Knowledge Base Question Answering through Recursive Hypergraphs

Naganand Yadati[1], Dayanidhi R[2], Vaishnavi S[2], Indira G[2], Srinidhi G[2]
[1] Indian Institute of Science, Bangalore
[2] Channabasaveshwara Institute of Technology, Tumkur
naganand@iisc.ac.in
{dayanidhirs.12, vaishnavisathish98, indiragowtham99, srinidhi.g71}@gmail.com

Abstract

Knowledge Base Question Answering (KBQA) is the problem of predicting an answer for a factoid question over a given knowledge base (KB). Answering questions typically requires reasoning over multiple links in the given KB. Humans tend to answer questions by *grouping* different objects to perform reasoning over acquired knowledge. Hypergraphs provide a natural tool to model group relationships. In this work, inspired by typical human intelligence, we propose a new method for KBQA based on hypergraphs. Existing methods for KBQA, though effective, do not *explicitly incorporate the recursive relational group structure* in the given KB. Our method, which we name RecHyperNet (Recursive Hypergraph Network), exploits a new way of modelling KBs through recursive hypergraphs to organise such group relationships in KBs. Experiments on multiple KBQA benchmarks demonstrate the effectiveness of the proposed RecHyperNet. We have released the code.

1 Introduction

Knowledge Base Question Answering (KBQA) (Xu et al., 2016)), a task that tests the ability of a machine to understand knowledge like a human, is a challenging, central, and popular task in natural language processing. KBQA is the problem of predicting an answer for a factoid question over a given knowledge base (KB) containing facts such as (Inception, written by, Christopher Nolan). Answering questions typically requires reasoning over multiple links in the given KB (Zhang et al., 2018).

When a typical human answers a question (e.g. What are the genres of movies written by Christopher Nolan?), they tend to group objects (e.g. movies written by Christopher Nolan such as Inception, Interstellar, etc. are grouped together) over their acquired knowledge. Hypergraphs are mathematical tools that naturally encode group relationships. Directed hypergraphs have been recently used to model KBs for KBQA (Han et al., 2020) but we argue that they fail to naturally model group relationships without loss of information. Spcifically, two KB triples such as (Inception, written by, Christopher Nolan) and (Person of Interest, written by, Jonathan Nolan) would be modelled by a directed hyperedge {Inception, Person of Interest} -> {Christopher Nolan, Jonathan Nolan} clearly resulting in loss of information (not clear who directed which). Though objects of similar types are grouped together and directions are encoded, directed hypergraphs fail to model KBs without loss of information.

We argue that recursive hypergraphs (Joslyn and Nowak, 2017; Menezes and Roth, 2019) provide a flexible way to model KBs without loss of information. Specifically, all movies written by Christopher Nolan are grouped together and those written by Jonathan Nolan are *separately* grouped together. To summarise, we make the following contributions

- We model KBs as recursive hypergraphs. To the best of our knowledge, this is the first such attempt for natural language processing.
- We propose RecHyperNet (Recursive Hypergraph Network), a novel graph neural network-based model to exploit recursive hypergraphs and apply it for KBQA.
- We show the effectiveness of RecHyperNet on multiple KBQA benchmarks. We have released the source code in the supplementary.

2 Related Work

Knowledge Base Question Answering (KBQA) has been widely investigated especially through

Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics, pages 448–454

end-to-end deep neural networks since the release of popular datasets such as MetaQA (Zhang et al., 2018). One of the first attempts at KBQA was key-value memory network (KVMN) (Miller et al., 2016) that maintained memory to store KB facts and text as key-value pairs. Graphs of Relations Among Facts and Text Networks (GraftNet) (Sun et al., 2018) uses a heterogeneous GCN-based method to fuse information from heterogeneous sources (KBs and text). SubgraphReader (SGReader) (Xiong et al., 2019) employs a graph attention-based method to combine unstructured text and structured KB. PullNet (Sun et al., 2019) also uses a GCN to identify subgraph nodes that should be retrieved ('pull') from text and KB. More recently, EmbedKGQA (Saxena et al., 2020) uses ideas from knoweldge graph embedding literature to improve knowledge base question answering esp. on sparser incomplete knowledge graphs. Two-Phase Hypergraph Based Reasoning with Dynamic Relations (2HR-DR) (Han et al., 2020) is a directed hypergraph-based method for KBQA. The reader is referred to a comprehensive literature review on this topic (Fu et al., 2020).

