

Communicative Message Passing for Inductive Relation Reasoning

Sijie Mai*, Shuangjia Zheng*, Yuedong Yang[†], Haifeng Hu[†]

Sun Yat-sen University

{maisj, zhengshj9}@mail2.sysu.edu.cn, {yangyd25, huhaf}@mail.sysu.edu.cn

Abstract

Relation prediction for knowledge graphs aims at predicting missing relationships between entities. Despite the importance of inductive relation prediction, most previous works are limited to a transductive setting and cannot process previously unseen entities. The recent proposed subgraph-based relation reasoning models provided alternatives to predict links from the subgraph structure surrounding a candidate triplet inductively. However, we observe that these methods often neglect the directed nature of the extracted subgraph and weaken the role of relation information in the subgraph modeling. As a result, they fail to effectively handle the asymmetric/anti-symmetric triplets and produce insufficient embeddings for the target triplets. To this end, we introduce a **Communicative Message Passing** neural network for **Inductive reLation rEasoning**, **CoMPiLE**, that reasons over local directed subgraph structures and has a vigorous inductive bias to process entity-independent semantic relations. In contrast to existing models, CoMPiLE strengthens the message interactions between edges and entities through a communicative kernel and enables a sufficient flow of relation information. Moreover, we demonstrate that CoMPiLE can naturally handle asymmetric/anti-symmetric relations without the need for explosively increasing the number of model parameters by extracting the directed enclosing subgraphs. Extensive experiments show substantial performance gains in comparison to state-of-the-art methods on commonly used benchmark datasets with variant inductive settings.

Introduction

Knowledge graphs (KGs) are collections of factual information represented in the form of relational triplets. Each relation triplet can be organized as (h, r, t) where h and t represents the head and tail entities and r is the relation between h and t . KGs have played a critical role across variety of tasks like Question Answering (Zhang et al. 2017), Semantic Search (Xiong, Power, and Callan 2017), Dialogue Generation (He et al. 2017) and many more. However, because of the limitations of human knowledge and extraction algorithms, they tend to suffer from incompleteness, that is, absent links in the KGs. Numerous methods

have been developed to fulfill the gap between KGs and real-world knowledge, which are referred as link prediction or knowledge graph completion tasks. After condensing entities and relations into a low-dimensional vector space, these models predict missing facts by operating the entity and relation embeddings.

Despite the impressiveness of model performance, most of KG embeddings methods inherently assume a fixed set of entities in the graph and ignore the evolving nature of KGs. In fact, many real-world KGs are ever-evolving (Trivedi et al. 2017), with novel entities being added over time e.g., new users in social media platforms or new molecules in biomedical knowledge graphs. As these entities were unknown to the model during training, the model does not obtain their embeddings, and hence have no means to infer the relations for these entities. Thus, the ability to make predictions on new entities avoiding costly re-training from scratch is desirable for production-ready machine learning models.

Some efforts have been made to obtain the inductive embeddings for new entities using external resources (Wang et al. 2014; Xie et al. 2016b; Zhong et al. 2015). Although these approaches may be useful, they required extra computation over massive resources, which are time-consuming and may not always be feasible. An alternative is to induce probabilistic logic rules by enumerating statistical regularities and patterns embedded in the knowledge graph (Meilicke et al. 2018; Yang, Yang, and Cohen 2017), which is entity-independent and hence inherently inductive. However, these methods suffer from scalability issues and are less expressive due to their rule-based nature.

Taking inspiration from the success of graph neural networks (GNNs) in graph data modeling, extensive studies applied GNNs to learn the embeddings for KGs (Schlichtkrull et al. 2018; Shang et al. 2019; Nathani et al. 2019), among which most of them were adopted in transductive settings. By extending a subgraph-based link prediction method that performs in the basic network (Zhang and Chen 2018), recent work GraIL (Teru, Denis, and Hamilton 2020) shed light on the evolution issue for KG. The basic strategy behind GraIL can be split into three steps: (i) extracting the enclosing subgraph surrounding the target triplets, (ii) annotating the relative position of each entity in the extracted subgraph, and (iii) scoring the annotated subgraph with a GNN. It can be beneficial to the inference of entirely novel entities

*These two authors contributed equally.

[†]Corresponding authors.

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

that not be surrounded by known nodes, and there is no need of the domain-related initial embedding of these emerging entities. The experimental results of applying GraIL demonstrate the feasibility of utilizing a subgraph-based model on inductive reasoning in the context of KGs.

One drawback of GraIL is that it ignores the directionality nature of the KG when extracting the enclosing subgraph for the target triplet. Although the direction is nominally preserved during message passing by aggregating the incoming information for an entity, the subsequent symmetric scoring function make it incapable of effectively handling binary relationships in KGs, especially for asymmetric (e.g., nominatedfor) and anti-symmetric relations (e.g., filiation). Moreover, the adopted vertex-based message passing network (Xu et al. 2018) weakens the role of relation embeddings, which violates the nature of the inductive relation reasoning because inductive setting is entity-independent and relies on the relation information to conduct inference.

Based on the above observations, we propose a novel Communicative Message Passing neural networks for Inductive reLation rEasoning (CoMPILE). In CoMPILE, we first extract directed enclosing subgraph for each triplet instead of an undirected one. A communicative message passing network framework is then extended to strengthen the information interactions between entities and relations while update both the edge and entity embeddings simultaneously. We also apply an edge-aware attention mechanism to aggregate the local neighborhood features and gather the global entity information to enrich the entity/relation representation. In contrast to previous efforts (Nickel, Rosasco, and Poggio 2015) that require an explosion of the number of parameters to handle binary relations, our model is naturally capable of dealing with asymmetric and anti-symmetric relations by communicative message passing within directed subgraphs, without the unnecessary explosive increase of the number of parameters.

