

# Is Graph Structure Necessary for Multi-hop Question Answering?

Nan Shao<sup>†</sup>, Yiming Cui<sup>††</sup>, Ting Liu<sup>‡</sup>, Shijin Wang<sup>†§</sup>, Guoping Hu<sup>†</sup>

<sup>†</sup>State Key Laboratory of Cognitive Intelligence, iFLYTEK Research, China

<sup>‡</sup>Research Center for Social Computing and Information Retrieval (SCIR),  
Harbin Institute of Technology, Harbin, China

<sup>§</sup>iFLYTEK AI Research (Hebei), Langfang, China

<sup>†§</sup>{nanshao, ymcui, sjwang3, gpku}@iflytek.com

<sup>‡</sup>{ymcui, tliu}@ir.hit.edu.cn

## Abstract

Recently, attempting to model texts as graph structure and introducing graph neural networks to deal with it has become a trend in many NLP research areas. In this paper, we investigate whether the graph structure is necessary for multi-hop question answering. Our analysis is centered on HotpotQA. We construct a strong baseline model to establish that, with the proper use of pre-trained models, graph structure may not be necessary for multi-hop question answering. We point out that both graph structure and adjacency matrix are task-related prior knowledge, and graph-attention can be considered as a special case of self-attention. Experiments and visualized analysis demonstrate that graph-attention or the entire graph structure can be replaced by self-attention or Transformers.

## 1 Introduction

Different from single-hop question answering, where the answer can be derived from a single sentence in a single paragraph, more and more studies focus on multi-hop question answering across multiple documents or paragraphs (Welbl et al., 2018; Talmor and Berant, 2018; Yang et al., 2018).

To solve this problem, the majority of existing studies constructed a graph structure according to co-occurrence relations of entities that scattered across multiple sentences or paragraphs. Dhingra et al. (2018) and Song et al. (2018) designed a DAG-styled recurrent layer to model the relations between entities. De Cao et al. (2019) first used GCN (Kipf and Welling, 2017) to tackle entity graph. Qiu et al. (2019) proposed a dynamic entity graph for span-based multi-hop QA. Tu et al. (2019b) extended the entity graph to a heterogeneous graph by introducing document nodes and query nodes.

Previous works argue that a fancy graph structure is a vital part of their models and demonstrate that

by ablation experiments. However, in experiments, we find when we use the pre-trained models in the fine-tuning approach, removing the entire graph structure may not hurt the final results. Therefore, in this paper, we aimed to answer the following question: **How much does graph structure contribute to multi-hop question answering?**

To answer the question above, we choose the widely used multi-hop question answering benchmark, HotpotQA (Yang et al., 2018), as our testbed. We reimplement a graph-based model, Dynamically Fused Graph Network (Qiu et al., 2019), as our baseline model. The remainder of this paper is organized as follows.

- In Section 2, we first describe our baseline model. Then, we show that the graph structure can play an important role only when the pre-trained models are used in a feature-based manner. While the pre-trained models are used in the fine-tuning approach, the graph structure may not be helpful.
- To explain the results, in Section 3.1, we point out that graph-attention (Veličković et al., 2018) is a special case of self-attention. The adjacency matrix based on manually defined rules and the graph structure can be regarded as prior knowledge, which could be learned by self-attention or Transformers (Vaswani et al., 2017).
- In Section 3.2, we design experiments to show when we model text as an entity graph, both graph-attention and self-attention can achieve comparable results. When we treat texts as a sequence structure, only a 2-layer Transformer could achieve similar results as DFGN.
- In Section 3.4, visualized analysis show that there are diverse entity-centered attention patterns exist in pre-trained models, indicating the redundancy of entity-based graph structure.
- Section 4 gives the conclusion.

## 2 The Approach

We choose the widely used multi-hop QA dataset, HotpotQA as our testbed. We reimplement DFGN (Qiu et al., 2019) and modify the usage of the pre-trained model. The model first leverage a retriever to select relevant passages from candidate set and feed them into a graph-based reader. All entities in the entity graph are recognized by an independent NER model.

### 2.1 HotpotQA Dataset

HotpotQA is a widely used large-scale multi-hop QA dataset. There are two different settings in HotpotQA. In *distractor* setting, each example contains 2 gold paragraphs and 8 distractor paragraphs retrieved from Wikipedia. In *full wiki* setting, a model is asked to retrieve gold paragraphs from the entire Wikipedia. In this paper, all experiments are conducted in the *distractor* setting.

