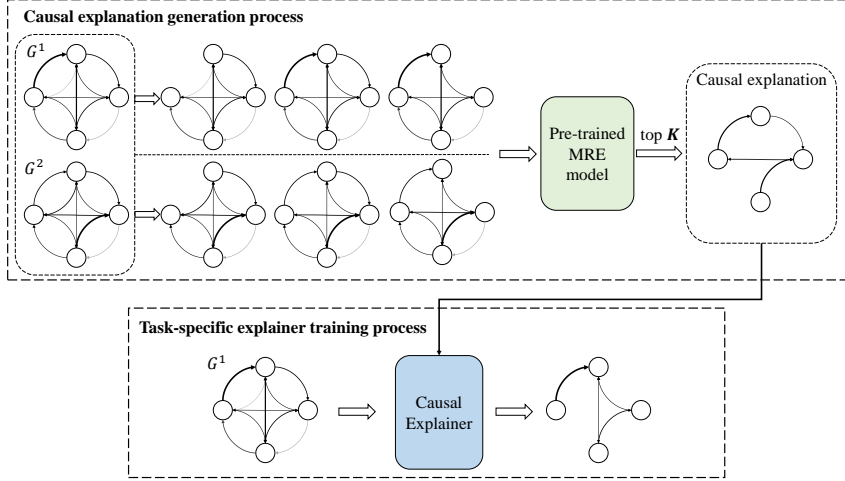


Figure 3: Illustration of training a task-specific causal explainer. Causal explanation generation process generates the causal explanations for the dependency forests using a pre-trained MRE model and the designed rules. Task-specific explainer training process trains a task-specific causal explainer with the *generated* causal explanation dataset.

is constructed with dependency forests and the corresponding causal explanations. Therefore, such a dataset is relevant to the specific MRE task.

Task-specific explainer training process. This process generates a task-specific explainer based on the causal explanation dataset. Following (Lin et al., 2021), we use an encoder-decoder architecture as the explainer. The encoder consists of several graph convolutional layers to aggregate information between neighbors in the dependency forest and learn node features. The decoder uses the inner product operation to obtain the explanation matrix. Specifically, the explanation matrix for G^i can be obtained by the explainer as:

$$\mathbf{X}^i = \sigma \left(f_{\text{GCN}}(\mathbf{F}^i, \mathbf{H}_0) \times f_{\text{GCN}}(\mathbf{F}^i, \mathbf{H}_0)^\top \right) \quad (10)$$

where $f_{\text{GCN}}(\cdot)$ denotes graph convolutional layers, \mathbf{X}^i is the explanation matrix and each value in \mathbf{X}^i represents the contribution of its corresponding edge to the prediction relation r , and σ is the activation function. **Causal pruning process.** Based on the pre-trained explainer, task-relevant explanation of the dependency forest can be obtained. Given the \mathbf{F}^p calculated by Eq. 4 and the pre-trained explainer, the explanation matrix \mathbf{X}^p corresponding to \mathbf{F}^p can be calculated via Eq. 10. Causal pruning for \mathbf{F}^p can be formulated as:

$$\hat{\mathbf{F}}^p = \text{softmax}(\mathbf{F}^p \odot (1 + \beta \mathbf{X}^p)) \quad (11)$$

where \odot is the element-wise multiplication, and $\beta \in [0, 1]$ is a hyper-parameter to control the coefficient of the explanation matrix \mathbf{X}^p .

4.2 Dependency Forest Encoder

Given N pruned dependency forests, DCGCNs are used to encode information from the forest structure. For the p -th pruned dependency forest, which is represented by the adjacency matrix $\hat{\mathbf{F}}^p$. We use DCGCNs with L layers to aggregate information about neighbors in $\hat{\mathbf{F}}^p$, and the representation of node i at the l -th layer can be calculated as:

$$\mathbf{h}_{p_i}^{(l)} = \rho \left(\sum_j^n \hat{\mathbf{F}}_{ij}^p \left(\mathbf{W}_p^{(l)} \times \mathbf{g}_{p_j}^{(l)} + \mathbf{b}_p^{(l)} \right) \right) \quad (12)$$

where $\hat{\mathbf{F}}_{ij}^p$ denotes the weight between node i and node j in $\hat{\mathbf{F}}^p$. $\mathbf{g}_{p_j}^{(l)}$ denotes the information about node j in the p -th pruned dependency forest from all the preceding layers and can be obtained by the same way as Eq. 2.