In brief, the main contributions are listed below:

- Introducing a competitive inductive knowledge graph embedding model, CoMPILE, that fits the directional nature of knowledge graph and can naturally deal with asymmetric and anti-symmetric relations.
- Introducing a novel node-edge communicative message passing mechanism to strengthen the role of relation information in the subgraph and fit the nature of the inductive setting.
- Evaluating CoMPILE and several previously proposed models on three inductive datasets: our model achieves state-of-the-art AUC-PR and Hits@10 across most of them. We also extract new inductive datasets by filtering out the triplets that have no enclosing subgraph to evaluate the inductive relation reasoning more accurately.

Related Work

Transductive Relation Prediction

Representation learning on KGs has been an active research area due to the wide applications of the resultant entity and relation embeddings. Typical KG embedding models




Figure 1: Comparisons of different subgraph extraction strategies. The left one is the undirected subgraph that considers all the h -hop neighbors of both the target head and tail. The subgraph on the right refer to the enclosing subgraph that only considers the h -hop common neighbors of target head and target tail.

include TransE (Bordes et al. 2013), Distmult (Yang et al. 2014), ComplEx (Trouillon et al. 2017), ConvE (Dettmers et al. 2017), to name a few. While these kinds of methods process each triplet independently without the consideration of semantic and structure information embedded in rich neighborhoods. Recently, GNNs have been used to capture global structural information inherently stored in KGs and have been shown to achieve state-of-the-art performance on a variety of datasets (Schlichtkrull et al. 2018; Shang et al. 2019; Nathani et al. 2019). However, the above approaches on KG embedding mostly work in transductive settings.

Inductive Relation Prediction

One research line of inductive representation learning is to introduce additional attributes such as description text or images to embed unseen entities (Xie et al. 2016b,a; Shi and Weninger 2017). Although the resultant embeddings can be utilized for KG completion, these methods heavily relied on the presence of external resources which is not present in many KGs. To relieve this issue, several inductive KG embedding models (Hamaguchi et al. 2017; Wang et al. 2019) are proposed to aggregate the information of existing neighbors of an emerging entity with graph neural networks. However, both of these approaches demand the new nodes to be surrounded by known nodes and cannot handle entirely new graphs.

Another research line is rule-based approaches that use observed co-occurrences of frequent patterns in the knowledge graph to recognize logical rules (Galárraga et al. 2015). They are inherently inductive since the logical-rules are independent of entities, but these approaches suffer from scalability issues and lack expressive power due to their rule-based nature. Inspired by these statistical rule-induction ap-

proaches, several differentiable rule learners including NeuralLP (Yang, Yang, and Cohen 2017), RuleN (Meilicke et al. 2018), and DRUM (Sadeghian et al. 2019) are proposed to learn the logical rules as well as confidence scores from KGs in an end-to-end paradigm. However, they did not take account of the neighbor structure surrounding the predicted relations, hence is not expressive enough when the paths between head and tail entities are sparse. This set of methods, together with GraIL (Teru, Denis, and Hamilton 2020) constitute our baselines.

CoMPILE

Denotations and Task Definition

A target triplet in the knowledge graph is denoted as (s, r, t) where s , r , and t refers to the head entity, relation, and tail entity, respectively. Inductive relation reasoning aims to score the plausibility of target triplet (h_T, r_T, t_T) , where the representations of h_T and t_T are not available during prediction. In this work, we use an enclosing directed subgraph to represent the target triplet (h_T, r_T, t_T) . The enclosing subgraph between the target head and tail is denoted as $G = (V, E)$ where V and $E \in V \times V$ denotes the set of nodes and observed edges in the subgraph G , respectively. We use N_e to represent the number of edges in the subgraph. The embeddings of nodes is denoted as $N \in \mathbb{R}^{N_n \times d_n}$ where N_n is the number of the nodes in the subgraph. The relation embedding is denoted as $R \in \mathbb{R}^{N_r \times d}$ (N_r is the number of relations), which is parameterized as a learnable matrix updated by gradient descent and is shared across train and test graphs. We define the head-to-edge, relation-to-edge, and tail-to-edge adjacency matrix as $A^{he} \in \mathbb{R}^{N_n \times N_e}$, $A^{re} \in \mathbb{R}^{N_r \times N_e}$, and $A^{te} \in \mathbb{R}^{N_n \times N_e}$, which maps the tail, relation, and head to the corresponding edge, respectively. The values in the adjacency matrices are either 0 or 1 where 0 denotes no connections.

Directed Subgraph Extraction

In this section, we illustrate the procedure to extract the directed enclosing subgraph. Different from GraIL (Teru, Denis, and Hamilton 2020) that extracts undirected enclosing subgraph between the target head and tail which ignores the direction of the target triplet, inspired by the mechanism that human uses to infer logical rules, we demonstrate the superiority of directed enclosing subgraph. A triplet has direction, representing the relation from head to tail. Considering two triplets (h_T, r_T, t_T) and (t_T, r_T, h_T) , if we use undirected enclosing subgraph, then the predictions are likely to be very close for these two triplets because the enclosing subgraphs are the same. However, only one of them is true if the relation r_T is asymmetric. Therefore, we need to use directed enclosing subgraph to handle these kinds of relation more effectively. Moreover, by this means, we can solve the direction problem in KGs without increasing the model complexity (the time complexity even decreases since the directed enclosing subgraph is obtained by pruning the undirected enclosing subgraph, as shown in Fig. 1).

