

Open Domain Question Answering based on Text Enhanced Knowledge Graph with Hyperedge Infusion

Jiale Han, Bo Cheng, Xu Wang

State Key Laboratory of networking and switching technology,
Beijing University of Posts and Telecommunications
{hanjl, chengbo, wxx}@bupt.edu.cn

Abstract

The incompleteness of knowledge base (KB) is a vital factor limiting the performance of question answering (QA). This paper proposes a novel QA method by leveraging text information to enhance the incomplete KB. The model enriches the entity representation through semantic information contained in the text, and employs graph convolutional networks to update the entity status. Furthermore, to exploit the latent structural information of text, we treat the text as hyperedges connecting entities among it to complement the deficient relations in KB, and hypergraph convolutional networks are further applied to reason on the hypergraph-formed text. Extensive experiments on the WebQuestionsSP benchmark with different KB settings prove the effectiveness of our model.

1 Introduction

Open domain question answering (QA) is a challenging task that requires answering the factual questions in natural language. According to the structure of supporting information, QA system can be divided into knowledge-based QA (KBQA) (Bordes et al., 2015; Zhang et al., 2018) and text-based QA (TBQA) (Welbl et al., 2018; Yang et al., 2018). KBQA obtains the answers by a structured knowledge base, which is easy to query and reason with but limited by the incompleteness of well-designed triples. TBQA’s supporting information is plain text containing rich semantic and latent structural information, however, it’s difficult for a machine to understand. The complementary properties inspire us to fuse these two kinds of data to enhance the incomplete KB and further improve the QA system’s performance.

Some work has already been proposed. Das et al. (2017) represent KB and text using universal schema and apply memory networks, but lack the

Question: Which university did *Cleary* graduate from?

Text: *Cleary* obtained her *master’s degree* in *Library Science* at the *University of Washington*.




Figure 1: Example of a question with its related KB and text. The KB is incomplete to answer the question, which lacks relation "graduated university" and entity "University of Washington". By completing the missing information with adding text as hyperedge, we can handle the question more effectively.

association between KB and text. Sun et al. (2018) build a heterogeneous graph with entities and text as nodes and employ a graph based method. Xiong et al. (2019) first encode entities in KB by graph attention networks and then read text with the help of accumulated entity knowledge. Although good results have been achieved, the text information is not fully utilized, especially the relation information among the entities contained in the text. Figure 1 shows an example that the KB is insufficient to answer the question. This question can be adequately answered by using the structural information of the text to bring high-level relationships.

In this paper, we propose a novel QA model based on text enhanced knowledge graph, which enriches entity representation by text semantic information and complements the relations in KB through structural information of the text. Specifically, the model firstly encodes entities in KB combining text information and applies graph convolutional networks (GCN) (Wu et al., 2020) to reason across KB. Note that a document usually mentions multiple entities, we convert the unstructured text into a structured hypergraph by regarding text

as hyperedge connecting entities among text, and then employ hypergraph convolutional networks (HGCN) (Feng et al., 2019; Yadati et al., 2019) to further update the entity states. Finally, the model predicts the final answers.

Our highlights are summarized as follows: 1) We novelly treat documents as high-order relations (hyperedges) connecting entities mentioned in them. 2) We apply Hypergraph Convolutional Networks to reason and propose the dual-step attention to catch the importance of different entities and documents. 3) Extensive experiments conducted on the widely used WebQuestionsSP (Yih et al., 2016) with different KB settings demonstrate our model is effective.

2 Related Work

The combination of knowledge base and text in QA is a challenging task, which has attracted many researchers’ attention. The work of (Das et al., 2017) extends universal schema to question answering and employs Key-Value Memory networks to process to text and KB. Sun et al. (2018) regard documents as heterogeneous nodes and combine them with entities in KB to form a uniform graph. The model proposed by Xiong et al. (2019) contains a graph-attention based KB reader and a knowledge-aware text reader. Some other work focuses on retrieving a small graph that contains just the question-related information (Sun et al., 2019) and the interpretability of QA on KB and text (Sydorova et al., 2019). These methods lack considering the high-order relationship among the entities contained in the text. This paper regards the text as hyperedge and further employs hypergraph convolutional networks.