Then, we concatenate the representations obtained from the N dependency forests and fuse them using a linear layer. This process can be formulated as follows:

$$\mathbf{H}_b = \text{Linear}([\mathbf{H}^1; \mathbf{H}^2; \dots; \mathbf{H}^N]) \quad (13)$$

where \mathbf{H}^i is the node representations obtained by DCGCNs for the i -th dependency forest, and \mathbf{H}_b denotes the node representations of each block. M identical blocks are combined in the same way as above to obtain the final node representations for sentence S , denoted as \mathbf{H} .

4.3 Relation Prediction

To predict the relations among entities, the max pooling mechanism is used. We obtain the global

	CPR	PGR
TRAIN	16107	11780
DEV	10030	-
TEST	14269	219

Table 1: The number of instances of CPR and PGR.

sentence representation \mathbf{h}_S by applying the max pooling function to all the words in sentence S :

$$\mathbf{h}_S = \text{MaxPooling}(\{\mathbf{h}_1, \dots, \mathbf{h}_n\}) \quad (14)$$

where \mathbf{h}_i is the feature vector of word w_i , and then obtain the representation of each entity by applying the max pooling function to the words that belongs to an entity mention (i.e., E_q). Therefore, the entity representation of E_q can be obtained by:

$$\mathbf{h}_{E_q} = \text{MaxPooling}(\{\mathbf{h}_i | w_i \in E_q\}) \quad (15)$$

The sentence representation and entity representations are concatenated and fed into a feed-forward neural network (FFNN), and then we transform it into an m -dimensional vector \mathbf{h}_R using a linear layer to make a prediction:

$$\mathbf{h}_R = \text{Linear}(\text{FFNN}([\mathbf{h}_S; \mathbf{h}_{E_1}; \dots; \mathbf{h}_{E_Q}])) \quad (16)$$

where Q is 2 in the binary relational extraction task and is 3 in the ternary relational extraction task, m denotes the number of relations R .

5 Experiment

5.1 Datasets

We evaluate our model on three datasets with two types of tasks: cross-sentence n-ary relation extraction and sentence-level relation extraction following (Guo et al., 2021).

For the cross-sentence n-ary relation extraction task, we use the dataset extracted by (Peng et al., 2017) based on PubMed. Most of the instances in this dataset contain multiple sentences, and the entities in the instances are cross-sentence. In detail, this dataset contains 6987 instances of ternary relations and 6087 instances of binary relations, each of them is divided into five folders according to (Song et al., 2018). The relation between entities in each instance belongs to one of the relation sets, including “resistance or non-response”, “sensitivity”, “response”, “resistance”, and “None”. Following (Guo et al., 2021), we define two sub-tasks on this dataset: multi-class and binary-class relation extraction. For multi-class relation extraction,

we keep the original dataset unchanged, and for binary-class relation extraction, we define the first four relations as “Yes” and the “None” as “No”.

For the sentence-level relation extraction task, we use two datasets for Medical Relation Extraction, namely, BioCreative Vi CPR (CPR) (Krallinger et al., 2017) and Phenotype-Gene relation (PGR) (Sousa et al., 2019). CPR focuses on the relations between chemical components and human proteins, which contains six relation types (“CPR:3”, “CPR:3”, “CPR:4”, “CPR:5”, “CPR:6”, “CPR:9”, “None”). PGR focuses on whether human phenotypes are related to human genes, which contains two relation types (“TRUE” for related and “FALSE” for unrelated). The number of instances for train/dev/test sets of CPR and PGR datasets is shown in Table 1.

5.2 Implementation

During the causal explanation generation process, we use a pre-trained CP-GCN model *without* the task-specific causal pruning module as $f_{\text{MRE}}(\cdot)$ and choose 1/5 of the training set for the (Peng et al., 2017) dataset while the full training set for other datasets to generate the full dependency forests. Then, we set $\mathcal{K} = 20$ to construct causal explanation datasets.