The enclosing subgraph is restricted by the hop number h , and $h + 1$ is the maximum distance from target head to target

tail. To extract the h -hop directed enclosing subgraph, we first introduce the incoming neighbors and outgoing neighbors. Given a triplet (s, r, t) , we define s as the 1-hop incoming neighbor of t , and t as the 1-hop outgoing neighbor of s , and h -hop incoming/outgoing neighbors vice versa. Firstly, we extract h -hop outgoing neighbors of target head and the h -hop incoming neighbors of target tail. Then we extract the 1-hop incoming neighbor of target tail. If the h -hop outgoing neighbors of target head and the 1-hop incoming neighbor of target tail have common entities, then the directed subgraph between target head and tail is existed. Afterwards, We find the common entities (nodes) of h -hop outgoing neighbors of target head and the h -hop incoming neighbor of target tail (if target head or tail are not in the common entities, we need to add them into common entities). Finally we add the edges (triplets) whose head and tail both belong to the common entities to construct the subgraph. By this means, the maximum distance from target head to tail would become $h + 1$ even though we only extract h -hop neighbors of them.

Node/Edge Embedding Initialization

Since the entities are unseen, we need to define an entity-independent embedding for each entity (node) in the subgraph. Similar to GraIL (Teru, Denis, and Hamilton 2020), we initialize the node embedding by the distances to the target head and target tail to capture the relative position of each node in the subgraph. The node embedding for node i is defined as $N_i = \text{one-hot}(d_{hi}) \oplus \text{one-hot}(d_{it}) \in \mathbb{R}^{2(h+2)}$ where d_{hi} denotes the minimum distance from the target head to node i , and d_{it} denotes the minimum distance from the node i to target tail (note that the distance definition here also meets the ‘directed’ requirement). For edge i , i.e. (h_i, r_i, t_i) , the initialized edge embedding is defined as $E_i = N_{h_i} \oplus R_{r_i} \oplus N_{t_i} \in \mathbb{R}^{4(h+2)+d}$.

Directed Subgraph Modeling

The message passing model in GraIL (Teru, Denis, and Hamilton 2020) is a simple R-GCN (Schlichtkrull et al. 2018) with edge attention, which ignores separately modeling edge embedding and ignores the bidirectional communication between edges and nodes. Moreover, GraIL uses a node-to-node message passing mechanism where the relation information is only used for computing the weights for the neighboring nodes. However, the inductive relation reasoning should be entity-free (i.e., node-free in the subgraph case), where relation plays a dominant role while the entities cannot provide deterministic information during inference. Thus, the node-to-node message passing mechanism in GraIL weakens the role of relations and violates the nature of the inductive knowledge graph. To this end, we design a brand new message passing architecture to model the inductive enclosing subgraph by iteratively communicating and enhancing the edge and node embeddings (see Fig. 2). Since the edge information is mainly provided by the corresponding relation, the nodes can learn to better aggregate the relation information in the subgraph during the node-edge interactions, such that the model can learn to infer the relations between the target head and tail based on the relation

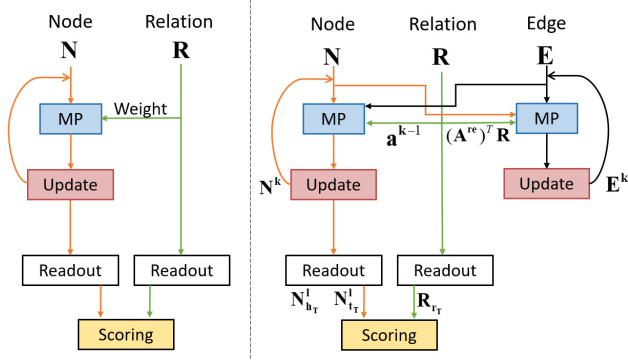


Figure 2: Comparison of Message Passing Mechanism in GraIL and Our Model. ‘MP’ in the figure denotes ‘message passing’, and the left and right figure refers to the message passing mechanism in GraIL and our model, respectively. Our method explicitly models edge embedding to strengthen the flow of the relation information and bidirectionally communicate the nodes and edges.

information presented in the subgraph.

Our node-edge interaction mechanism is inspired by CMPNN (Song et al. 2020). Nevertheless, during the node-edge interactions of CMPNN, the representation of each edge is updated by the head node embedding and its inverse edge (which is not available in our task), but neglecting the tail node embedding and the relation embedding. In our framework, we introduce a new node-edge communication mechanism that considers both head, relation, and tail to update the edge embedding. We describe our message passing model in detail in the following sections.

Firstly, the node and edge representations are mapped to the same dimensionality d using the following equations (d is also the dimensionality of the relation embedding):

$$\mathbf{N}^0 = f_1(\mathbf{N}\mathbf{W}_n^0), \mathbf{E}^0 = f_1(\mathbf{E}\mathbf{W}_e^0) \quad (1)$$

where f_1 denotes the nonlinear activation function to increase the nonlinear expressive power of the model, $\mathbf{W}_n^0 \in \mathbb{R}^{2(h+2) \times d}$ and $\mathbf{W}_e^0 \in \mathbb{R}^{(4(h+2)+d) \times d}$ are the learnable parametric matrix, $\mathbf{N}^0 \in \mathbb{R}^{N_n \times d}$ and $\mathbf{E}^0 \in \mathbb{R}^{N_e \times d}$ is the transformed node and edge embedding, respectively.