Graph Neural Networks (GNNs): While deep neural networks such as convolutional neural networks and recurrent neural networks are specially designed for grids and sequences respectively, GNNs (Kipf and Welling, 2017; Hamilton et al., 2017; Veličković et al., 2018) are specially designed to learn representations from graph-structured data. GNNs have been recently extended to hypergraphs (Feng et al., 2019; Yadati et al., 2019). The first published works of GNNs for NLP investigated the tasks of semantic-role labelling (Marcheggiani and Titov, 2017), and neural machine translation (Bastings et al., 2017). A recent tutorial touches upon the recent advances of GNNs in NLP (Vashishth et al., 2019).

GNNs for Question Answering (QA): A popular task closely related to question answering is multi-hop reasoning across documents (context passages) where GNNs have been extensively used (Qiu et al., 2019b; Ding et al., 2019; Tu et al., 2020). Graph Attention Networks (Veličković et al., 2018) have been shown to be effective for modelling the multi-grained structure of documents for machine reading comprehension (MRC) (Zheng et al., 2020). GNNs have also recently been used on heterogeneous data sources for MRC (Kim et al., 2019; Tu et al., 2019), knowledge graphs for QA (De Cao et al., 2019; Cao et al., 2019; Lin et al., 2019; Qiu et al., 2019a; Feng et al., 2020; Ji et al., 2020; Shao et al., 2020), answer sentence selection (Tian et al., 2020), and QA on tables (Zhang, 2020) and math (Chen et al., 2020). In all these publications except 2HR-DR (Han et al., 2020), the input is a graph and limited to modelling pairwise relationships. We address this fundamental limitation by modelling a knowledge graph as a recursive hypergraph 2HR-DR models knowledge graphs through directed hypergraphs but we argue that they are more naturally represented by recursive hypergraphs.

3 Method

We describe the KBQA problem, and how recursive hypergraphs (Menezes and Roth, 2019; Joslyn and Nowak, 2017) can be used to model knowledge bases. We then describe our method of exploiting recursive hypergraphs for KBQA.

3.1 KBQA Problem

The KBQA problem considered in this work is as follows. We are given a knowledge base $\mathcal{K} \subseteq \mathcal{E} \times \mathcal{R} \times \mathcal{E}$ with entity set $\mathcal{E}$ and relation set $\mathcal{R}$. We are also given a natural language question q with a topic entity $e \in \mathcal{E}$. The task is to output an answer $a \in \mathcal{E}$ that correctly answers the question q.

3.2 Recursive Hypergraph

In this subsection, we precisely define a recursive hypergraph by first defining the following.

Definition 1 (Depth k Powerset)**.** *For a set S, let us use $\mathcal{S}(S)$ to denote the powerset of S i.e. $\mathcal{S}(S) := \{\dot{S} : \varnothing \subseteq \dot{S} \subseteq S\}$. Then, the the depth k powerset of S is*

$$2^{S,k} := \mathcal{S}\left(\bigcup_{i=0}^{k} S_i\right) \tag{1}$$

where $S_0 = S, S_i = \mathcal{S}\left(\bigcup_{j=0}^{i-1} S_j\right)$, for $i \geq 1$.

Note that $2^{S,0} = \mathcal{S}(S)$ i.e. $2^{S,0}$ is powerset of S.

Definition 2 (k-Recursive Hypergraph)**.** *A pair $H = (V, E)$, where V is a set of n vertices, and $E \subseteq \left(2^{V,k} - \varnothing\right)$ is a set of recursive hyperedges.*

Note that a hypergraph in the traditional sense is a 0-recursive hypergraph. We call a hyperedge $e \in E$ as a depth k hyperedge if $e \subseteq 2^{V,k}$ but $e \nsubseteq 2^{V,k}$.

3.3 Modelling knowledge Base as a Recursive Hypergraph

One of our main contributions in this work is to model a knowledge base as a 1-recursive hypergraph with $V = \mathcal{E} \cup \mathcal{R}$ as the set of vertices. Each head entity can be seen as a depth 1 hyperedge connecting all its relations. Each relation can be further seen as a depth 0 hyperedge connecting all the tail entities. For example, if a knowledge base contains the movie "Inception" as a head entity with (relation, object) pairs as 1) (starred actors, Leonardo DiCaprio), 2) (starred actors, Ellen Page), 3) (starred actors, Tom Hardy), 4) (genre, action), 5) (genre, adventure), and 6) (genre, science fiction), and 7) (written by, Christopher Nolan) then we can view them as a recursive hypergraph with Inception as a depth 1 hyperedge connecting relation vertices starred actors, genre, and written by. Each of these relation vertices is in turn a depth 0 hyperedge connecting object entities. For example, the object entities Leonardo DiCaprio, Ellen Page, and Tom Hardy are contained in a single hyperedge (that represents the pair Inception, starred actors).