### 2.2 Model Description

**Retriever.** We use RoBERTa large model (Liu et al., 2019) to calculate the relevant score between the query and each candidate paragraphs. We filter the paragraphs whose score is less than 0.1, and the maximum number of selected paragraphs is 3. Selected paragraphs are concatenated as context  $\mathbf{C}$ .  
**Encoding Layer.** We concatenate the query  $\mathbf{Q}$  and context  $\mathbf{C}$  and feed the sequence into another RoBERTa large model. The results are further fed into a bi-attention layer (Seo et al., 2016) to obtain the representations from the encoding layer.

**Graph Fusion Block.** Given context representations  $\mathbf{C}_{t-1}$  at hop  $t - 1$ , the tokens representations are passed into a mean-max pooling layer to get nodes representations in entity graph  $\mathbf{H}_{t-1} \in \mathbb{R}^{2d \times N}$ , where  $N$  is the number of entity. After that, a graph-attention layer is applied to update nodes representations in the entity graph:

$$\beta_{i,j}^{(t)} = \text{LeakyReLU}(\mathbf{W}_t^\top [\mathbf{h}_i^{(t-1)}, \mathbf{h}_j^{(t-1)}]) \quad (1)$$

$$\alpha_{i,j}^{(t)} = \frac{\exp(\beta_{i,j}^{(t)})}{\sum_{k \in \mathcal{N}_i} \exp(\beta_{i,k}^{(t)})} \quad (2)$$

$$\mathbf{h}_i^{(t)} = \text{ReLU}(\sum_{k \in \mathcal{N}_i} \alpha_{i,k}^{(t)} \mathbf{h}_k^{(t-1)}) \quad (3)$$

where  $\mathcal{N}_i$  is the set of neighbors of node  $i$ . We follow the same Graph2Doc module as Qiu et al. (2019) to transform the nodes representations into the tokens representations. Besides, there are sev-

Setting	Joint EM	Joint F1
Baseline (Yang et al., 2018)	10.83	40.16
QFE (Nishida et al., 2019)	34.63	59.61
DFGN (Qiu et al., 2019)	33.62	59.82
TAP2 (Glass et al., 2019)	39.77	69.12
HGN (Fang et al., 2019)	43.57	71.03
SAE (Tu et al., 2019a)	<b>45.36</b>	71.45
Our Model	44.67	<b>72.73</b>

Table 1: Results on the test set of HotpotQA.

Setting	Joint EM	Joint F1
Baseline (Fine-tuning)	45.91	<b>73.93</b>
w/o Graph	<b>45.98</b>	73.78
Baseline (Feature-based)	36.45	63.75
w/o Graph	32.26	59.76

Table 2: Ablation of graph structure under different settings.

eral extra modules in the graph fusion block, including query-entity attention, query update mechanism, and weak supervision.

**Prediction Layer.** We follow the same cascade structure as Qiu et al. (2019) to predict the answers and supporting sentences.

**Entity Graph Construction.** We fine-tune a pre-trained BERT base model on the dataset of the CoNLL’03 NER shared task (Tjong Kim Sang and De Meulder, 2003) and use it to extract entities from candidate paragraphs. Connections between entities are defined as following rules:

- Entities with the same mention text in context are connected.
- Entities appear in the same sentence are connected.

### 2.3 Model Results

In Table 1, we show the performance comparison with different models on the blind test set of HotpotQA. Our strong baseline model achieves state-of-the-art results on the official leaderboard.

In order to analyze how much the graph structure contributes to the entire model, we perform a set of ablation experiments. We remove the whole graph fusion block, and the outputs of the pre-trained model are directly fed into the prediction layer. By the reason that the main difference between our baseline model and DFGN is that we use a large pre-trained model in the fine-tuning approach instead of the feature-based approach, we perform the experiments in two different settings.

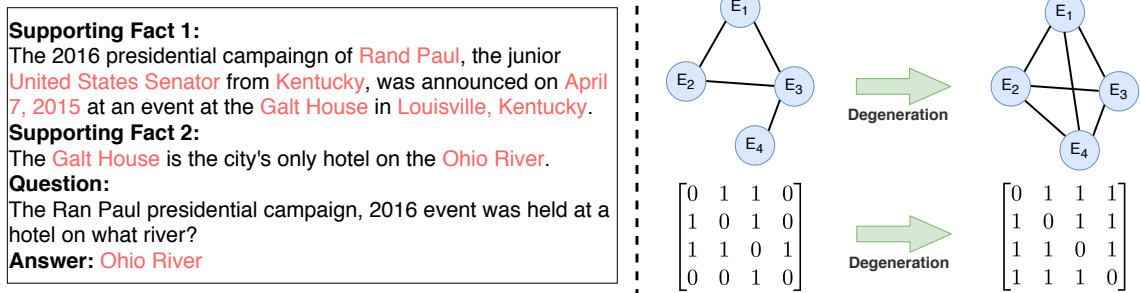


Figure 1: Entities in raw texts are modeled as an entity graph and handled by graph attention networks. When the entity graph are fully connected, a graph-attention layer will degenerate into a vanilla self-attention layer.