Hypergraph convolutional networks (Feng et al., 2019; Yadati et al., 2019) utilize hypergraph structure rather than a general graph to represent the high-order correlation among data entirely, and hypergraph attention (Bai et al., 2019) further enhances the ability of representation learning by using an attention module.

3 Model

3.1 Task Definition

To maintain consistency and fairness, we adopt the same setting as Sun et al. (2018) that builds a subgraph for each question. Specifically, given a question $q = (w_1, w_2, \dots, w_{|q|})$, the related sub knowledge graph $\mathcal{K} = (\mathcal{V}, \mathcal{E}, \mathcal{T})$ is extracted by the

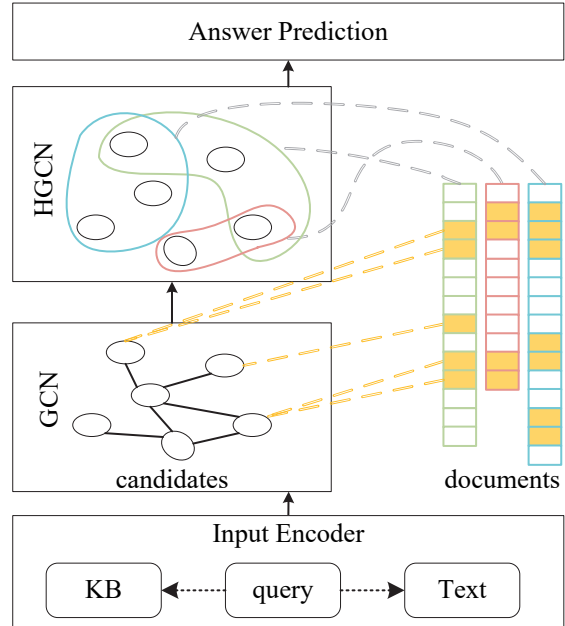


Figure 2: The overview of the model. We utilize the semantic information mentioned in the text to enrich the entity representation, and novelly treat text as hyperedges to complement the relation in incomplete KB.

Personalized PageRank (Haveliwala, 2002), where \mathcal{V} is the entity set, \mathcal{E} is the relation set, and \mathcal{T} contains a set of triples (v_h, r, v_t) indicated there is a relation $r \in \mathcal{E}$ between $v_h \in \mathcal{V}$ and $v_t \in \mathcal{V}$. Also a relevant text corpus $\mathcal{D} = \{d_1, d_2, \dots, d_{|\mathcal{D}|}\}$ is retrieved from Wikipedia by an off-the-shelf document retriever (Chen et al., 2017), which $d_i = (w_1, w_2, \dots, w_{|d_i|})$ represents a document and the entities mentioned in documents have been linked. The task requires to extract answers from all KB and document entities. The overview of our model is shown in Figure 2.

3.2 Input Encoder

Query and Text Encoder: Let $\mathbf{X}_q \in \mathbb{R}^{|q| \times n}$ and $\mathbf{X}_d \in \mathbb{R}^{|d| \times n}$ be the embedding matrices of query q and document $d \in \mathcal{D}$, where n is the embedding dimension. Bi-LSTM networks (Hochreiter and Schmidhuber, 1997) are applied to encode the query and document separately and get the hidden states $\mathbf{H}_q \in \mathbb{R}^{|q| \times h}$ and $\mathbf{H}_d \in \mathbb{R}^{|d| \times h}$, h is the hidden dimension of bi-LSTM. Then we compute the representation of query h_q and document h_d with attention mechanism.

$$h_q = \mathbf{H}_q^T \text{softmax}(f_q(\mathbf{H}_q)) \in \mathbb{R}^{h \times 1}$$

$$h_d = \mathbf{H}_d^T \text{softmax}(f_d(\mathbf{H}_d \mathbf{H}_d^T)) \in \mathbb{R}^{h \times 1}$$

where T represents matrix transposition, f_q is a linear network which converts h dimension to 1 dimension, and f_d converts $|q|$ dimension to 1 dimension.