3.4 RecHyperNet for KBQA

In this section, we describe our proposed methodology for KBQA. The main ingredients/modules of RecHyperNet are three modules: KB Embedding Module, Topic and Question Embedding Module, and Answer Retrieval Module.

1) KB Embedding Module uses knowledge base embedding methods to initialise entity embeddings of the input KB. We denote these initial embeddings by $x_e, e \in \mathcal{E}$. Popular knowledge base embedding methods include TransE (Bordes et al., 2013), and ComplEX (Trouillon et al., 2016).

2) Topic and Question Embedding Module embeds the question q to a fixed dimension vector $\boldsymbol{q}$ of dimension d. We use a feed-forward neural network that first represents q using LSTM/ RoBERTa (Liu et al., 2019). The topic entity e that is also present in the question is embedded using a multi-relational graph convolutional network on the KB (Vashishth et al., 2020). The entity update rule for Composition-based multi-relational graph convolutional network (CompGCN) is as follows:

$$\boldsymbol{h}_v = f\left(\sum_{(u,r)\in\mathcal{N}(v)} \boldsymbol{W}\phi(\boldsymbol{x}_u, \boldsymbol{z}_r) \right), \qquad (2)$$

where $\boldsymbol{x}_u$ and $\boldsymbol{z}_r$ are initial features of vertex u and relation r respectively. ϕ is a composition function that is dictated by the knowledge base embedding method and f is a non-linear activation function such as Rectified Linear Unit. $\mathcal{N}(v)$ is the (relation, object) neighbourhood of the vertex v (for example, the 7 pairs listed for "inception" movie entity in the previous subsection).

Our key modification to CompGCN update rule is motivated through the proposed recursive hypergraph view of knowledge base. Specifically, since each relation can be seen as a hyperedge containing all object entities, we can modify the update rule as follows

$$\boldsymbol{h}_v = f\left(\sum_{r\in N(v)} \boldsymbol{W}\phi\left(\max_{u\in\mathcal{N}(v,r)} (\boldsymbol{x}_u), \boldsymbol{z}_r \right) \right), \qquad (3)$$

where $N(v)$ is the set of all relations in the neighbourhood of v, and $\mathcal{N}(v, r)$ is the set of all object entities connected by the (subject, relation) pair (v, r). $\max$ is the element-wise maximum of a set of vectors. We can replace $\max$ by other aggregator functions such as mean, and sum but we experimentally observed that $\max$ gives the best performance. We embed the topic head entity $e \in \mathcal{E}$ using Equation 3 and obtain the hidden representation $\boldsymbol{h}_e$ with the same dimension d. We finally pass the concatenated representation $\boldsymbol{h}_e||\boldsymbol{q}$ to a 2-layer feed-forward neural network to get the topic and question representation.

3) Answer Selection Module Given a KB embedding scoring function Φ and a set of answer entities $\mathcal{A} \subseteq \mathcal{E}$, we learn the topic and question representation $\boldsymbol{h}_e||\boldsymbol{q}$ so that $\Phi(x_e,\ \boldsymbol{h}_e||\boldsymbol{q},\ x_a) > 0,\ \forall a \in \mathcal{A}$ and $\Phi(x_e,\ \boldsymbol{h}_e||\boldsymbol{q},\ x_{\bar{a}}) > 0,\ \forall \bar{a} \notin \mathcal{A}$. We leartn the model parameters by minimising the binary cross entropy loss between the sigmoid of the scores and the target labels (1 for correct, 0 for wrong).

During test time, we score the given (topic head entity, question) pair against all possible answers $a' \in \mathcal{E}$. So, the answer is given by $\arg\max_{a'\in\mathcal{E}} \Phi(x_e,\ \boldsymbol{h}_e||\boldsymbol{q},\ x_{a'})$.

