

PipeNet: Question Answering with Semantic Pruning over Knowledge Graphs

Ying Su, Jipeng Zhang, Yangqiu Song, Tong Zhang

Hong Kong University of Science and Technology

{ysuay, jzhanggr}@connect.ust.hk, yqsong@cse.ust.hk,
tongzhang@ust.hk

Abstract

It is well acknowledged that incorporating explicit knowledge graphs (KGs) can benefit question answering. Existing approaches typically follow a grounding-reasoning pipeline in which entity nodes are first grounded for the query (question and candidate answers), and then a reasoning module reasons over the matched multi-hop subgraph for answer prediction. Although the pipeline largely alleviates the issue of extracting essential information from giant KGs, efficiency is still an open challenge when scaling up hops in grounding the subgraphs. In this paper, we target at finding semantically related entity nodes in the subgraph to improve the efficiency of graph reasoning with KG. We propose a grounding-pruning-reasoning pipeline to prune noisy nodes, remarkably reducing the computation cost and memory usage while also obtaining decent subgraph representation. In detail, the pruning module first scores concept nodes based on the dependency distance between matched spans and then prunes the nodes according to score ranks. To facilitate the evaluation of pruned subgraphs, we also propose a graph attention network (GAT) based module to reason with the subgraph data. Experimental results on CommonsenseQA and OpenBookQA demonstrate the effectiveness of our method.

1 Introduction

Question answering requires related background knowledge. A line of research resorts to combining pre-trained language models (LMs) and knowledge graphs (KG) to utilize both the implicit knowledge in LMs and explicit knowledge in structured KGs (Schlichtkrull et al., 2018; Lin et al., 2019; Feng et al., 2020; Yasunaga et al., 2021).

The researches towards utilizing knowledge from KGs typically follow a grounding-and-reasoning pipeline, namely schema graph grounding and schema graph reasoning (Lin et al., 2019).

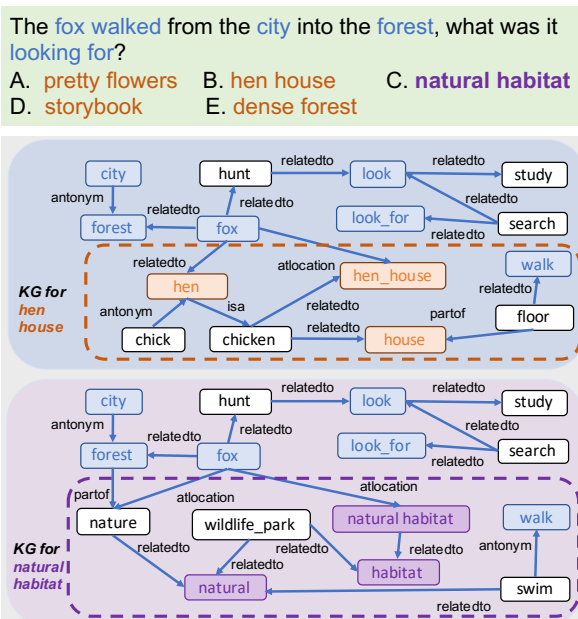


Figure 1: An example of a query and grounded knowledge graph for two candidate answers. The external KG nodes are more diverse around the answer concepts than question concepts.

In the grounding module, multi-hop neighbors of matched concept nodes in the query from KG form a subgraph. Recent works focus on improving reasoning ability by enhancing the representation of multi-hop nodes in grounded subgraphs with graph neural networks (GNNs) (Feng et al., 2020; Yasunaga et al., 2021) or interaction between representations of query context and subgraphs (Zhang et al., 2022b; Sun et al., 2022). While pre-trained LMs are powerful at extracting plain text features for the query context, the quality of subgraph feature extracted from GNNs is still prone to noisy nodes in grounded subgraphs. Specifically, there are two challenges in fusing KGs with GNNs. First, the computation and memory cost would increase with the hops increase. Second, the noisy nodes induced with increasing hops deteriorate the quality of the subgraph feature, and further decrease the performance of the reasoning module.

In this paper, we tackle the problems brought by noisy nodes with a grounding-pruning-reasoning pipeline framework, PipeNet. Previous researches show that the improvement of graph-based reasoning systems is minor, though with the number of grounded hops increasing, many more new nodes are induced (Santoro et al., 2017; Wang et al., 2019a; Feng et al., 2020). As shown in Figure 1, many of them are the same for different candidate answers, especially near the question concepts. Diverse nodes are mainly brought in due to the difference in answer concepts. This diversity is critical to the subgraph representation learning with GNNs.

Our pruning module prunes noisy nodes before the reasoning module, reducing the computation cost and memory usage while keeping the diversity of subgraphs in the meantime. Specifically, we propose a dependency structure based pruning method to prune the nodes with dependency parsing (DP) tools. The DP-pruning strategy is inspired by relation extraction in automatic ontology building, in which the dependency tree is applied to find possible relations between concepts according to the distance on the tree (Fellbaum and Miller, 1998; Sombatsrisomboon et al., 2003; Ciaramita et al., 2005; Kang et al., 2015). Similarly, we assume the dependency tree provides reasonable linguistic links between grounded concepts in a natural language context. We further convert the dependency distances between grounded concepts into concept node scores and propagate the node scores onto the grounded multi-hop subgraph to prune external noisy nodes.

To facilitate the evaluation of pruned subgraph, we also propose a simplified version of GAT (Veličković et al., 2018) for graph representation learning. We redesign the message passing mechanism in (Yasunaga et al., 2021). Our contributions are as follows:

- We propose a grounding-pruning-reasoning pipeline PipeNet for question answering with KG, in which a DP-pruning module improves efficiency by pruning the noisy nodes.
- We propose a simplified GAT module for fusing KG with GNNs. The module simplifies the message flow while achieving comparable or higher performance in the meantime;

Experiments on two standard benchmarks, CommonsenseQA (Talmor et al., 2019) and Open-

bookQA (Mihaylov et al., 2018), demonstrate the effectiveness of our proposed method. The code is open-sourced¹.