The results are shown in Table 2. By using the fine-tuning approach, model with and without graph fusion block can reach almost equal results. When we fix parameters of the pre-trained model, the performances significantly degrade by 9% for EM and 10% for F1. If we further remove graph fusion block, both EM and F1 drop about 4%.

Taken together, only when pre-trained models are used in the feature-based approach, graph neural networks can play an important role. Nevertheless, when pre-trained models are used as a fine-tuning approach, which is a common practice, graph structure does not contribute to the final results. In other words, the graph structure may not be necessary for multi-hop question answering.

### 3 Understanding Graph Structure

Experimental results in Section 2.3 imply that self-attention or Transformer may have superiority in multi-hop question answering. To understand this, in this section, we will first discuss the connection between graph structure, graph-attention, and self-attention. We then verify the hypothesis by experiments and visualized analysis.

#### 3.1 Graph Attention vs. Self Attention

The key to solving the multi-hop question is to find the corresponding entity in the original text through the query. Then one or more reasoning paths are constructed from these start entities toward other identical or co-occurring entities. As shown in Figure 1, previous works usually extract entities from multiple paragraphs and model these entities as an entity graph. The adjacency matrix is constructed by manually defined rules, which usually the co-occurrence relationship of entities. From this point of view, both the graph structure and the adjacency matrix can be regarded as task-related prior knowledge. The entity graph structure restricts the model

can only do reasoning based on entities, and the adjacency matrix assists the model to ignore non-adjacent nodes in a hop. However, it is probably that the model without any prior knowledge can still learn the entity-to-entity attention pattern.

In addition, considering Eq.1-3, it is easy to find that graph-attention has a similar form as self-attention. In forward propagation, each node in the entity graph calculates attention scores with other connected nodes. As shown in Figure 1, graph-attention will degenerate into a vanilla self-attention layer when the nodes in the graph are fully connected. Therefore, the graph-attention can be considered as a special case of self-attention.

#### 3.2 Graph Structure May Not Be Necessary

According to the discussion above, we aimed to evaluate whether the graph structure with an adjacency matrix is superior to self-attention.

To this end, we use the model described in Section 2 as our baseline model. The pre-trained model in the baseline model is used in the feature-based approach. Several different modules are added between the encoding layer and the prediction layer. **Model With Graph Structure.** We apply graph-attention or self-attention on the entity graph and compare the difference in the final results. In order to make a fair comparison, we choose the self-attention that has the same form as graph-attention. The main difference is that the self-attention does not keep an adjacency matrix as prior knowledge and the entities in the graph are fully connected. Moreover, we define that the density of a binary matrix is the percentage of ‘1’ in it. We sort each example in development set by the density of its adjacency matrix and divide them by different quantiles. We evaluate how different density of the adjacency matrix affects the final results.

**Model Without Graph Structure.** In this experi-

Setting	Joint EM	Joint F1
Baseline	32.26	59.76
+ Graph Fusion Block	<b>36.45</b>	63.75
+ Self Attention	35.41	61.77
+ Graph Attention	35.79	61.91
+ Transformer	36.23	<b>63.82</b>

Table 3: Performance comparison in terms of joint EM and F1 scores under different module settings.

Quantile	0.2	0.4	0.6	0.8	1.0	avg
Density	18.7	23.6	29.0	36.8	100	28.8

Table 4: The Adjacency Matrix density at different quantiles.

ment, we verify whether the whole graph structure can be replaced by Transformers. We directly feed the context representations from the encoding layer into the Transformers.