KB Encoder: Each entity $v \in \mathcal{V}$ is initialized by pre-trained knowledge graph embedding $x_v \in \mathbb{R}^{n \times 1}$. And relation is initialized by semantic vector and KG embedding. Specifically, for relation $r \in \mathcal{E}$ and its KG embedding $x_r \in \mathbb{R}^{n \times 1}$, we tokenize it as $r = (w_1, w_2, \dots, w_{|r|})$ and feed into bi-LSTM layer with word embedding to get the hidden states $\mathbf{H}_r \in \mathbb{R}^{|r| \times h}$, then calculate the representation h_r as follows.

$$\begin{aligned} \mathbf{H}_{rq} &= \text{softmax}(\mathbf{H}_r \mathbf{H}_q^{\text{T}}) \mathbf{H}_q \in \mathbb{R}^{|r| \times h} \\ \mathbf{H}'_r &= [\mathbf{H}_r; \mathbf{H}_{rq}] \in \mathbb{R}^{|r| \times 2h} \\ h'_r &= \mathbf{H}'_r{}^{\text{T}} \text{softmax}(f_{r1}(\mathbf{H}'_r)) \in \mathbb{R}^{2h \times 1} \\ h_r &= f_{r2}([h'_r; x_r]) \in \mathbb{R}^{h \times 1} \end{aligned}$$

where $[\cdot]$ denotes column-wise concatenation, f_{r1} is a linear network which converts $2h$ dimension to 1 dimension, and f_{r2} converts $2h + n$ dimension to h dimension.

3.3 Reasoning over Text Enhanced Knowledge Graph

This component utilizes text information to improve the incomplete KB by enriching entity representation and adding hyperedges, and applies GCN and HGNCN to reasoning.

GCN for Entity-Enriched KB: To utilize the rich semantic information contained in the text, we construct a binary matrix \mathbf{M} , where $\mathbf{M}_d^v \in \mathbb{R}^{|d| \times 1}$ indicates the span of entity v in document d , and pass information from documents to entities to form text-aware entity representation x'_v , then concatenate with x_v as initial node state $h_v^{(0)}$.

$$\begin{aligned} x'_v &= \sum_{d \in D_v} \mathbf{H}_d^{\text{T}} \mathbf{M}_d^v \in \mathbb{R}^{h \times 1} \\ h_v^{(0)} &= f_v([x_v; x'_v]) \in \mathbb{R}^{h \times 1} \end{aligned}$$

where D_v is the linked documents set of entity v , f_v converts $h + n$ dimension to h dimension. Then the model learns the entity representation by aggregating the connected entity feature.

$$h_v^{(l_1+1)} = \mathbf{W}_1 h_v^{(l_1)} + \sum_{(v_i, r_i) \in N_v} \alpha_i \mathbf{W}_2 [h_{v_i}^{(l_1)}; h_{r_i}] \in \mathbb{R}^{h \times 1}$$

$$\alpha_i = \sigma(h_q^{\text{T}} f_a([h_{v_i}^{(l_1)}; h_{r_i}]))$$

where $\mathbf{W}_1 \in \mathbb{R}^{h \times h}$, $\mathbf{W}_2 \in \mathbb{R}^{h \times 2h}$ are learnable parameters, N_v represents the adjacent triple set of entity v , f_a converts $2h$ dimension to h dimension, l_1 represents the current GCN layer, which has a total of L_1 layers, and σ is the sigmoid function.

HGCN for Hypergraph-Formed Text: The model regards plain text as hyperedges connecting the entities among the text to complement the lack of relations in KB. HGNCN is employed to encode on the hypergraph-formed text. What's more, dual-step attention catches the importance of different entities and documents. Formally, at layer l_2 , the model first transfers the entity feature to the connected hyperedges to form the document representation,

$$\begin{aligned} h_d^{(l_2+1)} &= \mathbf{W}_3 h_d^{(l_2)} + \sum_{v_i \in N_d} \beta_i \mathbf{W}_4 h_{v_i}^{(l_2)} \in \mathbb{R}^{h \times 1} \\ \beta_i &= \sigma(h_q^{\text{T}} h_{v_i}^{(l_2)}) \end{aligned}$$

where $\mathbf{W}_3, \mathbf{W}_4 \in \mathbb{R}^{h \times h}$ are learnable parameters, $h_v^{(0)} = h_v^{(L_1)}$, $h_d^{(0)} = h_d$, and N_d represents the connected entity set of document d . Then the model gathers the documents' information to update the connected entity states.

$$\begin{aligned} h_v^{(l_2+1)} &= \mathbf{W}_5 h_v^{(l_2)} + \sum_{d_i \in D_v} \gamma_i \mathbf{W}_6 h_{d_i}^{(l_2+1)} \in \mathbb{R}^{h \times 1} \\ \gamma_i &= \sigma(h_q^{\text{T}} h_{d_i}^{(l_2+1)}) \end{aligned}$$

where $\mathbf{W}_5, \mathbf{W}_6 \in \mathbb{R}^{h \times h}$ are learnable parameters.