2 Related Work

2.1 QA with LM+KG

With the development of benchmarking question answering, more and more hard question answering datasets are developed, which require background knowledge to solve (Mihaylov et al., 2018; Talmor et al., 2019, 2021). Pretrained LMs and KGs are commonly used knowledge sources, research typically adopts an LM+KG framework as to acquire relevant knowledge for commonsense QA (Feng et al., 2020; Yasunaga et al., 2021; Zhang et al., 2022b; Su et al., 2022; Park et al., 2023; Huang et al., 2023; Ye et al., 2023; Wang et al., 2023; Taunk et al., 2023; Zhao et al., 2023; Dong et al., 2023; Mazumder and Liu; Kang et al., 2024; Zhao et al., 2024)

Schlichtkrull et al. (2018) first adopts RGCN to model relational data in KG, which specifically models the node representation as the aggregation from neighboring nodes. GconAttn (Wang et al., 2019a) adds inter-attention between the concepts in premise and hypothesis to find the best-aligned concepts between the respective graphs. KagNet (Lin et al., 2019) further proposes an LSTM-based path encoder to model knowledge paths in the schema graph on top of GCNs. RN (Santoro et al., 2017) uses MLPs to encode the one-hop paths and pooling over the path embedding to get the schema graph representation. MHGRN (Feng et al., 2020) stresses modeling multi-hop paths and utilized an attention mechanism to weigh the importance of multi-hop paths. QAGNN (Yasunaga et al., 2021) adopts GAT for type and relation-aware messages to update the node representations. GreaseLM (Zhang et al., 2022b) further improves the knowledge fusion quality between context and subgraph representation by adding an information fusion module.

Unlike these works, we focus on effectively finding informative subgraph nodes from the raw output of the grounding module. We adopt a pruning module to find such nodes, which benefits the subgraph representation learning from GNNs.

¹<https://github.com/HKUST-KnowComp/PipeNet>

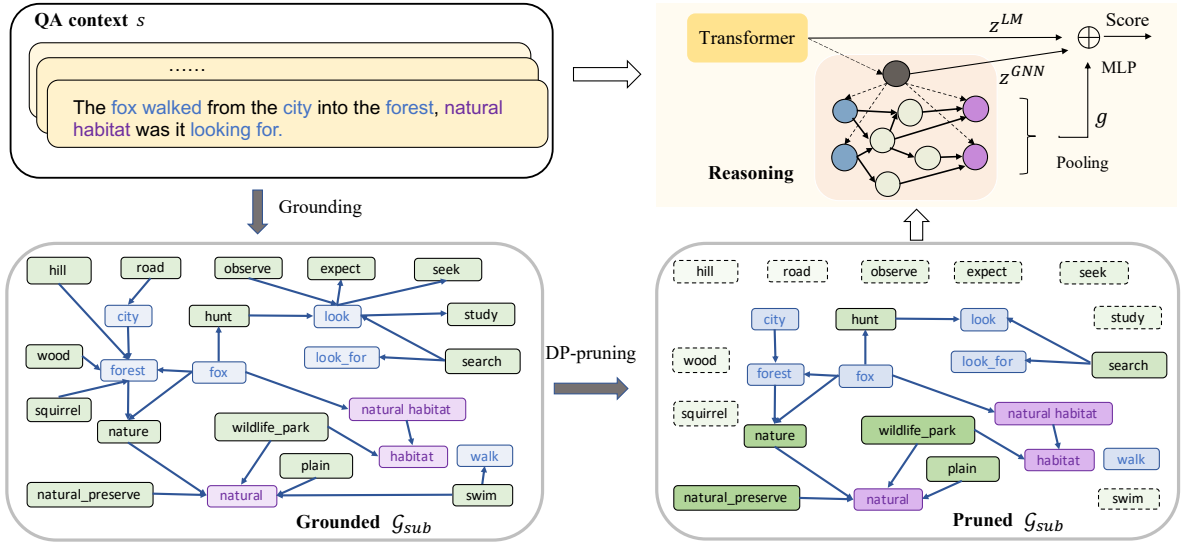


Figure 2: The overall framework of grounding-pruning-reasoning pipeline PipeNet. Concept nodes are first grounded in the KG to form a subgraph \mathcal{G}_{sub} related to question and answer context s . A pruning module prunes noisy nodes according to node score before the reasoning module. The final answer score is calculated based on the context representation z^{LM} and subgraph representation z^{GNN} .

2.2 Efficient Computation for GNN

Though the application of GNN has become popular in many graph-based scenarios, it is still challenging to apply GNN to large-scale graphs with massive numbers of nodes and edges (Hamilton et al., 2017; Yu et al., 2022; Zhang et al., 2022a) due to expensive computation cost and high memory usage. Categories of research towards tackling this problem are mainly sampling-based (Chen et al., 2018; Zeng et al., 2019; Chiang et al., 2019; Zeng et al., 2021; Fey et al., 2021) and precomputing-based (Wu et al., 2019; Rossi et al., 2020; Liu and Ji, 2022).

Previous pruning method JointLK (Sun et al., 2022) dynamically prunes noisy nodes during training, which still takes the raw output of the grounding module as inputs and does not decrease memory or computation cost. GSC (Wang et al., 2022) reduces parameters in the GNN layer by separately viewing the reasoning process as counting, which reduces model size while ignoring the semantic interaction between context and subgraph. Unlike them, we focus on extracting informative subgraph nodes of much smaller size from the grounded subgraph in a precomputing stage.

3 Methodology

Our grounding-pruning-reasoning framework, PipeNet, consists of three stages: subgraph grounding, subgraph pruning, and reasoning. The overall framework is shown in Figure 2.

3.1 Problem Formulation

Given a context query q and a set of candidate answers $\{a_1, a_2, \dots, a_k\}$, the task is to choose the most plausible answer from the set. Related background knowledge can be retrieved from a relevant KG $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ given the query and answer set. \mathcal{V} represents the set of entity nodes and \mathcal{E} represents the set of relational edges in the KG.

Following the definition in Yasunaga et al. (2021), specifically for a question q and a candidate answer a , we define the grounded concept nodes from \mathcal{G} as \mathcal{V}_q and \mathcal{V}_a respectively. The question and each answer are further composed as a QA context s . External concept nodes from \mathcal{G} during the multi-hop expansion are defined as \mathcal{V}_e . The grounded nodes and edges between them form the grounded subgraph \mathcal{G}_{sub} .

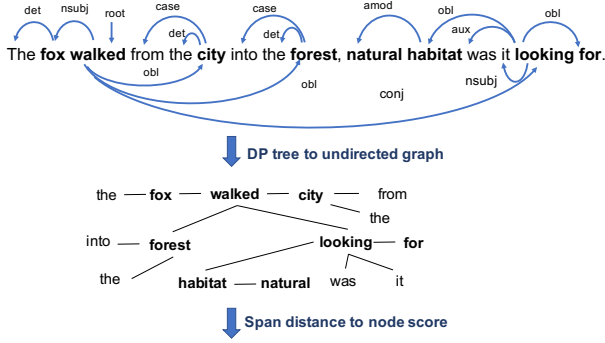
As we aim to explore the impacts of the external nodes on the learning efficiency of GNNs with KG, we define the one-hop and two-hop settings as:

One-hop. The grounded subgraph consists of entity nodes from \mathcal{V}_q and \mathcal{V}_a , and the linked edges between the nodes.