For evaluation, we follow previous studies to use the test accuracy averaged over five cross validation folds for the cross-sentence n-ary task and F1 scores for the sentence-level task. Refer to the supplementary files for the details.

See Appendix A.1 for the hyper-parameter experiment on N , α , and β .

5.3 Results on Cross-Sentence N-Ary Relation Extraction Task

For the cross-sentence n-ary relation extraction task, We compare CP-GCN against two kinds of models and report the average test accuracies on the (Peng et al., 2017) dataset in table 2.

Tree: models use the 1-best dependency tree. DAG LSTM, GRN, and GCN(Full) use the full dependency tree directly, while GCN(Pruned) generates a pruned dependency tree with some rules (Zhang et al., 2018). Besides, DAG LSTM uses graph-structure LSTM to encode the dependency tree, while GRN and GCN use graph recurrent networks and graph convolutional networks, respectively.

Forest: models construct dependency forests. ACGCN treats a fully connected graph obtained by

Syntax Type	Model	Binary-class				Multi-class			
		Ternary		Binary		Ternary		Binary	
		Single	Cross	Single	Cross	Single	Cross	Single	Cross
Tree	DAG LSTM (Peng et al., 2017)	77.9	80.7	74.3	76.5	-	-	-	-
	GRN (Song et al., 2018)	80.3	83.2	83.5	83.6	-	71.7	-	71.7
	GCN(Full) (Zhang et al., 2018)	84.3	84.8	84.2	83.6	-	77.5	-	74.3
	GCN(Pruned) (Zhang et al., 2018)	85.8	85.8	83.8	83.7	-	78.1	-	73.6
Forest	AGGCN (Guo et al., 2019a)	87.1	87	85.2	85.6	-	79.7	-	77.4
	AGGCN* (Guo et al., 2019a)	86.3	87.2	86.3	85.8	77.7	78.7	77.7	77.3
	LF-GCN (Guo et al., 2021)	88	88.4	86.7	87.1	-	81.5	-	79.3
	LF-GCN* (Guo et al., 2021)	88.2	88.3	87	86.3	82.9	83.9	80	79.6
	AC-GCN (Qian et al., 2021)	88.8	88.8	86.8	86.5	-	84.6	-	81
	CP-GCN(ours)	89.5	89.1	87.3	86.5	84.3	84.9	81	80.1

Table 2: Average test accuracies on the (Peng et al., 2017) dataset for binary-class n-ary relation extraction and multi-class n-ary relation extraction. ‘‘Ternary’’ denotes drug-gene-mutation tuple and ‘‘Binary’’ denotes drug-mutation pair. ‘‘Single’’ means considering the instances within a single sentence, while ‘‘Cross’’ means considering all instances. Models with * indicate the accuracy of our reimplemention on their released implementation.

multi-head attention as a forest. LF-GCN automatically generates latent forests using multi-head attention and MMT. AC-GCN generates dependency forests with multi-head attention and encodes them with a 2D convolutional network.

As shown in Table 2, our proposed CP-GCN model achieves state-of-the-art performance in most settings. Specifically, the model using the pruned dependency tree performs better than those using the full dependency tree, suggesting that noisy information does exist in the 1-best dependency tree. In addition, the forest structure shows an advantage on this task, while CP-GCN surpasses the current state-of-the-art forest-structured model (AC-GCN) by 0.7 and 0.3 points on the binary-class ternary relation extraction task. The multi-class n-ary relation extraction task in (Peng et al., 2017) dataset is more challenging due to the unbalanced distribution of each relation, and CP-GCN can consistently achieve comparable performance.

5.4 Results on Sentence-Level Relation Extraction Task

For the sentence-level relation extraction task, we implement our approach on the CPR and PGR datasets and compare it against state-of-the-art models. We classify these models into three groups according to their syntax type.

None: models do not use tree or forest structures. Att-GRU adds a self-attention layer to GRU, and Bran uses a bi-affine self-attention model to capture interactions in sentences. BioBERT is a biomedical pre-trained language representation model.