Node Embedding Updating: The node embedding is updated for totally l iterations. At each iteration (we take iteration k as an example in the following equations), edge embedding is required for updating the node embedding in our node-edge interaction mechanism. Firstly, to highlight the edges that are highly related to the target triplet, we design an **enhanced edge attention**. In the edge attention of GraIL (Teru, Denis, and Hamilton 2020), only target relation is utilized for predicting the importance of edges. In contrast, we utilize all the target head, target relation, and target tail to highlight the edges that have a close connection to the target triplet, which is more comprehensive. We use the whole triplet to conduct attention because the nodes can aggregate the relation information during node-edge interactions and thereby the updated node embeddings are also informative. For edge i ($1 \leq i \leq N_e$), the equations for the enhanced

edge attention are presented as follows:

$$\mathbf{a}_i^{k-1} = (\mathbf{N}_{h_i}^{k-1} + \mathbf{R}_{r_i} - \mathbf{N}_{t_i}^{k-1}) \oplus (\mathbf{N}_{h_T}^{k-1} + \mathbf{R}_{r_T} - \mathbf{N}_{t_T}^{k-1}) \quad (2)$$

$$a_i^{k-1} = \sigma(f_1(\mathbf{a}_i^{k-1} \mathbf{W}_{a_1}^{k-1}) \mathbf{W}_{a_2}^{k-1}) \quad (3)$$

$$\mathbf{E}_i^{(k-1)a} = a_i^{k-1} \mathbf{E}_i^{k-1} \quad (4)$$

where $\mathbf{E}_i^{(k-1)a}$ is the attentive embedding for edge i , a_i^{k-1} is a scalar that represents the weight for edge i , \oplus denotes feature concatenation, $\mathbf{N}_{h_T}^{k-1} + \mathbf{R}_{r_T} - \mathbf{N}_{t_T}^{k-1}$ denotes the embedding for target triplet, and we use $\mathbf{N}_{h_i}^{k-1} + \mathbf{R}_{r_i} - \mathbf{N}_{t_i}^{k-1}$ to denote edge i so as to be consistent with the representation of the target triplet.

Then we use the attentive edge embedding to update the node representation:

$$\mathbf{N}_{agg}^k = \mathbf{A}^{te} \mathbf{E}^{(k-1)a} \quad (5)$$

$$\mathbf{N}^k = f_1((\mathbf{N}_{agg}^k + \mathbf{N}^{k-1}) \mathbf{W}_n^k) \quad (6)$$

where \mathbf{N}_{agg}^k denotes the node aggregation information, $\mathbf{W}_n^k \in \mathbb{R}^{d \times d}$ represents the parametric matrix for node embedding at iteration k , $\mathbf{E}^{(k-1)a}$ is the attentive edge embedding, \mathbf{A}^{te} is the tail-to-edge adjacency matrix that connects each edge to its tail. Note that to preserve the directed nature of our model, for an edge (h_i, r_i, t_i) , its embedding is only used to update the tail t_i but not h_i . Coupling with the directed subgraph, it ensures that the information only flows from target head to target tail and the reverse flow is forbidden. By using the edge embedding to update node embedding, the target tail embedding can aggregate all the relations along with their relative positions occurred from the target head to target tail in the subgraph (the relative positions are provided by the node embedding), which provides powerful relational inference ability.

In the last iteration of the node embedding updating, similar to CMPNN (Song et al. 2020), we use a multi-layer perception network followed by a Gated Recurrent Unit (GRU) (Cho et al. 2014) to replace Eq. 6 to increase the expressive power of the network, as shown in the following equations:

$$\mathbf{N}^l = \text{CommunicationMLP}(\mathbf{N}_{agg}^l \oplus \mathbf{N}^{l-1} \oplus \mathbf{N}^0) \quad (7)$$

$$\mathbf{N}^l = \text{GRU}(\mathbf{N}^l) \quad (8)$$

where CommunicationMLP is the multi-layer perception network that communicates the node aggregation information \mathbf{N}_{agg}^l , node embedding \mathbf{N}^{l-1} , and the original transformed node embedding \mathbf{N}^0 (we add \mathbf{N}^0 to perform residual learning (He et al. 2016)). Note that GRUs requires the inputs to be ordered, while our node embeddings are inherently partly ordered since the nodes are arranged according to the distance to target head in the ascending order.

Edge Embedding Updating: The edge embedding is updated for totally $l - 1$ iterations. To update the edge embedding, node embedding is required in our node-edge interaction mechanism. We define inverse mappings from node to edge and relation to edge, which are denoted as:

$$\mathbf{E}_{agg}^k = (\mathbf{A}^{he})^T \mathbf{N}^k + (\mathbf{A}^{re})^T \mathbf{R} - (\mathbf{A}^{te})^T \mathbf{N}^k \quad (9)$$

where T denotes matrix transpose, $(A^{he})^T N^k$ aggregates the head information to edge, $(A^{re})^T R$ aggregates the relation information to edge, and $(A^{te})^T N^k$ aggregates the tail information to edge. The definition of edge aggregation information E_{agg}^k in Eq. 9 meets the directed requirement and is consistent across the model. Then we use the aggregation information to update edge representation:

$$E^{k'} = f_1(E^{k-1} + f_2(E_{agg}^k)) \quad (10)$$

$$E^k = f_1(E^{k'} W_e^k + E^0) \quad (11)$$

where f_1 and f_2 denotes nonlinear activation function to increase the nonlinear modeling capacity of the model. We add the E^0 to update the edge embedding in Eq. 11 to perform residual learning. Edge dropout is also performed on E^k .