Two-hop. The grounded subgraph consists of entity nodes from \mathcal{V}_q , \mathcal{V}_a and \mathcal{V}_e , and the linked edges between the nodes. \mathcal{V}_e is the set of one-hop neighbors from \mathcal{V}_q and \mathcal{V}_a .

3.2 DP-pruning

Our DP-pruning strategy on grounded subgraphs is based on dependency links between matched spans in the QA context s . Dependency analysis



\mathcal{V}_q : {fox, walked, city, forest, looking, looking for},
 \mathcal{V}_a : {natural, habitat, natural habitat}
Span distance: $Dist(\text{fox, natural})=4$, $Dist(\text{fox, natural habitat})=3$,
 $Dist(\text{fox, habitat})=3$
Node score: $\mathcal{D}_q[\text{fox}] = -\frac{(4+3+3)}{3} = -3.33$

Figure 3: Dependency tree of a QA context example. Words in bold are matched spans of concepts in ConceptNet.

helps find relations between terms using dependency information present in parsing trees. Explicit syntax-aware knowledge has shown effective usages in downstream tasks, such as machine translation (Bastings et al., 2017; Marcheggiani et al., 2018), information extraction (Sahu et al., 2019), and semantic role labeling (Zhang et al., 2020).

DP tree and span distance. We adopt the widely used open-source tool stanza² for dependency analysis on the QA context. The dependency parsing (DP) tree \mathcal{T} is then converted into an undirected graph \mathbb{G} . On the graph, we can calculate the shortest path lengths as the span distance between span words. An example is shown in Figure 3. We align the results of concept matching and dependency parsing on the word level. If the matched span covers more than one word, the distance is calculated as the minimum distance of covered words to other spans.

Span distance to node score. As we focus on refining the matched subgraph \mathcal{G}_{sub} , we calculate the node score of matched concepts in q and a based on the corresponding span distance. For each concept c_q in \mathcal{V}_q , the node score is:

$$\mathcal{D}_q[c_q] = -\frac{\sum_{i=1}^{|\mathcal{V}_a|} Dist(c_q, c_a)}{|\mathcal{V}_a|}, \quad (1)$$

where $Dist$ is the corresponding span distance of matched concepts. For each concept c_a in \mathcal{V}_a , the node score is calculated in the same way.

Propagate node score. Our algorithm aims to

²https://stanfordnlp.github.io/stanza/corenlp_client.html

Algorithm 1 Grounding and Pruning

Require: q, a
Require: Hop n
Require: KG \mathcal{G}
Require: Prune rate p

$\mathcal{V}_q, \mathcal{V}_a, s \leftarrow q, a, \mathcal{G}$
 $\mathcal{T} \leftarrow s$
 $\mathbb{G} \leftarrow \mathcal{T}$
 $\mathcal{D}_q, \mathcal{D}_a \leftarrow \mathbb{G}$
 $i \leftarrow 1, \mathcal{V}_t \leftarrow \mathcal{V}_q \cup \mathcal{V}_a, \mathcal{D}_t \leftarrow \mathcal{D}_q \cup \mathcal{D}_a$
while $i \leq n$ **do**
 $\mathcal{V}_e \leftarrow Neighbor(\mathcal{V}_t)$
 $\mathcal{D}_e \leftarrow Avg(\mathcal{D}_t)$
 $\mathcal{V}_t \leftarrow \mathcal{V}_t \cup \mathcal{V}_e$
 $\mathcal{D}_t \leftarrow \mathcal{D}_t \cup \mathcal{D}_e$
end while
 $threshold \leftarrow \mathcal{D}_t, p$
for $v \in \mathcal{V}_t \setminus \{\mathcal{V}_q, \mathcal{V}_a\}$ **do**
 if $\mathcal{D}_t[v] \leq threshold$ **then**
 Delete v
 end if
end for
return \mathcal{V}_t

prune the external nodes \mathcal{V}_e in the subgraph for two-hop or above because noisy nodes are mainly induced with the hops growing. The pseudo-code for pruning the external nodes is listed in Algorithm 1. In initialization, grounded concept sets $\mathcal{V}_t = \mathcal{V}_q \cup \mathcal{V}_a$, and score set for grounded concept sets $\mathcal{D}_t = \mathcal{D}_q \cup \mathcal{D}_a$. External nodes having neighbors in the grounded concept sets are added to expand the grounded subgraph \mathcal{G}_{sub} . The node score of external nodes is assigned as the average of their neighbor node scores during expansion. Until the expansion hops end, the nodes except \mathcal{V}_q and \mathcal{V}_a are pruned according to their score ranks. The nodes with smaller node scores are pruned.

Our algorithm propagates the dependency structure information from QA context s onto the retrieved static subgraph \mathcal{G}_{sub} . We keep concept nodes with higher scores because they generally have closer distances to the concept nodes in \mathcal{V}_a , which increases the diversity of pruned subgraph. Finally, the $(|\mathcal{V}_t| - |\mathcal{V}_q| - |\mathcal{V}_a|) * p$ will be pruned with pruning rate p .

3.3 Reasoning

We design a reasoning module fusing the QA context feature and subgraph feature. The dimension of subgraph feature generated from L -layer GNN