The experimental results are shown in Table 3. Compared with the baseline, the model with the graph fusion block obtains a significant advantage. We add the entity graph with self-attention to the baseline model, and the final results significantly improved. Compared with self-attention, graph-attention does not show a clear advantage. The density of examples at different quantile are shown in Table 4, the adjacency matrix in multi-hop QA is relatively dense, which may causes graph-attention can not make a significant difference. The results of graph-attention and self-attention in the different intervals of density are shown in Figure 2. Despite the different density of the adjacency matrix, graph-attention consistently achieves similar results as self-attention. This signifies that self-attention can learn to ignore irrelevant entities. Besides, examples with a more dense adjacency matrix are simpler for both graph-attention and self-attention, this probably because these adjacency matrices are constructed from shorter documents. Moreover, as shown in Table 3, Transformers show a powerful reasoning ability. Only stacking two layers of the Transformer can achieve comparable results as DFGN.

### 3.3 Training Details

For all experiments in this paper, the number of layers of different modules is two, and the hidden dimensions are set to 300. In feature-based setting, all models are trained for 30 epochs with a batch size of 24. In fine-tuning setting, models are trained

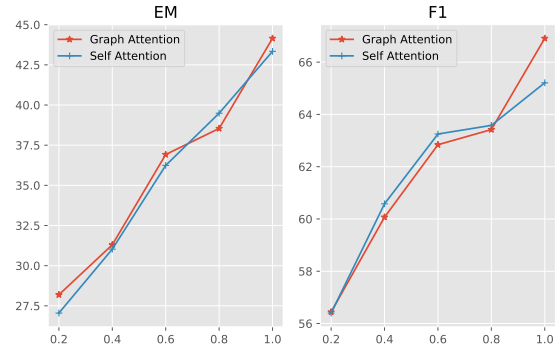


Figure 2: Results of graph-attention and self-attention on examples with different adjacency matrix density.

for 3 epochs with a batch size of 8. The initial learning rate is  $2e-4$  and  $3e-5$  in the feature-based setting and fine-tuning setting respectively.

### 3.4 Entity-centered Attention Pattern in Pre-trained Model

Inspired by Kovaleva et al. (2019), we leverage an approximate method to find which attention head contains entity-centered attention patterns. We employ an NER model to identify tokens belong to a certain entity span. Then, for each attention head in the pre-trained model, we sum the absolute attention weights among those tokens belong to an entity and tokens not belong to an entity. The score of an attention head is the difference between the sum of weights from entities and non-entities tokens. We then average the derived scores over all the examples. Finally, the attention head with the maximum score is the desired head that contains entity-centered attention patterns.

We find four typical attention patterns and visualized it in Figure 3. In case 1-3, we visualized the attention weights of each token attending to the subject entity. In case 4, we visualized the attention weights of each token attending to the last token of the sentence. The results show pre-trained models are pretty skillful at capturing relations between entities and other constituents in a sentence.

**Entity2Entity.** We find entity-to-entity attention pattern is very widespread in pre-trained models. In this case, ‘American Physicist’ and ‘Czech’ attend to ‘Emil Wolf’ with very high attention weights. Such attention pattern plays the same role as graph attention plays.

**Attribute2Entity.** In this case, ‘filmmaker’, ‘film critic’ and ‘teacher’ obtain higher weights, indicating the occupation of ‘Thom Andersen’. Note that

Case1: Entity2Entity

Emil	Wolf	(	born	July	30
,	1922	)	is	a	Czech
born	American	Physicist	who	made	.....

Case2: Attribute2Entity

Thom	Andersen	(	born	1943	,
Chicago	)	is	a	filmmaker	,
film	critic	and	teacher	.	

Case3: Coreference2Entity

Julian	Bolling	(	June	19	,
1966	.....	He	has	won	15
gold	medals	for	Sri	Lanka	at
South	Asian	Games	between	.....	

Case4: Entity2Sentence

WAVN	is	a	Gospel	formatted	broadcast
radio	station	licensed	to	Southaven	,
Mississippi	,	servicing	Metro	Memphis	.

Figure 3: Results of visualized attention patterns in the pre-trained model.

these tokens are not part of a certain entity, hence deem to be ignored by graph structure.

**Coreference2Entity.** We also find that coreference will not make the pre-trained model confused. In case 3, the entity ‘Sir Lanka’ in second sentence attends to ‘Julian Bolling’ in the first sentence, which means the pre-trained model understand ‘He’ refers to ‘Julian Bolling’ even though they belong to different sentences.

**Entity2Sentence.** We find many entities attend to the last token of sentence. In the prediction layer, the representations of the first and last token in a sentence are combined to determine whether a particular sentence is a supporting fact. Therefore, we suppose this is another attention pattern that entities attend to the whole sentence.

It is obvious that graph attention can not cover the last three attention patterns. Therefore, we draw a conclusion that self attention has advantages on generality and flexibility.