Scoring Function Definition: GraIL designed an asymmetric scoring function for subgraph inductive learning by concatenating four related vectors:

$$S = W(h_G^l \oplus h_{h_T}^l \oplus h_{t_T}^l \oplus e_{r_T}) \quad (12)$$

where h_G^l denotes the subgraph representation, $h_{h_T}^l$ and $h_{t_T}^l$ denote the hidden vectors of head and tail entities, e_{r_T} is a learned embedding of the target relation. This scoring function is symmetric as both the relation embedding and subgraph embedding are undirected. To alleviate this problem, we adopt the idea of TransE (Bordes et al. 2013) to design the scoring function so as to preserve directed nature of our model as well as to be consistent with the definition of the edge information. The scoring function is defined as:

$$S = f_2(N_{h_T}^l + R_{r_T} - N_{t_T}^l) \quad (13)$$

where $N_{h_T}^l$, R_{r_T} , and $N_{t_T}^l$ denotes the final representation of target head, target relation, and target tail, respectively. We than use a two-layer fully-connected network on S to infer the score of the target triplet (h_T, r_T, t_T) .

Experiments

In our experiments, we aim to answer the following questions: 1) Does our CoMPILE outperforms state-of-the-art methods on the commonly used datasets? 2) Does our proposed message passing network better than the RGCN + edge attention in GraIL? 3) Does the directed subgraph outperform the undirected one? 4) What are the importance of the components in CoMPILE? 5) Can CoMPILE deals with asymmetric relations better than the other methods?

Datasets

WN18RR (Dettmers et al. 2017), FB15k-237 (Toutanova et al. 2015), and Nell-995 (Xiong, Hoang, and Wang 2017) are commonly used datasets that are originally developed for transductive relation prediction. Teru et al. (Teru, Denis, and Hamilton 2020) extracts four versions of inductive datasets for each dataset. Each inductive dataset constitutes of train and test graphs, where the test graph contains entities that are not presented in train graph. The subgraph for each triplet is extracted from the train or test graph.

Our Post-processed Datasets: In our experiment, we found that the inductive datasets constructed by GraIL (Teru,

Denis, and Hamilton 2020) have many empty-subgraph triplets (no valid edge exists in the enclosing subgraph of target head and tail under hop h), especially for WN18RR. Moreover, since in GraIL, the negative triplets are randomly sampled, leading to a significant number of empty-subgraphs in negative triplets. This results in inaccurate evaluation of the performance, for the reason that under the subgraph reasoning structure, it is almost impossible to infer the relation between two entities if there are no valid edges in the subgraph (these triplets also cannot work in rule-based structures such as RuleN (Meilicke et al. 2018)).

Therefore, we construct new inductive datasets by filtering out the triplets in the original inductive datasets that have no subgraph under hop h and constructing negative triplets that have subgraphs between the fake target head and fake target tail. Specifically, we extract three versions of inductive datasets each for FB15k-237 and NELL-995 datasets. Since there are no enough non-empty negative triplets for each positive triplet to perform Hits@10 experiment in the original inductive WN18RR datasets, we only extract the inductive datasets for NELL-995 and FB15k-237. The statistics of the datasets is shown in Appendix. Note that the subgraph here refers to undirected enclosing subgraph extracted in GraIL. An undirected subgraph is existed for two entities does not necessarily means that a directed subgraph is existed. If a directed subgraph does not existed, we will extract the undirected subgraph for these two entities.

Experimental Details

To be consistent with the prior methods, we use AUC-PR and Hits@10 to evaluate the models. Similar to GraIL (Teru, Denis, and Hamilton 2020), to compute AUC-PR, we sample one negative triplet for each test triplet and evaluate which triplet has larger score. For Hits@10, we compare the true triplet with the sampled negative triplets in terms of the scores, to see whether the true triplet can rank the top 10. The negative triplets are obtained by replacing the head or tail of the test triplets with other entities. For the original inductive datasets, the negative triplets are randomly sampled and do not consider whether they have enclosing subgraph. For our extracted inductive datasets, we ensure that the negative triplets can also have an enclosing subgraph.

We implement our model on Pytorch. We use Adam (Kingma and Ba 2015) as optimizer with learning rate being 0.001. The hop number h is set to 3 which is consistent with GraIL. We train the model for four times and average the testing results to obtain the final performance. The number of iterations l is set to 3. For more details, please refer to our codes at: https://github.com/TmacMai/CoMPILE_Inductive_Knowledge_Graph.

We evaluate our model with the following baselines: Neural-LP (Yang, Yang, and Cohen 2017), DRUM (Sadeghian et al. 2019), RuleN (Meilicke et al. 2018), and GraIL (Teru, Denis, and Hamilton 2020). For the detailed introduction of these methods, please refer to the introduction and related work sections.

Methods	WN18RR				FB15k-237				NELL-995			
	v1	v2	v3	v4	v1	v2	v3	v4	v1	v2	v3	v4
Neural-LP	86.02	83.78	62.90	82.06	69.64	76.55	73.95	75.74	64.66	83.61	87.58	85.69
DRUM	86.02	84.05	63.20	82.06	69.71	76.44	74.03	76.20	59.86	83.99	87.71	85.94
RuleN	90.26	89.01	76.46	85.75	75.24	88.70	91.24	91.79	84.99	88.40	87.20	80.52
GraIL	94.32	94.18	85.80	92.72	84.69	90.57	91.68	94.46	86.05	92.62	93.34	87.50
CoMPILE	98.23	99.56	93.60	99.80	85.50	91.68	93.12	94.90	80.16	95.88	96.08	85.48

Table 1: Compared with Baselines on the Original Inductive Datasets (AUC-PR). The best performance is highlighted.