3.4 Answer Prediction

After L_1 GCN layers and L_2 HGNCN layers, the model finally predicts the probability of each entity being the answer,

$$p_v = \sigma(f_{out}(h_v^{(L_2)}))$$

where f_{out} converts h dimension to 1 dimension.

4 Experiments

4.1 Dataset

WebQuestionsSP (Yih et al., 2016) is a multi-answer QA dataset which contains 4737 questions. In our experiments we adopt the dataset¹ preprocessed by Sun et al. (2018). Table 1 shows the

¹<https://github.com/OceanskySun/GraftNet>

Dataset	questions	avg candidates	avg linked	avg answers	avg entities in documents
	train / dev / test	KB / Text / KB+Text	documents		
WebQSP	2848 / 250 / 1639	384.6 / 141.6 / 515.1	43.6	11.2	4.6

Table 1: Statistical results of the dataset WebQuestionsSP.

Model	KB-only	Text-only	KB+Text
KVMem	46.7 / 38.6	23.2 / 13.0	40.5 / 30.9
GraftNet	66.7 / 62.4	25.3 / 15.3	67.8 / 60.4
SG-KA	66.5 / 58.0	- / -	67.2 / 57.3
PullNet	68.1 / -	24.8 / -	- / -
Ours	66.9 / 60.1	27.2 / 17.1	68.4 / 60.6

Table 2: Hits@1 / F1 scores on WebQSP.

statistics of the dataset and retrieved subgraphs for the questions, including KB and linked text. In particular, the average number of linked entities in the documents is 4.6, which illustrates the rationality of adopting hyperedges.

4.2 Baseline Methods

We compare our methods with the following models:

- **KVMemNet** (Miller et al., 2016) is an end-to-end memory network which stores KB facts and text into key-value pairs.
- **GraftNet** (Sun et al., 2018) combines KB and text with the early fusion strategy and applies a graph-based model.
- **SG-KA Reader** (Xiong et al., 2019) proposes two components to reason over KB and incorporate entity information to text.
- **PullNet** (Sun et al., 2019) is a QA framework for learning how to retrieve small sub-graph related to answering the question.

4.3 Training Details

The model is implemented in PyTorch (Paszke et al., 2019) and trained on one Nvidia Tesla P40 GPU. We apply 100-dimensional TransE embeddings (Bordes et al., 2013) for entities and relations, and 300-dimensional GloVe embeddings (Pennington et al., 2014) for question and text words. The word numbers of questions and documents are limited to be 10 and 50. The hidden size is set to 100. We select the hyperparameter values by manual tuning to perform the best results on the validation dataset. The dropout is 0.2, and the batch size is 8.

The GCN layer L_1 and HGCN layer L_2 are 1 and 2 separately. The average runtime for one epoch is 5 minutes, and we set the max number of epochs to 200. The number of parameters is 69 million. The Adam optimizer (Kingma and Ba, 2015) is applied to minimize the binary cross-entropy loss with a learning rate of 0.0005. The threshold for F1 is set to 0.05.

4.4 Results

Main Results: The metrics adopted in the experiments are Hits@1, which is the accuracy of the top answer predicted by the model, and F1, which represents the ability to predict all answers. As shown in table 2, we experiment with our model under KB-only, Text-only, and KB+Text settings and compare them with baseline methods. Our model gets competitive performance in the KB-only setting and achieves the best results in the other two settings, especially in the Text-only setting, Hits@1 and F1 are 1.9% and 1.8% higher than the second-best method respectively, which shows the validity of treating documents as hyperedges. The promising performance may inspire us to handle similar tasks that build plain text to hypergraph and apply efficient HGCN. In KB+Text’s setting, our method also achieves the best performance, proving that our proposed enhancement strategy can effectively enhance incomplete KB by fully introducing the semantic and structural information implied in the text. In particular, our model improves a lot compared with KB-only, more than the work of (Sun et al., 2018), which demonstrates our way that treating documents as hyperedges is more productive than regarding them as heterogeneous nodes.