4 Experiments

We evaluate our proposed method on MetaQA (Zhang et al., 2018) and WebQuestionsSP (Yih et al., 2016) datasets. We closely follow the experimental setup of a prior work (Saxena et al., 2020) for the preprocessed versions of these datasets. MetaQA consists of $400k$ questions ($1-hop$, 2-hop, and 3-hop), and a knowledge graph of 135k triples, 43k entities, and 9 relations. WebQuestionsSP consists of 4700 questions and a knowledge

Model	MetaQA KG-Full			MetaQA KG-50			WebQSP KG-Full	WebQSP KG-50
	1-hop	2-hop	3-hop	1-hop	2-hop	3-hop		
KVMN (Miller et al., 2016)	96.2	82.7	48.9	75.7	48.4	37.6	46.7	32.7
VRN (Zhang et al., 2018)	97.5	89.9	62.5	–	–	–	–	–
GraftNet (Sun et al., 2018)	97.0	94.8	77.7	91.5	69.5	66.4	66.8	49.7
SGReader (Xiong et al., 2019)	96.7	80.7	68.6	79.2	77.1	63.5	–	–
PullNet (Sun et al., 2019)	97.0	**99.9**	91.4	**92.4**	90.4	**85.2**	**68.1**	51.9
EmbedKGQA (Saxena et al., 2020)	97.5	98.8	94.8	83.9	91.8	70.3	66.6	53.2
2HR-DR (Han et al., 2020)	98.8	93.7	91.4	80.8	89.3	65.1	67.0	52.2
RecHyperNet (ours)	**99.1**	99.2	**95.0**	84.4	**92.3**	71.1	**68.4**	**53.7**

Table 1: Results (higher is better) on MetaQA, WebQuestionsSP datasets. Baselines such as PullNet utilise external corpus to answer questions (unrealistic as it is not always readily available) while our method does not.

Model	1-hop	2-hop	3-hop
GraftNet	64.0	52.6	59.2
PullNet	65.1	52.1	59.7
RecHyperNet	**84.4**	**92.3**	**71.1**

Table 2: Results (higher is better) of baselines without text corpus on MetaQA KG-50.

Model	MetaQA	WebQSP
RecHyperNet	**92.3**	**53.7**
replace Eq. 3 by Eq. 2	90.4	52.1
use $\boldsymbol{q}$ in place of $\boldsymbol{h}_e \| \boldsymbol{q}$	85.6	50.7

Table 3: Ablation study on MetaQA KG-50 (2-hop) and WebQSP KG-50

graph of 1.8 million entities, and 5.7 million triples. Following prior work (Saxena et al., 2020), we experimented on two different settings (for both datasets) - **KG Full** (in which the KG is left untouched), and the more realistic **KG-50** setting in which 50% links are randomly removed. We compared against 7 different baselines as shown in Table 1. Please see Section 2 (subsection: Knowledge Base Question Answering) for brief descriptions of the baseline methods.

Model details Following prior work (Saxena et al., 2020), we used a long short term memory (LSTM) network to learn embeddings for words in the questions with an embedding size of 256 for MetaQA and RoBERTa (768 dimensional embeddings) (Liu et al., 2019) for WebQuestionsSP datasets. The hidden dimension size for graph convolutional network was also set to 256. A dropout rate of 0.2 was used for all neural layers.. All models were implemented in PyTorch (Paszke et al., 2019) and trained with ADAM as the optimiser (Kingma and Ba, 2015), a learning rate of 5×10^{-4}, weight decay of 1.0, a batch size of 128 trained for 100 epochs (with patience of 5).

Results. Table 1 shows the results. We evaluated our proposed method and all baselines through Hits@1 metric. As we can see in Table 1, exploiting recursive hypergraphs through our proposed method help. Methods such as PullNet and GraftNet utilise additional text corpus (unrealistic setting as it is not always readily available) to answer questions while our method does not. We exploit only the given knowledge base structure. Futhermore, as shown in Table 2, baselines such as GraftNet and PullNet perform poorly in the absence of text.

Ablation Analysis We conducted an ablation study by removing essential components from RecHyperNet. Specifically we replaced our proposed Equation 3 to exploit recursive hypergraphs by the exisdting Equation 2 of CompGCN. As another basline, we removed the embedding of topic head entity obtained through GCN and used only the question representation while scoring. As shown in Table 3, both these components are essential for our proposed RecHyperNet.

Conclusion and Future Work In this work, we exploit recursive hypergraphs in NLP for the task of KBQA. We have proposed a novel method based on graph convolutional networks. We have demonstrated the effectiveness of RecHyperNet on KBQA benchmarks. In future, we exploit recursive structures for other tasks where graph neural nets are effective such as question generation (Pan et al., 2020), sentiment analysis (Wang et al., 2020), etc.

Acknowledgements

This work is supported by the Ministry of Human Resource Development (Government of India).