Model	Time	Space
\mathcal{G} is a dense graph		
L-h KagNet	$\mathcal{O}(\mathcal{R} ^L \mathcal{V} ^{L+1} L)$	$\mathcal{O}(\mathcal{R} ^L \mathcal{V} ^{L+1} L \cdot D)$
L-h MHGRN	$\mathcal{O}(\mathcal{R} ^2 \mathcal{V} ^2 L)$	$\mathcal{O}(\mathcal{R} \mathcal{V} L \cdot D)$
L-1 QAGNN	$\mathcal{O}(\mathcal{V} ^2 L)$	$\mathcal{O}(\mathcal{R} \mathcal{V} L \cdot D)$
L-1 GreaseLM	$\mathcal{O}(\mathcal{V} ^2 L)$	$\mathcal{O}(\mathcal{R} \mathcal{V} L \cdot D)$
L-1 JointLK	$\mathcal{O}(\mathcal{V} ^2 L)$	$\mathcal{O}(\mathcal{R} \mathcal{V} L \cdot D)$
L-1 GSC	$\mathcal{O}(\mathcal{V} L)$	$\mathcal{O}(\mathcal{R} \mathcal{V} L)$
L-1 PipeNet	$\mathcal{O}((\frac{ \mathcal{V} }{k})^2 L)$	$\mathcal{O}(\mathcal{R} \frac{ \mathcal{V} }{k} L \cdot D)$
\mathcal{G} is a sparse graph with maximum node degree $\Delta \ll \mathcal{V} $		
L-h KagNet	$\mathcal{O}(\mathcal{R} ^L \mathcal{V} L \Delta^L)$	$\mathcal{O}(\mathcal{R} ^L \mathcal{V} L \Delta^L \cdot D)$
L-h MHGRN	$\mathcal{O}(\mathcal{R} ^2 \mathcal{V} L \Delta)$	$\mathcal{O}(\mathcal{R} \mathcal{V} L \cdot D)$
L-1 QAGNN	$\mathcal{O}(\mathcal{V} L \Delta)$	$\mathcal{O}(\mathcal{R} \mathcal{V} L \cdot D)$
L-1 GreaseLM	$\mathcal{O}(\mathcal{V} L \Delta)$	$\mathcal{O}(\mathcal{R} \mathcal{V} L \cdot D)$
L-1 JointLK	$\mathcal{O}(\mathcal{V} L \Delta)$	$\mathcal{O}(\mathcal{R} \mathcal{V} L \cdot D)$
L-1 GSC	$\mathcal{O}(\mathcal{V} L)$	$\mathcal{O}(\mathcal{R} \mathcal{V} L)$
L-1 PipeNet	$\mathcal{O}(\frac{ \mathcal{V} }{k} L \Delta)$	$\mathcal{O}(\mathcal{R} \frac{ \mathcal{V} }{k} L \cdot D)$

Table 1: L-h means L-hop and L-1 means L-layer. \mathcal{G} is a graph with relation set \mathcal{R} . k is the reduction rate in the PipeNet pruning stage.

is D . Theoretically, the efficiency analysis in time and space for the GNN is shown in Table 1. Note the definition of reduction rate k in the table is slightly different from the pruning rate p :

$$\frac{1}{k} = 1 - \frac{(|\mathcal{V}_t| - |\mathcal{V}_q| - |\mathcal{V}_a|)}{|\mathcal{V}_t|} * p. \quad (2)$$

For the QA context feature, the input is QA context s . A pre-trained language model first encodes the context into the vector representation z as:

$$z^{LM} = f_{enc}(s), \quad (3)$$

where z is the hidden state of [CLS] token in the last hidden layer.

Following (Yasunaga et al., 2021), the QA context is induced as an additional node to the grounded subgraph \mathcal{G}_{sub} and assigned to connect the nodes in \mathcal{V}_q and \mathcal{V}_a . The representation of this additional context node in the subgraph is initialized as z^{LM} .

For the subgraph feature, the embeddings of entity nodes in the subgraph are initialized as D -dim vectors. Similar to (Yasunaga et al., 2021; Sun et al., 2022; Zhang et al., 2022b), a standard GNN structure is applied to learn entity node representations via iterative message passing between neighbors on the subgraph. Specifically, in the $(l+1)$ -layer, the hidden state of the node on the subgraph is updated by:

$$\mathbf{h}_t^{(l+1)} = f_n\left(\sum_{s \in \mathcal{N}_t \cup \{t\}} \alpha_{st} \mathbf{m}_{st}\right), \quad (4)$$

where \mathcal{N}_t represents the neighborhood of target node t and $\mathbf{m}_{st} \in \mathbb{R}^D$ denotes the message from each neighbor node s to t . $f_n : \mathbb{R}^D \rightarrow \mathbb{R}^D$ is a 2-layer multilayer perceptron (MLP) function.

Specifically, for the message on the edge, we encode the connected node types and the edge type into embedding forms. As shown in (Wang et al., 2022), these two types of information in the subgraph are important.

$$\mathbf{r}_{st} = f([e_{st}, u_s, u_t]), \quad (5)$$

where u_s, u_t are one-hot vectors of node type and e_{st} is one-hot vector of edge type. f is a 2-layer MLP converting the concatenated feature into a D dimension edge representation. The message on the relational edges propagated from source node s to target node t is:

$$\mathbf{m}_{st} = f_m(\mathbf{h}_s^l, \mathbf{r}_{st}), \quad (6)$$

where $f_m : \mathbb{R}^{2D} \rightarrow \mathbb{R}^D$ is a linear transformation.

We adopt an attention-based message passing module based on GAT (Veličković et al., 2018). Different from (Yasunaga et al., 2021), the attention is calculated based on the node types and relation type. First, the query and key vectors are computed as:

$$\mathbf{q}_s = f_q(\mathbf{h}_s^l), \quad (7)$$

$$\mathbf{k}_t = f_k(\mathbf{h}_t^l, \mathbf{r}_{st}), \quad (8)$$

where $f_q : \mathbb{R}^D \rightarrow \mathbb{R}^D$ and $f_k : \mathbb{R}^{2D} \rightarrow \mathbb{R}^D$ are linear transformations. Finally, the attention weight α_{st} :

$$\alpha_{st} = \frac{\exp(\gamma_{st})}{\sum_{t \in N_s} \exp(\gamma_{st})}, \gamma_{st} = \frac{\mathbf{q}_s^T \mathbf{k}_t}{\sqrt{D}}. \quad (9)$$

At the final layer of the GNN network, we get the representation of the additional context node and pooled representation of KG nodes in the subgraph as z^{GNN} and \mathbf{g} .

Answer Prediction. Given question q and a candidate answer a , the plausibility score $p(a|q)$:

$$p(a|q) \propto \exp(MLP(z^{LM}, z^{GNN}, \mathbf{g})), \quad (10)$$

where an MLP layer encodes the context feature and graph feature into the final score. The answer among candidate answers with the highest plausibility score is the predicted answer.

4 Experiments

Our experiments are conducted on two standard question answering benchmarks, CommonsenseQA (CSQA) and OpenBookQA (OBQA). We also introduce details of baselines and implementations in this section.

4.1 Datasets

CommonsenseQA. CommonsenseQA (Talmor et al., 2019) is a 5-way multiple choice QA task that requires reasoning with commonsense knowledge, containing 12,102 questions which are created with entities from ConceptNet (Speer et al., 2017). Following (Lin et al., 2019), we conduct experiments on the in-house (IH) data split (8,500/1,221/1,241 for IHtrain/IHdev/IHtest respectively).

OpenBookQA. OpenBookQA (Mihaylov et al., 2018) is a 4-way multiple choice QA task, containing 5,957 questions (4,957/500/500 for train/dev/test respectively). It is an elementary science question together with an open book of science facts. Answering OpenBookQA requires commonsense knowledge beyond the provided facts.