Different KB Setting: Following the work of Sun et al. (2018) that the KB is downsampled to different extents, we experiment on 10%, 30%, and 50% KB settings, which represents the percentage of required evidence covered by KB to simulate the situation of incomplete KB, and analyze the impact of the text on model performance. As shown in table 3, our model obtains the promising performance in the KB-only setting, especially the F1 metric all achieves the highest values, demon-

Model	10%		30%		50%	
	KB-only	KB+Text	KB-only	KB+Text	KB-only	KB+Text
KVMem	12.5 / 4.3	24.6 / 14.4	25.8 / 13.8	27.0 / 17.7	33.3 / 21.3	32.5 / 23.6
GraftNet	15.5 / 6.5	31.5 / 17.7	34.9 / 20.4	40.7 / 25.2	47.7 / 34.3	49.9 / 34.7
SG-KA	17.1 / 7.0	33.6 / 18.9	35.9 / 20.2	42.6 / 27.1	49.2 / 33.5	52.7 / 36.1
PullNet	- / -	- / -	- / -	- / -	50.3 / -	51.9 / -
Ours	18.3 / 7.9	33.7 / 19.9	35.2 / 21.0	42.8 / 27.5	49.3 / 34.3	52.8 / 37.1

Table 3: Hits@1 / F1 scores under different KB settings.

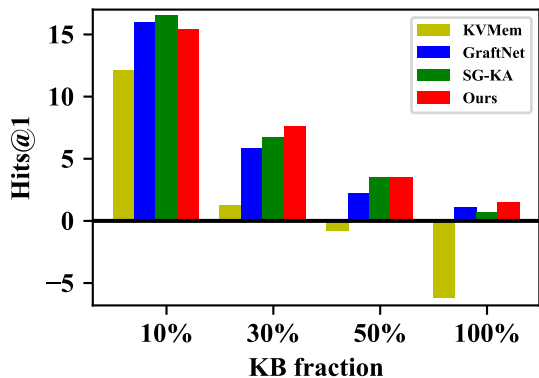


Figure 3: Improvement of KB+Text over KB-only under different KB fraction setting.

strating the ability of our method for multi-answer prediction. After combining the text, our model achieves the best results compared with the baseline methods. What’s more, the performance in the KB+Text setting has been significantly improved over the KB-only setting. The more incomplete the KB, the more obvious the performance improvement, which shows it can effectively use the document information to complete and enhance the KB, so as to further improve the performance of QA. In order to intuitively visualize the improvement, figure 3 displays the increment of all models after adding text under different settings (KB+Text–KB-only). We can observe that our method achieves the largest or almost the largest increment. What’s more, we notice the text information improves the performance obviously in the case of incomplete KB, but may cause the extra interference when the KB is sufficient to support answering questions, which even lead to performance degradation. This makes us think about how to effectively use text to further improve the performance of question answering under the full KB setting.

Ablation Study: An ablation study is conducted to evaluate the benefits of different components in

Model	10%KB+Text	
	Hits@1	F1
Full Model	33.7	19.9
–GCN attention	33.3	19.3
–dual-step attention	32.5	18.9
–entity-enriched KB	32.8	18.7

Table 4: Experimental results of ablation study.

the model. Table 4 shows the results under 10% KB setting. From the second and third rows, the attention mechanism adopted by the model is effective, especially the dual-step attention proposed at the HGCN layer, which brings 1.2% improvement of Hits@1. The strategy of entity-enriched KB also increases Hits@1 by 0.9%, proving its validity.

5 Conclusion

We propose a QA method that aims to enhance the incomplete KB by text information, which fully explored the semantic and latent structural information in the text. In particular, the text is treated as hyperedges to complement the lack of relations in KB. The model first applies GCN to encode the entity-enriched KB, then employs HGCN to further reason over hypergraph-formed text, and predicts the final answers. Experimental results on the WebQuestionsSP benchmark prove the effectiveness of our model and each component.