Methods	WN18RR				FB15k-237				NELL-995			
	v1	v2	v3	v4	v1	v2	v3	v4	v1	v2	v3	v4
Neural-LP	74.37	68.93	46.18	67.13	52.92	58.94	52.90	55.88	40.78	78.73	82.71	80.58
DRUM	74.37	68.93	46.18	67.13	52.92	58.73	52.90	55.88	19.42	78.55	82.71	80.58
RuleN	80.85	78.23	53.39	71.59	49.76	77.82	87.69	85.60	53.50	81.75	77.26	61.35
GraIL	82.45	78.68	58.43	73.41	64.15	81.80	82.83	89.29	59.50	93.25	91.41	73.19
CoMPILE	83.60	79.82	60.69	75.49	67.64	82.98	84.67	87.44	58.38	93.87	92.77	75.19

Table 2: Compared with Baselines on the Original Inductive Datasets (Hits@10).

Results and Discussions

Compare with Baselines on Original Inductive Datasets

In this section, we compare our proposed message passing network with other baselines on the original inductive datasets proposed by GraIL. Note that this experiment is performed to evaluate the effectiveness of our node-edge message passing network, and the other settings such as the enclosing subgraph, loss function, and node embedding initialization are the same as those in GraIL. As presented in Table 1 and Table 2, our CoMPILE demonstrates improvement on the majority of the original inductive datasets in terms of both the AUC-PR and Hits@10 evaluation metrics. Specifically, the improvement is significant on the AUC-PR for WN18RR inductive datasets, on which our CoMPILE outperforms SOTA method GraIL by a large margin. The results indicate that the proposed message passing network shows marked improvement over the R-GCN with attention in GraIL, which highlights the necessity of the bidirectional communication between the node and edge embeddings as well as the effectiveness of strengthening the role of relation information in the subgraph modeling.

Compare with Baseline on Our Inductive Datasets

As demonstrated in the Dataset section, the original inductive dataset contains empty-subgraph triplets under hop h which leads to inaccurate evaluation of the performance. Therefore, we further evaluate CoMPILE on our post-processed inductive datasets that have filtered out the empty-subgraph triplets. For comparison, we implement and evaluate the R-GCN with edge attention which is the message passing method in GraIL, and the results are presented in Table 3. Note that the enclosing subgraphs are directed in both the models. As shown in Table 3, CoMPILE still performs better than the GraIL on the majority of evaluation metrics in our extracted inductive datasets, which further suggests the superiority of the message passing method in CoMPILE. For **model complexity**, the results suggest that our CoMPILE has a total number of 35,969 parameters, while the num-

ber of parameters of GraIL is 29,264. CoMPILE has slightly more parameters than GraIL, mainly because we model both the edge and node embeddings to enhance the flow of relation information in our CoMPILE.

Directed Subgraph vs. Undirected Subgraph

In this section, we investigate whether the performance has improved when the undirected enclosing subgraph is replaced by the directed one. From the results in Table 4 we can infer that the version of CoMPILE with directed subgraph significantly outperforms the version of CoMPILE with undirected subgraph almost in all the inductive datasets and evaluation metrics, demonstrating the necessity of effectively handling the direction problem in knowledge graph. This is in line with our expectations as both the FB15k-237 and NELL-995 contain a significant number of asymmetric and anti-symmetric relations (Wang et al. 2018). It is also worth mentioning that dealing with the directed subgraph is more efficient in terms of the **time complexity**, for the reason that the directed subgraph can be viewed as the cropped version of the undirected one. Although being more efficient, our directed subgraph still outperforms undirected subgraph in terms of AUC-PR and Hits@10 metrics, indicating that a directed subgraph is a better choice to infer the relation between two unseen entities.

Ablation Studies

We perform ablation studies to investigate the contributions of different components in our architecture. For the enhanced edge attention, we remove it from our model to analyze its contribution (see the case of ‘W/O EEA’ in Table 5). Moreover, we compare our enhanced edge attention with the edge attention in GraIL which only uses the target relation to conduct attention (see the case of ‘Edge Attention’). Furthermore, we investigate whether the updating of the edge embedding gives rise to the overall performance (see the case of ‘W/O EEU’). Additionally, we also remove the relation information in the edge embedding updating to analyze the importance of relation information (see the case of ‘W/O Relation in EEU’, and Eq. 9 will become

	FB15k-237						NELL-995					
	v1		v2		v3		v1		v2		v3	
	AUC-PR	Hits@10	AUC-PR	Hits@10	AUC-PR	Hits@10	AUC-PR	Hits@10	AUC-PR	Hits@10	AUC-PR	Hits@10
GraIL	78.84	61.25	80.76	64.93	82.25	68.13	68.47	58.63	84.07	76.52	81.42	75.12
CoMPILE	78.06	63.96	81.09	68.17	83.21	71.67	68.13	61.38	85.67	76.11	82.88	76.83

Table 3: Compared with Baselines on Our Extracted Inductive Datasets.