4.2 Baselines

We use baselines for two experiments: baselines for the PipeNet framework with our designed reasoning module, and baselines for the DP-pruning.

4.2.1 Framework

We compare with other grounding-reasoning-based frameworks: (1) Relation Network (RN) (Santoro et al., 2017), (2) RGCN (Schlichtkrull et al., 2018), (3) GconAttn (Wang et al., 2019b), (4) KagNet (Lin et al., 2019), (5) MHGRN (Feng et al., 2020), (6) QA-GNN (Yasunaga et al., 2021), (7) GreaseLM (Zhang et al., 2022b).

4.2.2 Pruning

JointLK (Sun et al., 2022). JointLK automatically selects relevant nodes from noisy KGs by designing a dense bidirectional attention module to attend to the question tokens and KG nodes. A dynamic pruning module recursively prunes irrelevant KG nodes based on the attention weights.

GSC (Wang et al., 2022). GSC designs a simple graph neural model which regards the reasoning over knowledge graph as a counting process. It reduces the hidden dimension of GNN layers and

Methods	IHdev-Acc.(%)	IHtest-Acc.(%)
RoBERTa-Large	73.07 (± 0.45)	68.69 (± 0.56)
Framework		
RGCN	72.69 (± 0.19)	68.41 (± 0.66)
GconAttn	72.61 (± 0.39)	68.59 (± 0.96)
KagNet	73.47 (± 0.22)	69.01 (± 0.76)
RN	74.57 (± 0.91)	69.08 (± 0.21)
MHGRN	74.45 (± 0.10)	71.11 (± 0.81)
QA-GNN	76.54 (± 0.21)	73.41 (± 0.92)
GreaseLM	78.5 (± 0.5)	74.2 (± 0.4)
PipeNet	78.95 (± 0.55)	74.49 (± 0.26)
Pruning		
JointLK	77.88 (± 0.25)	74.43 (± 0.83)
GSC	79.11 (± 0.22)	74.48 (± 0.41)
PipeNet(DP)	78.13 (± 0.13)	74.75 (± 0.47)

Table 2: Results on the CSQA in-house split dataset. The mean and standard deviation value of three runs on the in-house Dev (IHdev) and Test (IHtest) datasets are reported. Pruning rate p is 90% in PipeNet(DP).

results in a reasoning module with a much smaller size.

For the experiments on the framework, we use the grounded two-hop knowledge subgraph. For the experiments on pruning, we conduct experiments on PipeNet with a DP-pruning strategy over two-hop subgraphs.

4.3 Implementation Details

For all the experiments on PipeNet, we set the dimension ($D = 200$) and the number of layers ($L = 5$) in the reasoning module. The parameters of the reasoning module (LM+GNN) are optimized by RAdam (Liu et al., 2019a) by cross-entropy loss. The learning rate for the LM encoder is set as $1e-5$. For the decoder with GNN, the learning rate is $1e-3$. For both benchmarks, we use ConceptNet (Speer et al., 2017) as the knowledge graph. For the pruning experiments on PipeNet, the DP-pruning strategy prunes the nodes by the ranks of node scores. Specifically, the *threshold* value is determined by the score of top $(1 - p)$ percent ranked node is $\mathcal{V}_t \setminus \{\mathcal{V}_q, \mathcal{V}_a\}$. In each experiment, we use two RTX 3090 GPUs, and the average running time is about 4 hours on CSQA and 24 hours on OBQA.

5 Results

In this section, we first present of main results of PipeNet as well as PipeNet with DP pruning strategy on standard benchmarks. Then we analyze the time and memory efficiency improvement brought by DP-pruning strategy. Finally, we conduct an ablation study over pruning strategy.

Methods	RoBERTa-large	AristoRoBERTa
w/o KG	64.80 (± 2.37)	78.40 (± 1.64)
Framework		
+RGCN	62.45 (± 1.57)	74.60 (± 2.53)
+GconAtten	64.75 (± 1.48)	71.80 (± 1.21)
+RN	65.20 (± 1.18)	75.35 (± 1.39)
+MHGRN	66.85 (± 1.19)	80.6
+QAGNN	67.80 (± 2.75)	82.77 (± 1.56)
+GreaseLM	-	84.8
+PipeNet	69.33 (± 1.60)	87.33 (± 0.19)
Pruning		
+JointLK	70.34 (± 0.75)	84.92 (± 1.07)
+GSC	70.33 (± 0.81)	86.67 (± 0.46)
+PipeNet(DP)	69.60 (± 0.47)	87.80 (± 0.43)

Table 3: Test accuracy comparison on OBQA. Methods with AristoRoBERTa (Clark et al., 2020) use the textual evidence as an additional input to the QA context. Pruning rate p is 90% in PipeNet(DP).

5.1 Accuracy of PipeNet and DP-pruning

The results on CSQA and OBQA are shown in Table 2 and 3 separately. From the results on both benchmarks, we can find PipeNet is an effective framework for combining the context feature learning and subgraph feature learning. Besides node type and edge type features, QAGNN (Yasunaga et al., 2021) also employs node embedding and relevance-score as external features. GreaseLM (Zhang et al., 2022b) stresses the modality interaction between context feature and subgraph feature. Unlike them, we adopt a simplified message flow for subgraph feature and merge the two kinds of features with an MLP layer. The final performance is comparable with previous methods on CSQA and better on OBQA. This is because that node embedding and relevance score gradually loses benefits to the reasoning module with training continuing as analyzed in GSC (Wang et al., 2022). Decreasing redundant subgraph features and modality interaction at the same time makes the reasoning module focus more on the subgraph learning, which further benefits the reasoning performance.

DP-pruning strategy can further improve the subgraph representation learning based on the PipeNet framework. Since the best answer is chosen from multiple candidate choices, DP-pruning strategy can help maintain the uniqueness of grounded subgraphs by pruning nodes which are far from the concept nodes in candidate answers. Comparing results of PipeNet and PipeNet with DP-pruning, DP-pruning can help PipeNet achieve better performances on both benchmarks under most circumstances, with a high pruning rate as 90%.

Methods	Test
RoBERTa (Liu et al., 2019b)	72.1
AristoRoBERTa (Clark et al., 2020)	77.8
AristoRoBERTa + MHGRN (Feng et al., 2020)	80.6
ALBERT (Lan et al., 2020) + KB	81.0
AristoRoBERTa + QA-GNN (Yasunaga et al., 2021)	82.8
T5 (Raffel et al., 2020)	83.2
AristoRoBERTa + GreaseLM (Zhang et al., 2022b)	84.8
AristoRoBERTa + JointLK (Sun et al., 2022)	85.6
UnifiedQA (Khashabi et al., 2020)	87.2
AristoRoBERTa + GSC (Wang et al., 2022)	87.4
GenMC (Huang et al., 2022)	89.8
AristoRoBERTa + PipeNet(DP)	88.2

Table 4: Test accuracy comparison on OBQA leaderboard. The parameter size is about 3B for T5, and 11B for UnifiedQA and GenMC. The parameter size of PipeNet is about 358M.