Tree: models use the 1-best dependency tree. GCN, Tree-DDCNN, and Tree-GRN encode the full tree with GCN, DDCNN, and GRN, respec-

Syntax Type	Model	F1
None	Att-GRU (Liu et al., 2017)	49.5
	Bran (Verga et al., 2018)	50.8
Tree	GCN (Zhang et al., 2018)	52.2
	Tree-DDCNN (Jin et al., 2020)	50.3
	Tree-GRN (Jin et al., 2020)	51.4
Forest	Edgewise-GRN (Song et al., 2019)	53.4
	KBest-GRN (Song et al., 2019)	52.4
	AGGCN (Guo et al., 2019a)	56.7
	ForestFT-DDCNN (Jin et al., 2020)	55.7
	LF-GCN (Guo et al., 2021)	58.9
	AC-GCN (Qian et al., 2021)	65.8
	CP-GCN(ours)	67.3

Table 3: Main results on CPR.

tively. BO-LSTM prunes the tree, retaining only the shortest dependency path.

Forest: models construct dependency forests. Edgewise-GRN chooses edges with weights greater than the pre-defined threshold to form the dependency forest. KBest-GRN constructs the forest by aggregating K-best trees. ForestFT-DDCNN generates forests with a learnable dependency parser.

The results of CPR and PGR datasets are shown in Table 3 and Table 4. CP-GCN achieves state-of-the-art performance on both datasets. F1 score increases by 1.5 and 0.5 on the CPR and PGR datasets, respectively. Compared to models with forest structure, CP-GCN performs significantly better than both models with a bias towards syntactic information (Edgewise-GRN, KBest-GRN, and ForestFT-DDCNN) and models using almost only semantic information (AGGCN, LF-GAN, and AC-GCN), which demonstrates the effectiveness of our proposed CSFG method.

5.5 Analysis and Discussion

Ablation study. To validate the effectiveness of the ingredients of CP-GCN, i.e., the semantic em-

Syntax Type	Model	F1
None	BioBERT (Lee et al., 2020)	67.2
	BO-LSTM (Lamurias et al., 2019)	52.3
Tree	GCN (Zhang et al., 2018)	81.3
	Tree-GRN (Jin et al., 2020)	78.9
	Edgewise-GRN (Song et al., 2019)	83.6
Forest	KBest-GRN (Song et al., 2019)	85.7
	AGGCN (Guo et al., 2019a)	89.3
	ForestFT-DDCNN (Jin et al., 2020)	89.3
	LF-GCN (Guo et al., 2021)	91.9
	AC-GCN (Qian et al., 2021)	92.4
	CP-GCN(ours)	92.9

Table 4: Main results on PGR.

Model	F1
CP-GCN	67.3
-semantic embedding module	66.7
-task-specific causal pruning module	65.7

Table 5: An ablation study for CP-GCN on CPR dataset.

bedding module and the task-specific causal pruning module, we conduct the ablation study on CPR. We train the complete CP-GCN, an ablation model without the semantic embedding module, and another ablation model without the task-specific causal pruning module, respectively. Our experimental results are reported in Table 5. We observe that the performance of the model dropped (compared with complete CP-GCN) regardless of which module is removed, suggesting that both modules can help construct dependency forests that are more conducive to predicting relation. Comparing these two modules, the removal of the task-specific causal pruning module has a greater impact on performance, which suggests that the proposed causal pruning method can effectively distinguish vital information from noise.

Performance against sentence length. Figure 4 compares the F1 scores of our CP-GCN model and the LF-GCN model (Guo et al., 2021) under different sentence lengths. Following (Guo et al., 2021), We divide the test set of CPR into three groups ((0,25], (25,50], >50) based on sentence length. In general, CP-GCN outperforms LF-GCN against various sentence lengths. Otherwise, our model achieves a significant improvement on the more challenging long sentences, which demonstrates the ability of our model to capture long-range dependencies. Moreover, the dependency forests of the long sentences are more sophisticated, thus indicating that the task-specific causal explainer is able to extract task-relevant information from the sophisticated graph structure.