Acknowledgements

This work is supported by the Fundamental Research Funds for the Central Universities under grant 2020XD-A07-1, the National Key Research and Development Program of China under grant 2018YFB1003804, in part by the National Natural Science Foundation of China under grant 61921003, 61972043. Jiale Han is supported by BUPT Excellent Ph.D. Students Foundation (Grant No. CX2020102). We would like to thank all reviewers for careful and valuable comments.

References

- Song Bai, Feihu Zhang, and Philip H. S. Torr. 2019. [Hypergraph convolution and hypergraph attention](#). *CoRR*, abs/1901.08150.
- Antoine Bordes, Nicolas Usunier, Sumit Chopra, and Jason Weston. 2015. [Large-scale simple question answering with memory networks](#). *CoRR*, abs/1506.02075.
- Antoine Bordes, Nicolas Usunier, Alberto García-Durán, Jason Weston, and Oksana Yakhnenko. 2013. [Translating embeddings for modeling multi-relational data](#). In *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013. Proceedings of a meeting held December 5-8, 2013, Lake Tahoe, Nevada, United States*, pages 2787–2795.
- Danqi Chen, Adam Fisch, Jason Weston, and Antoine Bordes. 2017. [Reading wikipedia to answer open-domain questions](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017, Vancouver, Canada, July 30 - August 4, Volume 1: Long Papers*, pages 1870–1879. Association for Computational Linguistics.
- Rajarshi Das, Manzil Zaheer, Siva Reddy, and Andrew McCallum. 2017. [Question answering on knowledge bases and text using universal schema and memory networks](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017, Vancouver, Canada, July 30 - August 4, Volume 2: Short Papers*, pages 358–365. Association for Computational Linguistics.
- Yifan Feng, Haoxuan You, Zizhao Zhang, Rongrong Ji, and Yue Gao. 2019. [Hypergraph neural networks](#). In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019*, pages 3558–3565. AAAI Press.
- Taher H. Haveliwala. 2002. [Topic-sensitive pagerank](#). In *Proceedings of the Eleventh International World Wide Web Conference, WWW 2002, May 7-11, 2002, Honolulu, Hawaii, USA*, pages 517–526. ACM.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. [Long short-term memory](#). *Neural Computation*, 9(8):1735–1780.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Alexander H. 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 2016, Austin, Texas, USA, November 1-4, 2016*, pages 1400–1409. The Association for Computational Linguistics.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, 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 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, 8-14 December 2019, Vancouver, BC, Canada*, pages 8024–8035.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. [Glove: Global vectors for word representation](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP 2014, October 25-29, 2014, Doha, Qatar, A meeting of SIGDAT, a Special Interest Group of the ACL*, pages 1532–1543. ACL.
- Haitian Sun, Tania Bedrax-Weiss, and William W. 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 2019, Hong Kong, China, November 3-7, 2019*, pages 2380–2390. Association for Computational Linguistics.
- 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, Brussels, Belgium, October 31 - November 4, 2018*, pages 4231–4242. Association for Computational Linguistics.
- Alona Sydorova, Nina Pörner, and Benjamin Roth. 2019. [Interpretable question answering on knowledge bases and text](#). 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 4943–4951. Association for Computational Linguistics.
- Johannes Welbl, Pontus Stenetorp, and Sebastian Riedel. 2018. [Constructing datasets for multi-hop reading comprehension across documents](#). *Trans. Assoc. Comput. Linguistics*, 6:287–302.
- Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and Philip S. Yu. 2020. [A comprehensive survey on graph neural networks](#). *IEEE Transactions on Neural Networks and Learning Systems*, pages 1–21.

- 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 Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 4258–4264. Association for Computational Linguistics.
- Naganand Yadati, Madhav Nimishakavi, Prateek Yadav, Vikram Nitin, Anand Louis, and Partha P. Talukdar. 2019. [Hypergn: A new method for training graph convolutional networks on hypergraphs](#). In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, 8-14 December 2019, Vancouver, BC, Canada*, pages 1509–1520.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William W. Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. 2018. [Hotpotqa: A dataset for diverse, explainable multi-hop question answering](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 2369–2380. Association for Computational Linguistics.
- Wen-tau Yih, Matthew Richardson, Christopher 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, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 2: Short Papers*. The Association for Computer Linguistics.
- 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-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, pages 6069–6076. AAAI Press.