References

Joost Bastings, Ivan Titov, Wilker Aziz, Diego Marcheggiani, and Khalil Simaan. 2017. Graph convolutional encoders for syntax-aware neural machine translation. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1957–1967.

Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. 2013. Translating embeddings for modeling multi-relational data. In *Advances in Neural Information Processing Systems (NeurIPS) 26*, pages 2787–2795. Curran Associates, Inc.

Yu Cao, Meng Fang, and Dacheng Tao. 2019. BAG: Bi-directional attention entity graph convolutional network for multi-hop reasoning question answering. In *Proceedings of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, pages 357–362.

Kunlong Chen, Weidi Xu, Xingyi Cheng, Zou Xiaochuan, Yuyu Zhang, Le Song, Taifeng Wang, Yuan Qi, and Wei Chu. 2020. Question directed graph attention network for numerical reasoning over text. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.

Nicola De Cao, Wilker Aziz, and Ivan Titov. 2019. Question answering by reasoning across documents with graph convolutional networks. In *Proceedings of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, pages 2306–2317.

Ming Ding, Chang Zhou, Qibin Chen, Hongxia Yang, and Jie Tang. 2019. Cognitive graph for multi-hop reading comprehension at scale. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 2694–2703.

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)*.

Yifan Feng, Haoxuan You, Zizhao Zhang, Rongrong Ji, and Yue Gao. 2019. Hypergraph neural networks. In *Proceedings of the Thirty-Third Conference on Association for the Advancement of Artificial Intelligence (AAAI)*, pages 3558–3565.

Bin Fu, Yunqi Qiu, Chengguang Tang, Yang Li, Haiyang Yu, and Jian Sun. 2020. A survey on complex question answering over knowledge base: Recent advances and challenges. *CoRR*, abs/2007.13069.

Will Hamilton, Zhitao Ying, and Jure Leskovec. 2017. Inductive representation learning on large graphs. In *Advances in Neural Information Processing Systems (NeurIPS) 30*, pages 1024–1034. Curran Associates, Inc.

Jiale Han, Bo Cheng, and Xu Wang. 2020. Two-phase hypergraph based reasoning with dynamic relations for multi-hop kbqa. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence (IJCAI)*, pages 3615–3621.

Haozhe Ji, Pei Ke, Shaohan Huang, Furu Wei, Xiaoyan Zhu, and Minlie Huang. 2020. Language generation with multi-hop reasoning on commonsense knowledge graph. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.

Hanqi Jin, Tianming Wang, and Xiaojun Wan. 2020. Semsum: Semantic dependency guided neural abstractive summarization. In *Proceedings of the Thirty-Fourth Conference on Association for the Advancement of Artificial Intelligence (AAAI)*.

Cliff Joslyn and Kathleen Nowak. 2017. Ubergraphs: A definition of a recursive hypergraph structure. *Computing Research Repository (CoRR)*, abs/1704.05547.

Daesik Kim, Seonhoon Kim, and Nojun Kwak. 2019. Textbook question answering with multi-modal context graph understanding and self-supervised open-set comprehension. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 3568–3584.

Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*.

Thomas N Kipf and Max Welling. 2017. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations (ICLR)*.

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.

Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized BERT pretraining approach. *CoRR*, abs/1907.11692.

Diego Marcheggiani and Ivan Titov. 2017. Encoding sentences with graph convolutional networks for semantic role labeling. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1506–1515.

Telmo Menezes and Camille Roth. 2019. Semantic hypergraphs. *Computing Research Repository (CoRR)*, abs/1908.10784.

Alexander Miller, Adam Fisch, Jesse Dodge, Amir-Hossein Karimi, Antoine Bordes, and Jason Weston. 2016. Key-value memory networks for directly reading documents. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1400–1409.

Liangming Pan, Yuxi Xie, Yansong Feng, Tat-Seng Chua, and Min-Yen Kan. 2020. Semantic graphs for generating deep questions. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL)*.

Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems (NeurIPS) 32*, pages 8026–8037. Curran Associates, Inc.

Delai Qiu, Yuanzhe Zhang, Xinwei Feng, Xiangwen Liao, Wenbin Jiang, Yajuan Lyu, Kang Liu, and Jun Zhao. 2019a. Machine reading comprehension using structural knowledge graph-aware network. 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 5898–5903.

Lin Qiu, Yunxuan Xiao, Yanru Qu, Hao Zhou, Lei Li, Weinan Zhang, and Yong Yu. 2019b. Dynamically fused graph network for multi-hop reasoning. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 6140–6150.