Case study. To further validate the efficiency of CP-GCN, we conduct a case study on an exam-

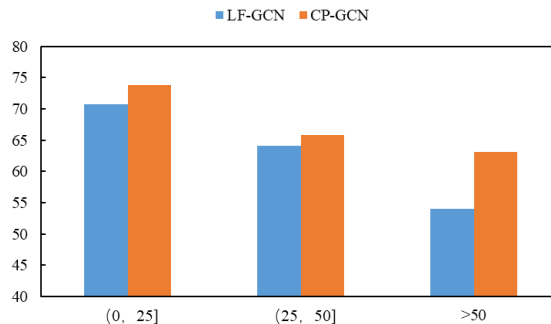


Figure 4: F1 scores against sentence length. The results on LF-GCN are reproduced based on its released implementation.

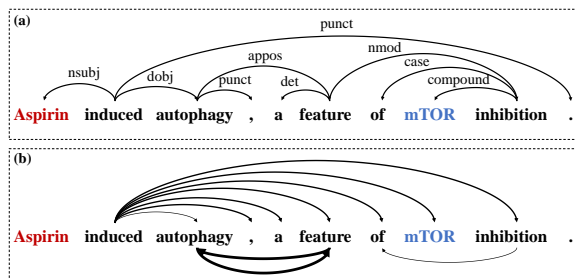


Figure 5: Visualizations of (a) 1-best dependency tree and (b) top 10 highest causal weight edges of the pruned dependency forest for the example input, where thicker lines denote the connections with higher causal weights.

ple sentence “*Aspirin* induced autophagy, a feature of *mTOR* inhibition”, which can be correctly predicted by our model to be a “down regulator” relation between *Aspirin*(E_1) and *mTOR*(E_2). Figure 5(a) shows its 1-best dependency tree, and Figure 5(b) visualizes the top 10 edges with the highest causal weights in the pruned dependency forest generated by the proposed CSFG and the thicker lines referring to higher causal weights. In this example, the connection between “*induced*” and “*feature*” enhances in the pruned dependency forest, and we reckon the latent reason is that CP-GCN can capture richer semantic information. We observe that there exists a strong connection between “*autophagy*” and “*feature*” in the pruned dependency forest, which improves the prediction of the relation between *Aspirin* and *mTOR*, supporting the effectiveness of the task-specific causal explainer.

6 Conclusion

In this paper, we introduce a novel approach for the medical relation extraction task, namely CP-GCN, which proposes a causality-pruned dependency forest enriched with semantic and syntactic information. We first construct dependency forests by incorporating semantic information into the dependency tree generated by the off-the-shelf parser.

Then, a task-specific causal explainer is trained to prune the dependency forests. Experiments on the benchmark medical datasets demonstrate the superiority of CP-GCN over the state-of-the-art methods for the medical relation extraction task.

Acknowledgements

This work is supported by the National Key Research and Development Program of China (No. 2019YFB1405100).