	FB15k-237						NELL-995					
	v1		v2		v3		v1		v2		v3	
	AUC-PR	Hits@10	AUC-PR	Hits@10	AUC-PR	Hits@10	AUC-PR	Hits@10	AUC-PR	Hits@10	AUC-PR	Hits@10
Undirected	77.80	59.79	80.89	65.98	81.45	68.82	63.25	58.88	79.33	68.92	83.73	76.66
Directed	78.06	63.96	81.09	68.17	83.21	71.67	68.13	61.38	85.67	76.11	82.88	76.83

Table 4: Compared between the Undirected and Directed Subgraphs on CoMPILE.

	v1		v2	
	AUC-PR	Hits@10	AUC-PR	Hits@10
W/O EEA	75.79	61.35	79.26	67.16
Edge Attention	77.07	61.04	78.53	66.46
W/O EEU	76.69	60.63	79.64	62.44
W/O Relation in EEU	77.74	60.73	79.25	62.60
CoMPILE	78.06	63.96	81.09	68.17

Table 5: Ablation Studies on Inductive FB15k-237 datasets. The ‘EEA’ and ‘EEU’ refers to Enhanced Edge Attention and Edge Embedding Updating, respectively.

$$E_{agg}^k = (A^{he})^T N^k - (A^{te})^T N^k \text{ in such circumstance.}$$

From Table 5 we find that CoMPILE significantly outperforms the version of CoMPILE whose enhanced edge attention is removed, demonstrating the effectiveness of our enhanced edge attention. Specifically, enhanced edge attention yields over 2% improvement on both AUC-PR and Hits@10 metrics of the inductive FB15k-237-v1 dataset. In contrast, the improvement of the edge attention in GraIL is relatively small, and even the model without attention outperforms the model with edge attention on inductive FB15k-237-v2 dataset. These results suggest that it is of great significance to consider all the information in target triplet to determine which edges are considered to be important.

From the result of ‘W/O EEU’ in Table 5, we can notice that the performance drops significantly on all the evaluation metrics when the edge embedding updating component is removed. Compared to completely removing the edge embedding updating module, removing the relation information in the edge embedding updating only yields slight improvement and is still clearly weaker than the case that the relation information is presented. These results demonstrate that the edge embedding updating module is a critical component that contributes to the marked improvement of our CoMPILE, and the relation information plays a very important role in the edge embedding updating. It further proves our claim that the relation is critical in inductive setting and highlights the necessity of strengthening the flow of relation information in the subgraph modeling.

Case Study for Asymmetric Relations To demonstrate that our model can naturally handle asymmetric/anti-symmetric relations to some extent (by using the specific scoring function and edge definition, etc) even provided with

the undirected enclosing subgraph, we select five relations that are asymmetric to evaluate our CoMPILE and GraIL. We use two negative triplets sampling strategies, where the first one is the standard operation that replaces the head or tail of the true triplets with other entities, and the second one is to exchange the head and tail of the test triplets. The results are presented in Table 6. Interestingly, we observe that the AUC scores of CoMPILE have no significant difference in this two negative triplets sampling strategies, indicating that CoMPILE can effectively distinguish the false triplet (t, r, h) from the true triplet (h, r, t) . In contrast, the performance of GraIL drops significantly when the sample strategy changes to the second one, which suggests that it cannot well distinguish these two triplets, mainly for the reason that the GraIL does not well handle the direction problem in KG.

	CoMPILE		GraIL	
	Standard	Exchanging h&t	Standard	Exchanging h&t
AUC	67.11	67.11	61.78	52.89
AUC-PR	63.75	61.53	64.44	56.17

Table 6: Case Study for Asymmetric Relations in Inductive FB15k-237-v1 dataset. ‘Standard’ denotes standard operation to sample negative triplets, and ‘Exchanging h&t’ denotes obtaining negative triplets by exchanging the head and tail of the test triplets.

Conclusions

We propose CoMPILE, a directed subgraph reasoning method for inductive relation prediction. We propose to use directed subgraph to infer the relation of two unseen entities which can handle asymmetric relations. Then we introduce a new message passing model that brings in the idea of enhanced edge attention and edge-node interactions to learn better node and edge embeddings and strengthen the role of relation information. Finally, we extract new inductive datasets by filtering out the triples that have no subgraph so as to evaluate each method more accurately. Experiments suggest that our CoMPILE achieves remarkable performance and outperforms state-of-the-art methods on the majority of evaluation metrics. In future work, we aim to develop a relation-only message passing network that completely discards the nodes in the enclosing subgraph and investigate the performance of relation enclosing subgraph.

Acknowledgments

This work was supported in part by the National Natural Science Foundation of China (62076262, 61673402, 61273270, 60802069), in part by the Natural Science Foundation of Guangdong Province (2017A030311029), in part by the Project of Guangzhou Science and Technology under Grant 201605122151511, in part by the National Key R&D Program of China (2020YFB020000), in part by the National Natural Science Foundation of China (61772566, 62041209), in part by the Guangdong Key Field R&D Plan (2019B020228001 and 2018B010109006), and in part by the Guangzhou S&T Research Plan (202007030010).