Apoorv Saxena, Aditay Tripathi, and Partha Talukdar. 2020. Improving multi-hop question answering over knowledge graphs using knowledge base embeddings. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 4498–4507.

Nan Shao, Yiming Cui, Ting Liu, Shijin Wang, and Guoping Hu. 2020. Is graph structure necessary for multi-hop reasoning? In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.

Haitian Sun, Tania Bedrax-Weiss, and William Cohen. 2019. PullNet: Open domain question answering with iterative retrieval on knowledge bases and text. 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 2380–2390.

Haitian Sun, Bhuwan Dhingra, Manzil Zaheer, Kathryn Mazaitis, Ruslan Salakhutdinov, and William W. Cohen. 2018. Open domain question answering using early fusion of knowledge bases and text. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4231–4242.

Zhixing Tian, Yuanzhe Zhang, Xinwei Feng, Wenbin Jiang, Yajuan Lyu, Kang Liu, and Jun Zhao. 2020. Capturing sentence relations for answer sentence selection with multi-perspective graph encoding. In *Proceedings of the Thirty-Fourth Conference on Association for the Advancement of Artificial Intelligence (AAAI)*.

Théo Trouillon, Johannes Welbl, Sebastian Riedel, Éric Gaussier, and Guillaume Bouchard. 2016. Complex embeddings for simple link prediction. In *Proceedings of the 33rd International Conference on International Conference on Machine Learning (ICML) - Volume 48*, page 2071–2080.

Ming Tu, Kevin Huang, Guangtao Wang, Jing Huang, Xiaodong He, and Bowen Zhou. 2020. Select, answer and explain: Interpretable multi-hop reading comprehension over multiple documents. In *Proceedings of the Thirty-Fourth Conference on Association for the Advancement of Artificial Intelligence (AAAI)*.

Ming Tu, Guangtao Wang, Jing Huang, Yun Tang, Xiaodong He, and Bowen Zhou. 2019. Multi-hop reading comprehension across multiple documents by reasoning over heterogeneous graphs. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 2704–2713.

Shikhar Vashishth, Soumya Sanyal, Vikram Nitin, and Partha Talukdar. 2020. Composition-based multi-relational graph convolutional networks. In *International Conference on Learning Representations (ICLR)*.

Shikhar Vashishth, Naganand Yadati, and Partha Talukdar. 2019. Graph-based deep learning in natural language processing. 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): Tutorial Abstracts*.

Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. Graph attention networks. In *International Conference on Learning Representations (ICLR)*.

Kai Wang, Weizhou Shen, Yunyi Yang, Xiaojun Quan, and Rui Wang. 2020. Relational graph attention network for aspect-based sentiment analysis. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL)*.

Wenhan Xiong, Mo Yu, Shiyu Chang, Xiaoxiao Guo, and William Yang Wang. 2019. Improving question answering over incomplete KBs with knowledge-aware reader. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 4258–4264.

Kun Xu, Siva Reddy, Yansong Feng, Songfang Huang, and Dongyan Zhao. 2016. Question answering on Freebase via relation extraction and textual evidence. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (ACL, Volume 1: Long Papers)*, pages 2326–2336.

Naganand Yadati, Madhav Nimishakavi, Prateek Yadav, Vikram Nitin, Anand Louis, and Partha Talukdar. 2019. HyperGCN: A new method of training graph convolutional networks on hypergraphs. In *Advances in Neural Information Processing Systems (NeurIPS) 32*, pages 1509–1520. Curran Associates, Inc.

Wen-tau Yih, Matthew Richardson, Chris Meek, Ming-Wei Chang, and Jina Suh. 2016. The value of semantic parse labeling for knowledge base question answering. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 201–206.

Xuanyu Zhang. 2020. Cfgnn: Cross flow graph neural networks for question answering on complex tables. In *Proceedings of the Thirty-Fourth Conference on Association for the Advancement of Artificial Intelligence (AAAI)*.

Yuyu Zhang, Hanjun Dai, Zornitsa Kozareva, Alexander J. Smola, and Le Song. 2018. Variational reasoning for question answering with knowledge graph. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence (AAAI)*, pages 6069–6076.

Bo Zheng, Haoyang Wen, Yaobo Liang, Nan Duan, Wanxiang Che, Daxin Jiang, Ming Zhou, and Ting Liu. 2020. Document modeling with graph attention networks for multi-grained machine reading comprehension. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL)*.