References

- Aditya Chattopadhyay, Piyushi Manupriya, Anirban Sarkar, and Vineeth N Balasubramanian. 2019. Neural network attributions: A causal perspective. In *International Conference on Machine Learning*, pages 981–990. PMLR.
- Guimin Chen, Yuanhe Tian, Yan Song, and Xiang Wan. 2021. Relation extraction with type-aware map memories of word dependencies. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 2501–2512.
- Anupam Datta, Shayak Sen, and Yair Zick. 2016. Algorithmic transparency via quantitative input influence: Theory and experiments with learning systems. In *2016 IEEE symposium on security and privacy (SP)*, pages 598–617. IEEE.
- Clive WJ Granger. 1969. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica: journal of the Econometric Society*, pages 424–438.
- Clive WJ Granger. 1980. Testing for causality: A personal viewpoint. *Journal of Economic Dynamics and control*, 2:329–352.
- Zhijiang Guo, Guoshun Nan, Wei Lu, and Shay B Cohen. 2021. Learning latent forests for medical relation extraction. In *Proceedings of the Twenty-Ninth International Conference on International Joint Conferences on Artificial Intelligence*, pages 3651–3657.
- Zhijiang Guo, Yan Zhang, and Wei Lu. 2019a. Attention guided graph convolutional networks for relation extraction. *arXiv preprint arXiv:1906.07510*.
- Zhijiang Guo, Yan Zhang, Zhiyang Teng, and Wei Lu. 2019b. Densely connected graph convolutional networks for graph-to-sequence learning. *Transactions of the Association for Computational Linguistics*, 7:297–312.
- Lifeng Jin, Lifeng Song, Yue Zhang, Kun Xu, Weiyun Ma, and Dong Yu. 2020. Relation extraction exploiting full dependency forests. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 8034–8041.
- Thomas N Kipf and Max Welling. 2016. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*.
- Terry Koo, Amir Globerson, Xavier Carreras Pérez, and Michael Collins. 2007. Structured prediction models via the matrix-tree theorem. In *Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL)*, pages 141–150.
- Martin Krallinger, Obdulia Rabal, Saber A Akhondi, Martín Pérez Pérez, Jesús Santamaría, Gael Pérez Rodríguez, Georgios Tsatsaronis, Ander Intxaurreondo, José Antonio López, Umesh Nandal, et al. 2017. Overview of the biocreative vi chemical-protein interaction track. In *Proceedings of the sixth BioCreative challenge evaluation workshop*, volume 1, pages 141–146.
- Andre Lamurias, Diana Sousa, Luka A Clarke, and Francisco M Couto. 2019. Bo-1stm: classifying relations via long short-term memory networks along biomedical ontologies. *BMC bioinformatics*, 20(1):1–12.
- Phong Le and Willem Zuidema. 2015. The forest convolutional network: Compositional distributional semantics with a neural chart and without binarization. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1155–1164.
- 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.
- Linfeng Li, Peng Wang, Jun Yan, Yao Wang, Simin Li, Jinpeng Jiang, Zhe Sun, Buzhou Tang, Tsung-Hui Chang, Shenghui Wang, et al. 2020. Real-world data medical knowledge graph: construction and applications. *Artificial intelligence in medicine*, 103:101817.
- Wanyu Lin, Hao Lan, and Baochun Li. 2021. Generative causal explanations for graph neural networks. In *International Conference on Machine Learning*, pages 6666–6679. PMLR.
- Sijia Liu, Feichen Shen, Yanshan Wang, Majid Rastegar-Mojarad, Ravikumar Komandur Elayavilli, Vipin Chaudhary, and Hongfang Liu. 2017. Attention-based neural networks for chemical protein relation extraction. *Training*, 1020(25.247):4157.
- Christopher D. Manning, Mihai Surdeanu, John Bauer, Jenny Rose Finkel, Steven Bethard, and David McClosky. 2014. *The stanford corenlp natural language processing toolkit*. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics, ACL 2014, June 22-27, 2014, Baltimore, MD, USA, System Demonstrations*, pages 55–60. The Association for Computer Linguistics.