DP-pruning strategy also has strengths over other pruning methods like JointLK and GSC. Compared to JointLK, PipeNet significantly reduces memory and computation costs during training as shown in Table 1. Moreover, on the OBQA benchmark where additional factual texts are induced to the QA context (with AristoRoBERTa (Clark et al., 2020)³), our PipeNet outperforms GSC by 1.13% on the accuracy score. AristoRoBERTa applies several methods to encode science-related knowledge into RoBERTa. PipeNet captures the semantic feature interaction between context and subgraph with an MLP layer while GSC separately models the subgraph representation as a counting process.

Furthermore, we also compare the performance of PipeNet with other methods on the OBQA test leaderboard, and the result is listed in Table 4. Compared to the pre-trained LM T5 (Raffel et al., 2020), PipeNet achieves 5% higher accuracy with much fewer parameters. It indicates that the knowledge in external KG is complementary to the implicit knowledge in LMs. Compared to UnifiedQA (Khashabi et al., 2020) which injects the commonsense knowledge from multiple QA sources into pre-trained LMs, PipeNet achieves 1% performance gain. It shows that knowledge graph is still an important and useful knowledge source for QA. The recent method GenMC outperforms PipeNet by inducing clues for generation based on T5-large. It may be worth exploring how to employ the clues to guide the subgraph selection for better representation.

³<https://huggingface.co/LIAMF-USP/aristo-roberta>




Figure 4: Distribution of grounded nodes and edges with pruning rate on external nodes from one-hop to two-hop on CSQA training dataset.

$p(\%)$	k	$M(G)$	$\uparrow(\%)$	$t(s)$	$\uparrow(\%)$	IHtest (%)
0	1.0	5.02	-	1.16	-	74.38
10	1.1	4.95	1	1.02	13	74.21
20	1.2	4.92	2	1.01	13	74.29
30	1.3	4.75	5	0.96	17	74.29
40	1.4	4.67	7	0.87	25	74.21
50	1.6	4.57	9	0.83	28	74.13
60	2.0	4.49	11	0.78	33	74.70
70	2.4	4.22	16	0.75	35	74.85
80	3.2	3.83	24	0.72	38	74.70
90	4.8	3.51	30	0.67	42	74.86

Table 5: Results on CSQA in-house split with PipeNet. GPU memory usage and time efficiency improvement are shown for pruning rate p on two-hop subgraph for GNN during training. The training batch size is 64.

5.2 Efficiency of PipeNet and DP-pruning

In this section, we conduct empirical studies to analyze the time and memory cost of our method. Besides, a corresponding theoretical analysis of the efficiency is presented in Section 3.3. Specifically, we implemented GAT using the tool *Pytorch Geometric* (Fey and Lenssen, 2019). Figure 4 illustrates that the average number of edges is linearly decreased with the number of nodes pruned.

Our method has demonstrated better time and memory efficiency. The result of running cost and performance on CSQA is presented in Table 5. The reduction rate k is calculated based on the number of nodes and edges in Figure 4. M is the GPU memory usage (max allocation memory) of GAT module and t is average batch time of the module during training. With pruning rate p growing, k is growing non-linearly, as well as memory usage M and time t efficiency. The memory and time efficiency exhibit different growing trends. Memory efficiency becomes evident when p is greater than 60 and time efficiency becomes evident when p is greater than 40%. Performance improvement becomes evident

h -hop	Prune method	Prune rate	IHtest-Acc(%)
One	-	0	73.27 (± 0.93)
Two	-	0	74.49 (± 0.26)
Two	Random	90%	73.51 (± 0.61)
Two	DP	90%	74.75 (± 0.47)

Table 6: Results on CSQA in-house split with PipeNet.

when p is greater than 60%. Specifically, when $p=90\%$, the memory and time efficiency achieve 30% and 42% improvement separately.

We also present the performance of CSQA test split with the pruning rate changes. It turns out that the pruning strategy leads to small variance in the performance change. Generally, larger p leads to better performances. The performance improvement keeps steady when p is greater than 60%. $p=90\%$ achieves the best efficiency by only increasing the number of nodes from 12 to 23 and the number of edges from 190 to 377 for each QA context, and also better than original two-hop subgraph. Overall, the performance demonstrates that the DP-pruning strategy can find informative nodes benefiting the subgraph representation learning with a great reduction in the memory and computation cost.

5.3 Ablation Study

We conduct experiments on pruning strategy over CSQA as the ablation study. For a fair comparison, we design a random pruning strategy with the same pruning rate of 90% to DP-pruning. The pruning is also applied to the additional KG nodes \mathcal{V}_e except for one-hop KG nodes.

The result is shown in Table 6. PipeNet with one-hop is the result of the grounded subgraph constructed by the matched concepts in question and answers. As shown in Figure 4, pruning rate 90% brings in almost same quantity of edges and nodes to one-hop subgraphs, while much less than original two-hop subgraph.

Random sampling can also bring performance gain because the induced nodes are relevant to the QA context. However, the gain is not as much as the DP-pruning method. This shows that finding semantically related nodes can benefit more in subgraph representation learning.

6 Conclusion

In this work, we propose PipeNet, a grounding-pruning-reasoning pipeline for question answering

with knowledge graph. The pruning strategy utilizes the dependency structure of query context to prune noisy entity nodes in the grounded subgraph, benefiting the subgraph representation learning with GNNs. We further design a GAT-based module for the subgraph representation learning with simplified message flow. Experiment results on two standard benchmarks demonstrate the effectiveness of semantic dependency of concept items benefits the subgraph representation learning.

7 Acknowledgement

The authors of this paper were supported by the NSFC Fund (U20B2053) from the NSFC of China, the RIF (R6020-19 and R6021-20) and the GRF (16211520 and 16205322) from RGC of Hong Kong. We also thank the UGC Research Matching Grants (RMGS20EG01-D, RMGS20CR11, RMGS20CR12, RMGS20EG19, RMGS20EG21, RMGS23CR05, RMGS23EG08). We would like to thank the Turing AI Computing Cloud (TACC) (Xu et al., 2021) and HKUST iSING Lab for providing us computation resources on their platform.