References

- Bordes, A.; Usunier, N.; Garcia-Duran, A.; Weston, J.; and Yakhnenko, O. 2013. Translating embeddings for modeling multi-relational data. In *Advances in neural information processing systems*, 2787–2795.
- Cho, K.; Van Merriënboer, B.; Gulcehre, C.; Bahdanau, D.; Bougares, F.; Schwenk, H.; and Bengio, Y. 2014. Learning Phrase Representations using RNN Encoder-Decoder for Statistical Machine Translation. In *EMNLP*, 1724–1734.
- Dettmers, T.; Minervini, P.; Stenetorp, P.; and Riedel, S. 2017. Convolutional 2d knowledge graph embeddings. *arXiv preprint arXiv:1707.01476*.
- Galárraga, L.; Teflioudi, C.; Hose, K.; and Suchanek, F. M. 2015. Fast rule mining in ontological knowledge bases with AMIE. *The VLDB Journal* 24(6): 707–730.
- Hamaguchi, T.; Oiwa, H.; Shimbo, M.; and Matsumoto, Y. 2017. Knowledge transfer for out-of-knowledge-base entities: A graph neural network approach. *arXiv preprint arXiv:1706.05674*.
- He, H.; Balakrishnan, A.; Eric, M.; and Liang, P. 2017. Learning symmetric collaborative dialogue agents with dynamic knowledge graph embeddings. *arXiv preprint arXiv:1704.07130*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep Residual Learning for Image Recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, 770–778.
- Kingma, D. P.; and Ba, J. 2015. Adam: A Method for Stochastic Optimization. In *ICLR*.
- Meilicke, C.; Fink, M.; Wang, Y.; Ruffinelli, D.; Gemulla, R.; and Stuckenschmidt, H. 2018. Fine-grained evaluation of rule-and embedding-based systems for knowledge graph completion. In *International Semantic Web Conference*, 3–20. Springer.
- Nathani, D.; Chauhan, J.; Sharma, C.; and Kaul, M. 2019. Learning attention-based embeddings for relation prediction in knowledge graphs. *arXiv preprint arXiv:1906.01195*.
- Nickel, M.; Rosasco, L.; and Poggio, T. 2015. Holographic embeddings of knowledge graphs. *arXiv preprint arXiv:1510.04935*.
- Sadeghian, A.; Armandpour, M.; Ding, P.; and Wang, D. Z. 2019. DRUM: End-To-End Differentiable Rule Mining On Knowledge Graphs. In *Advances in Neural Information Processing Systems*, 15347–15357.
- Schlichtkrull, M.; Kipf, T. N.; Bloem, P.; Van Den Berg, R.; Titov, I.; and Welling, M. 2018. Modeling relational data with graph convolutional networks. In *European Semantic Web Conference*, 593–607. Springer.
- Shang, C.; Tang, Y.; Huang, J.; Bi, J.; He, X.; and Zhou, B. 2019. End-to-end structure-aware convolutional networks for knowledge base completion. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 3060–3067.
- Shi, B.; and Wenginger, T. 2017. Open-world knowledge graph completion. *arXiv preprint arXiv:1711.03438*.
- Song, Y.; Zheng, S.; Niu, Z.; Fu, Z.; Lu, Y.; and Yang, Y. 2020. Communicative Representation Learning on Attributed Molecular Graphs. In *IJCAI*.
- Teru, K. K.; Denis, E.; and Hamilton, W. L. 2020. Inductive Relation Prediction by Subgraph Reasoning. In *ICML*.
- Toutanova, K.; Chen, D.; Pantel, P.; Poon, H.; Choudhury, P.; and Gamon, M. 2015. Representing Text for Joint Embedding of Text and Knowledge Bases. In *EMNLP*.
- Trivedi, R.; Dai, H.; Wang, Y.; and Song, L. 2017. Know-evolve: Deep temporal reasoning for dynamic knowledge graphs. *arXiv preprint arXiv:1705.05742*.
- Trouillon, T.; Dance, C. R.; Gaussier, É.; Welbl, J.; Riedel, S.; and Bouchard, G. 2017. Knowledge graph completion via complex tensor factorization. *The Journal of Machine Learning Research* 18(1): 4735–4772.
- Wang, P.; Han, J.; Li, C.; and Pan, R. 2019. Logic attention based neighborhood aggregation for inductive knowledge graph embedding. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 7152–7159.
- Wang, Y.; Ruffinelli, D.; Gemulla, R.; Broscheit, S.; and Meilicke, C. 2018. On evaluating embedding models for knowledge base completion. *arXiv preprint arXiv:1810.07180*.
- Wang, Z.; Zhang, J.; Feng, J.; and Chen, Z. 2014. Knowledge graph and text jointly embedding. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, 1591–1601.
- Xie, R.; Liu, Z.; Jia, J.; Luan, H.; and Sun, M. 2016a. Representation learning of knowledge graphs with entity descriptions. In *Thirtieth AAAI Conference on Artificial Intelligence*.
- Xie, R.; Liu, Z.; Luan, H.; and Sun, M. 2016b. Image-embodied knowledge representation learning. *arXiv preprint arXiv:1609.07028*.
- Xiong, C.; Power, R.; and Callan, J. 2017. Explicit semantic ranking for academic search via knowledge graph embedding. In *Proceedings of the 26th international conference on world wide web*, 1271–1279.
- Xiong, W.; Hoang, T.; and Wang, W. Y. 2017. DeepPath: A Reinforcement Learning Method for Knowledge Graph Reasoning. In *EMNLP*.

Xu, K.; Hu, W.; Leskovec, J.; and Jegelka, S. 2018. How powerful are graph neural networks? *arXiv preprint arXiv:1810.00826* .

Yang, B.; Yih, W.-t.; He, X.; Gao, J.; and Deng, L. 2014. Embedding entities and relations for learning and inference in knowledge bases. *arXiv preprint arXiv:1412.6575* .

Yang, F.; Yang, Z.; and Cohen, W. W. 2017. Differentiable learning of logical rules for knowledge base reasoning. In *Advances in Neural Information Processing Systems*, 2319–2328.

Zhang, M.; and Chen, Y. 2018. Link prediction based on graph neural networks. In *Advances in Neural Information Processing Systems*, 5165–5175.

Zhang, Y.; Dai, H.; Kozareva, Z.; Smola, A. J.; and Song, L. 2017. Variational reasoning for question answering with knowledge graph. *arXiv preprint arXiv:1709.04071* .

Zhong, H.; Zhang, J.; Wang, Z.; Wan, H.; and Chen, Z. 2015. Aligning knowledge and text embeddings by entity descriptions. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, 267–272.