- Judea Pearl. 1985. Bayesian networks: A model of self-activated memory for evidential reasoning. In *Proceedings of the 7th conference of the Cognitive Science Society, University of California, Irvine, CA, USA*, pages 15–17.
- Judea Pearl. 2009. *Causality*. Cambridge university press.
- 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.
- Mengyuan Qian, Jian Wang, Hongfei Lin, Di Zhao, Yijia Zhang, Wentai Tang, and Zhihao Yang. 2021. Auto-learning convolution-based graph convolutional network for medical relation extraction. In *China Conference on Information Retrieval*, pages 195–207. Springer.
- Chris Quirk and Hoifung Poon. 2016. Distant supervision for relation extraction beyond the sentence boundary. *arXiv preprint arXiv:1609.04873*.
- Maya Rotmensch, Yoni Halpern, Abdulhakim Tlimat, Steven Horng, and David Sontag. 2017. Learning a health knowledge graph from electronic medical records. *Scientific reports*, 7(1):1–11.
- Patrick Schwab and Walter Karlen. 2019. Cxplain: Causal explanations for model interpretation under uncertainty. *Advances in Neural Information Processing Systems*, 32.
- Linfeng Song, Yue Zhang, Daniel Gildea, Mo Yu, Zhiguo Wang, and Jinsong Su. 2019. Leveraging dependency forest for neural medical relation extraction. *arXiv preprint arXiv:1911.04123*.
- Linfeng Song, Yue Zhang, Zhiguo Wang, and Daniel Gildea. 2018. N-ary relation extraction using graph state lstm. *arXiv preprint arXiv:1808.09101*.
- Diana Sousa, André Lamúrias, and Francisco M Couto. 2019. A silver standard corpus of human phenotype-gene relations. *arXiv preprint arXiv:1903.10728*.
- Kai Sun, Richong Zhang, Yongyi Mao, Samuel Mensah, and Xudong Liu. 2020. Relation extraction with convolutional network over learnable syntax-transport graph. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 8928–8935.
- Yuanhe Tian, Guimin Chen, Yan Song, and Xiang Wan. 2021. Dependency-driven relation extraction with attentive graph convolutional networks. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4458–4471.
- 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.
- Petar Velickovic, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. 2017. Graph attention networks. *stat*, 1050:20.
- Patrick Verga, Emma Strubell, and Andrew McCallum. 2018. Simultaneously self-attending to all mentions for full-abstract biological relation extraction. *arXiv preprint arXiv:1802.10569*.
- Poorya Zareemoodi and Gholamreza Haffari. 2017. Incorporating syntactic uncertainty in neural machine translation with forest-to-sequence model. *arXiv preprint arXiv:1711.07019*.
- Yuhao Zhang, Peng Qi, and Christopher D Manning. 2018. Graph convolution over pruned dependency trees improves relation extraction. *arXiv preprint arXiv:1809.10185*.
- GuoDong Zhou, Jian Su, Jie Zhang, and Min Zhang. 2005. Exploring various knowledge in relation extraction. In *Proceedings of the 43rd annual meeting of the association for computational linguistics (acl'05)*, pages 427–434.

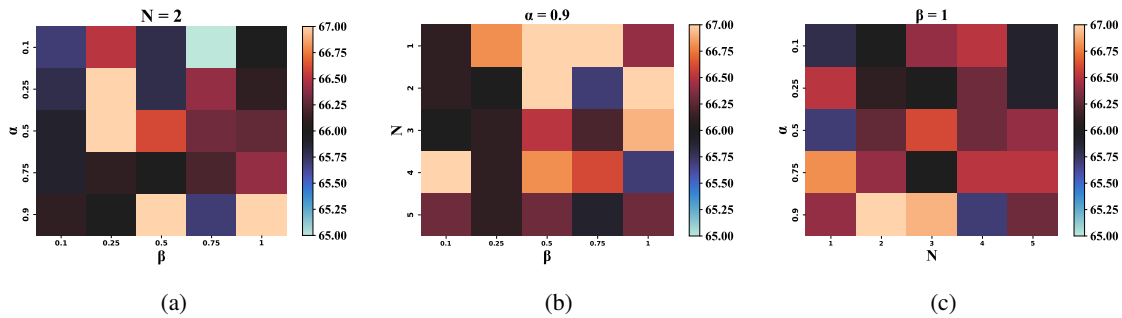


Figure 6: F1 scores with different hyper-parameter settings.

A Appendix

A.1 Hyper-Parameter Experiment

We perform several experiments on the CPR dataset to study the influence of the hyper-parameters in our proposed CP-GCN model, and the results are shown in Figure 6. The hyper-parameter α balances semantic and syntactic information in the dependency forest. The hyper-parameter β balances the impact of the task-specific causal pruning module. The hyper-parameter N represents the richness of semantic information. As (a), (b), and (c) are shown in Figure 6, our proposed CP-GCN model achieves comparable performance in most settings, which indicates the robustness of our model. Specifically, CP-GCN achieves the highest F1 score 67.3 with $N = 2$, $\alpha = 0.9$, and $\beta = 1$. As shown in Figure 6(c), when N decreases to 1, i.e. the semantic information decreases, CP-GCN performs best when the weight of the dependency tree, α , decreases as well, suggesting that our model is able to balance the syntactic and semantic information. As shown in Figure 6(a), when the weight of dependency tree α increases, CP-GCN performs best when the weight of task-specific causal pruning module β increases as well. This demonstrates that there is indeed some noise in the dependency tree and our proposed task-specific causal pruning module can remove task-irrelevant information.