References

- Jasmijn Bastings, Ivan Titov, Wilker Aziz, Diego Marcheggiani, and Khalil Sima'an. 2017. [Graph convolutional encoders for syntax-aware neural machine translation](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, EMNLP 2017, Copenhagen, Denmark, September 9-11, 2017*, pages 1957–1967. Association for Computational Linguistics.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Jie Chen, Tengfei Ma, and Cao Xiao. 2018. Fastgcn: Fast learning with graph convolutional networks via importance sampling. In *International Conference on Learning Representations*.
- Wei-Lin Chiang, Xuanqing Liu, Si Si, Yang Li, Samy Bengio, and Cho-Jui Hsieh. 2019. Cluster-gcn: An efficient algorithm for training deep and large graph convolutional networks. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*, pages 257–266.
- M. Ciaramita, A. Gangemi, E. Ratsch, J. Saric, and I. Rojas. 2005. Unsupervised learning of semantic relations between concepts of a molecular biology ontology. In *International Joint Conference on Artificial Intelligence*.
- Peter Clark, Oren Etzioni, Tushar Khot, Daniel Khashabi, Bhavana Mishra, Kyle Richardson, Ashish Sabharwal, Carissa Schoenick, Oyvind Tafjord, Niket Tandon, et al. 2020. From ‘f’ to ‘a’ on the ny regents science exams: An overview of the aristo project. *AI Magazine*, 41(4):39–53.
- Junnan Dong, Qinggang Zhang, Xiao Huang, Keyu Duan, Qiaoyu Tan, and Zhimeng Jiang. 2023. Hierarchy-aware multi-hop question answering over knowledge graphs. In *Proceedings of the ACM Web Conference 2023*, pages 2519–2527.
- Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Zhiyong Wu, Baobao Chang, Xu Sun, Jingjing Xu, and Zhifang Sui. 2022. A survey for in-context learning. *arXiv preprint arXiv:2301.00234*.
- C. Fellbaum and G. Miller. 1998. Automated discovery of wordnet relations. *Wordnet An Electronic Lexical Database*, 5:131–151.
- Yanlin Feng, Xinyue Chen, Bill Yuchen Lin, Peifeng Wang, Jun Yan, and Xiang Ren. 2020. Scalable multi-hop relational reasoning for knowledge-aware question answering. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1295–1309.
- Matthias Fey, Jan E Lenssen, Frank Weichert, and Jure Leskovec. 2021. Gnnautoscale: Scalable and expressive graph neural networks via historical embeddings. In *International Conference on Machine Learning*, pages 3294–3304. PMLR.
- Matthias Fey and Jan Eric Lenssen. 2019. Fast graph representation learning with pytorch geometric. *arXiv preprint arXiv:1903.02428*.
- Will Hamilton, Zhitao Ying, and Jure Leskovec. 2017. Inductive representation learning on large graphs. *Advances in neural information processing systems*, 30.
- Yongfeng Huang, Yanyang Li, Yichong Xu, Lin Zhang, Ruyi Gan, Jiaying Zhang, and Liwei Wang. 2023. Mvp-tuning: Multi-view knowledge retrieval with prompt tuning for commonsense reasoning. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13417–13432.
- Zixian Huang, Ao Wu, Jiaying Zhou, Yu Gu, Yue Zhao, and Gong Cheng. 2022. Clues before answers: Generation-enhanced multiple-choice qa. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*.
- Long Kang, Xiaoge Li, and Xiaochun An. 2024. Knowledge-aware adaptive graph network for commonsense question answering. *Journal of Intelligent Information Systems*, pages 1–20.
- S. K. Kang, L. Patil, A. Rangarajan, A. Moitra, and D. Dutta. 2015. Extraction of manufacturing rules from unstructured text using a semantic framework.