## 4 Conclusions

This study set out to investigate whether graph structure is necessary for multi-hop QA and what role it plays. We established that with the proper use of pre-trained models, graph structure may not be necessary. In addition, we point out the adjacency matrix and the graph structure can be regarded as some kind of task-related prior knowledge. Experiments and visualized analysis demonstrate both graph-attention and graph structure can be replaced by self-attention or Transformers. Our results suggest that future works introducing graph

structure into NLP tasks should explain their necessity and superiority.

## Acknowledgments

We would like to thank all anonymous reviewers for their thorough reviewing and providing insightful comments to improve our paper.

## References

- Nicola De Cao, Wilker Aziz, and Ivan Titov. 2019. [Question answering by reasoning across documents with graph convolutional networks](#). 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 2306–2317, Minneapolis, Minnesota. Association for Computational Linguistics.
- Bhuwan Dhingra, Qiao Jin, Zhilin Yang, William Cohen, and Ruslan Salakhutdinov. 2018. [Neural models for reasoning over multiple mentions using coreference](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 42–48, New Orleans, Louisiana. Association for Computational Linguistics.
- Yuwei Fang, Siqi Sun, Zhe Gan, Rohit Pillai, Shuohang Wang, and Jingjing Liu. 2019. Hierarchical graph network for multi-hop question answering. *arXiv preprint arXiv:1911.03631*.
- Michael Glass, Alfio Gliozzo, Rishav Chakravarti, Anthony Ferritto, Lin Pan, GP Bhargav, Dinesh Garg, and Avirup Sil. 2019. Span selection pre-training for question answering. *arXiv preprint arXiv:1909.04120*.
- Thomas N. Kipf and Max Welling. 2017. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations (ICLR)*.
- Olga Kovaleva, Alexey Romanov, Anna Rogers, and Anna Rumshisky. 2019. [Revealing the dark secrets of BERT](#). 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 4365–4374, Hong Kong, China. Association for Computational Linguistics.
- 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. *arXiv preprint arXiv:1907.11692*.

- Kosuke Nishida, Kyosuke Nishida, Masaaki Nagata, Atsushi Otsuka, Itsumi Saito, Hisako Asano, and Junji Tomita. 2019. [Answering while summarizing: Multi-task learning for multi-hop QA with evidence extraction](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2335–2345, Florence, Italy. Association for Computational Linguistics.
- Lin Qiu, Yunxuan Xiao, Yanru Qu, Hao Zhou, Lei Li, Weinan Zhang, and Yong Yu. 2019. [Dynamically fused graph network for multi-hop reasoning](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6140–6150, Florence, Italy. Association for Computational Linguistics.
- Minjoon Seo, Aniruddha Kembhavi, Ali Farhadi, and Hannaneh Hajishirzi. 2016. Bidirectional attention flow for machine comprehension. In *International Conference on Learning Representations (ICLR)*.
- Linfeng Song, Zhiguo Wang, Mo Yu, Yue Zhang, Radu Florian, and Daniel Gildea. 2018. Exploring graph-structured passage representation for multi-hop reading comprehension with graph neural networks. *arXiv preprint arXiv:1809.02040*.
- Alon Talmor and Jonathan Berant. 2018. [The web as a knowledge-base for answering complex questions](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 641–651, New Orleans, Louisiana. Association for Computational Linguistics.
- Erik F. Tjong Kim Sang and Fien De Meulder. 2003. [Introduction to the CoNLL-2003 shared task: Language-independent named entity recognition](#). In *Proceedings of the Seventh Conference on Natural Language Learning at HLT-NAACL 2003*, pages 142–147.
- Ming Tu, Kevin Huang, Guangtao Wang, Jing Huang, Xiaodong He, and Bowen Zhou. 2019a. Select, answer and explain: Interpretable multi-hop reading comprehension over multiple documents. *arXiv preprint arXiv:1911.00484*.
- Ming Tu, Guangtao Wang, Jing Huang, Yun Tang, Xiaodong He, and Bowen Zhou. 2019b. [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*, pages 2704–2713, Florence, Italy. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008.
- Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. [Graph attention networks](#). In *International Conference on Learning Representations*.
- Johannes Welbl, Pontus Stenetorp, and Sebastian Riedel. 2018. [Constructing datasets for multi-hop reading comprehension across documents](#). *Transactions of the Association for Computational Linguistics*, 6:287–302.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William 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*, pages 2369–2380, Brussels, Belgium. Association for Computational Linguistics.