- In *ASME 2015 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*.
- Daniel Khashabi, Sewon Min, Tushar Khot, Ashish Sabharwal, Oyvind Tafjord, Peter Clark, and Hannaneh Hajishirzi. 2020. Unifiedqa: Crossing format boundaries with a single qa system. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1896–1907.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2020. Albert: A lite bert for self-supervised learning of language representations. In *International Conference on Learning Representations*.
- Bill Yuchen Lin, Xinyue Chen, Jamin Chen, and Xiang Ren. 2019. Kagnet: Knowledge-aware graph networks for commonsense reasoning. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2829–2839.
- Liyuan Liu, Haoming Jiang, Pengcheng He, Weizhu Chen, Xiaodong Liu, Jianfeng Gao, and Jiawei Han. 2019a. On the variance of the adaptive learning rate and beyond. In *International Conference on Learning Representations*.
- Meng Liu and Shuiwang Ji. 2022. Neighbor2seq: Deep learning on massive graphs by transforming neighbors to sequences. In *Proceedings of the 2022 SIAM International Conference on Data Mining (SDM)*, pages 55–63. SIAM.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019b. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Diego Marcheggiani, Jasmijn Bastings, and Ivan Titov. 2018. Exploiting semantics in neural machine translation with graph convolutional networks. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT, New Orleans, Louisiana, USA, June 1-6, 2018, Volume 2 (Short Papers)*, pages 486–492. Association for Computational Linguistics.
- Sahisnu Mazumder and Bing Liu. Context-aware path ranking for knowledge base completion.
- Todor Mihaylov, Peter Clark, Tushar Khot, and Ashish Sabharwal. 2018. Can a suit of armor conduct electricity? a new dataset for open book question answering. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2381–2391.
- Jinyoung Park, Hyeong Kyu Choi, Juyeon Ko, Hyeonjin Park, Ji-Hoon Kim, Jisu Jeong, Kyungmin Kim, and Hyunwoo Kim. 2023. Relation-aware language-graph transformer for question answering. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 13457–13464.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, Peter J Liu, et al. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21(140):1–67.
- Emanuele Rossi, Fabrizio Frasca, Ben Chamberlain, Davide Eynard, Michael Bronstein, and Federico Monti. 2020. Sign: Scalable inception graph neural networks. *arXiv preprint arXiv:2004.11198*, 7:15.
- Sunil Kumar Sahu, Fenia Christopoulou, Makoto Miwa, and Sophia Ananiadou. 2019. Inter-sentence relation extraction with document-level graph convolutional neural network. In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 4309–4316. Association for Computational Linguistics.
- Adam Santoro, David Raposo, David G Barrett, Mateusz Malinowski, Razvan Pascanu, Peter Battaglia, and Timothy Lillicrap. 2017. A simple neural network module for relational reasoning. *Advances in neural information processing systems*, 30.
- Michael Schlichtkrull, Thomas N Kipf, Peter Bloem, Rianne van den Berg, Ivan Titov, and Max Welling. 2018. Modeling relational data with graph convolutional networks. In *European semantic web conference*, pages 593–607. Springer.
- R. Sombatsrisomboon, Y. Matsuo, and M. Ishizuka. 2003. Acquisition of hypernyms and hyponyms from the www.
- Robyn Speer, Joshua Chin, and Catherine Havasi. 2017. Conceptnet 5.5: An open multilingual graph of general knowledge. In *Thirty-first AAAI conference on artificial intelligence*.
- Ying Su, Zihao Wang, Tianqing Fang, Hongming Zhang, Yangqiu Song, and Tong Zhang. 2022. Mico: A multi-alternative contrastive learning framework for commonsense knowledge representation. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 1339–1351.
- Yueqing Sun, Qi Shi, Le Qi, and Yu Zhang. 2022. JointLK: Joint reasoning with language models and knowledge graphs for commonsense question answering. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5049–5060. Association for Computational Linguistics.
- Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. 2019. Commonsenseqa: A question answering challenge targeting commonsense knowledge. 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 4149–4158.
- Alon Talmor, Ori Yoran, Ronan Le Bras, Chandra Bhagavatula, Yoav Goldberg, Yejin Choi, and Jonathan Berant. 2021. Commonsenseqa 2.0: Exposing the limits of ai through gamification. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 1)*.
- Dhaval Taunk, Lakshya Khanna, Siri Venkata Pavan Kumar Kandru, Vasudeva Varma, Charu Sharma, and Makarand Tapaswi. 2023. Grapeqa: graph augmentation and pruning to enhance question-answering. In *Companion Proceedings of the ACM Web Conference 2023*, pages 1138–1144.
- Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. Graph attention networks. In *International Conference on Learning Representations*.
- Kuan Wang, Yuyu Zhang, Diyi Yang, Le Song, and Tao Qin. 2022. Gnn is a counter? revisiting gnn for question answering. In *International Conference on Learning Representations*.
- X. Wang, P. Kapanipathi, R. Musa, M. Yu, K. Talamadupula, I. Abdelaziz, M. Chang, A. Fokoue, B. Makni, and N. Mattei. 2019a. Improving natural language inference using external knowledge in the science questions domain. In *National Conference on Artificial Intelligence*.
- Xiaoyan Wang, Pavan Kapanipathi, Ryan Musa, Mo Yu, Kartik Talamadupula, Ibrahim Abdelaziz, Maria Chang, Achille Fokoue, Bassem Makni, Nicholas Mattei, et al. 2019b. Improving natural language inference using external knowledge in the science questions domain. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 7208–7215.
- Yujie Wang, Hu Zhang, Jiye Liang, and Ru Li. 2023. Dynamic heterogeneous-graph reasoning with language models and knowledge representation learning for commonsense question answering. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14048–14063.
- 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.
- Kaiqiang Xu, Xinchun Wan, Hao Wang, Zhenghang Ren, Xudong Liao, Decang Sun, Chaoliang Zeng, and Kai Chen. 2021. Tacc: A full-stack cloud computing infrastructure for machine learning tasks. *arXiv preprint arXiv:2110.01556*.
- Michihiro Yasunaga, Hongyu Ren, Antoine Bosselut, Percy Liang, and Jure Leskovec. 2021. Qa-gnn: Reasoning with language models and knowledge graphs for question answering. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 535–546.
- Qichen Ye, Bowen Cao, Nuo Chen, Weiyuan Xu, and Yuexian Zou. 2023. Fits: Fine-grained two-stage training for knowledge-aware question answering. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 13914–13922.
- Haiyang Yu, Limei Wang, Bokun Wang, Meng Liu, Tianbao Yang, and Shuiwang Ji. 2022. Graphfm: Improving large-scale gnn training via feature momentum. In *International Conference on Machine Learning*, pages 25684–25701. PMLR.
- Hanqing Zeng, Muhan Zhang, Yinglong Xia, Ajitesh Srivastava, Andrey Malevich, Rajgopal Kannan, Viktor Prasanna, Long Jin, and Ren Chen. 2021. Decoupling the depth and scope of graph neural networks. *Advances in Neural Information Processing Systems*, 34:19665–19679.
- Hanqing Zeng, Hongkuan Zhou, Ajitesh Srivastava, Rajgopal Kannan, and Viktor Prasanna. 2019. Graphsaint: Graph sampling based inductive learning method. In *International Conference on Learning Representations*.
- Bo Zhang, Yue Zhang, Rui Wang, Zhenghua Li, and Min Zhang. 2020. Syntax-aware opinion role labeling with dependency graph convolutional networks. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 3249–3258. Association for Computational Linguistics.
- Wentao Zhang, Zeang Sheng, Mingyu Yang, Yang Li, Yu Shen, Zhi Yang, and Bin Cui. 2022a. Nafs: A simple yet tough-to-beat baseline for graph representation learning. In *International Conference on Machine Learning*, pages 26467–26483. PMLR.
- X Zhang, A Bosselut, M Yasunaga, H Ren, P Liang, C Manning, and J Leskovec. 2022b. Greaselm: Graph reasoning enhanced language models for question answering. In *International Conference on Representation Learning (ICLR)*.
- Ruilin Zhao, Feng Zhao, Liang Hu, and Guandong Xu. 2024. Graph reasoning transformers for knowledge-aware question answering. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 19652–19660.
- Ziwan Zhao, Linmei Hu, Hanyu Zhao, Yingxia Shao, and Yequan Wang. 2023. Knowledgeable parameter efficient tuning network for commonsense question answering. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9051–9063.

Method	CSQA(IHdev)	OBQA(test)
w/o KG	73.07	78.40
GPT3.5-turbo	72.29	83.20
PipeNet(DP)	78.13	87.80

Table 7: Accuracy comparison between GPT3.5-turbo and PipeNet(DP) on CSQA(IHdev) and OBQA(test)

A Appendix

A.1 Comparison with LLM

Large language models such as GPT3 (Brown et al., 2020) and ChatGPT have recently received interest and achieved remarkable success over various question-answering tasks. We further adopt a 3-shot in-context learning (Dong et al., 2022) to prompt GPT3.5-turbo and present the results in Table 7. For OBQA, we add additional textual evidence in the prompt template for a fair comparison. It shows that GPT3.5-turbo achieves decent performances on both of the benchmarks, with comparable or better performances to the supervised finetuning method without KG (w/o KG). Nevertheless, PipeNet(DP) outperforms GPT3.5-turbo by a large margin though with a much smaller language model Roberta-large. This demonstrates that knowledge graph is still a meaningful knowledge source for question-answering tasks and our pruning method benefits such QA tasks with knowledge graph.