

A Graph-Guided Reasoning Approach for Open-ended Commonsense Question Answering

Zhen Han^{1*}, Yue Feng¹, Mingming Sun¹

¹Cognitive Computing Lab, Baidu Research
hanzhen02111@163.com, sunmingming01@baidu.com

Abstract

Recently, end-to-end trained models for multiple-choice commonsense question answering (QA) have delivered promising results. However, such question-answering systems cannot be directly applied in real-world scenarios where answer candidates are not provided. Hence, a new benchmark challenge set for open-ended commonsense reasoning (OpenCSR) has been recently released, which contains natural science questions without any predefined choices. On the OpenCSR challenge set, many questions require implicit multi-hop reasoning and have a large decision space, reflecting the difficult nature of this task. Existing work on OpenCSR sorely focuses on improving the retrieval process, which extracts relevant factual sentences from a textual knowledge base, leaving the important and non-trivial reasoning task outside the scope. In this work, we extend the scope to include a reasoner that constructs a question-dependent open knowledge graph based on retrieved supporting facts and employs a sequential subgraph reasoning process to predict the answer. The subgraph itself can be seen as a concise and compact graphical explanation of the prediction. Experiments on two OpenCSR datasets show that the proposed model achieves great performance on benchmark OpenCSR datasets.

1 Introduction

Commonsense reasoning has long been considered an essential topic in artificial intelligence. Most approaches work on the setting of *multiple-choice question answering* [8, 3], which selects an answer choice by scoring the question-choice pairs. However, the multiple-choice setting is not applicable in many real-world scenarios since many question-answering tasks do not provide answer candidates. As a step towards making commonsense reasoning

research more realistic and useful, open-ended commonsense reasoning (OpenCSR) has been introduced [9], which explores a commonsense knowledge corpus to answer commonsense questions. OpenCSR often requires multi-hop reasoning, i.e., the model should conclude the answer by reasoning over two or more facts from the knowledge corpus, which makes this task much more challenging. Lin et al. [9] proposed a retrieval-based method, called *DrFact*, by combining the maximum inner product search and symbolic links between facts. However, *DrFact* does not put much effort on the reasoning module to re-rank the retrieved facts. To this end, we proposed an integrated subgraph reasoning approach for OpenCSR with end-to-end learning, which iteratively employs a retriever to extract question-relevant facts from a knowledge corpus and a reasoner over the extracted facts. Given a commonsense question, the proposed approach applies DPR [6] to extract relevant facts from a textual knowledge corpus, converts the retrieved natural language facts into a graph-structured format using Open Information Annotation (OIA) [10] and performs subgraph reasoning on the constructed joint OIA-graph using a multi-relational graph attention network. Specifically, the reasoner first performs entity linking from the giving question to the joint OIA-graph. Then it starts from the linked entities (nodes), and iteratively samples relevant edges with a pruning procedure to form an enclosing subgraph around the question. The reasoning procedure takes into account both structural information, i.e., graph structure of the joint OIA graph, and semantic information, i.e., language representation of questions and facts. After several rounds of retrieval and pruning, the model predicts the answer from the concepts in the subgraph.

Our contributions are as follows: (1) we investigate how to perform a cooperative retrieval-and-reasoning in open-ended commonsense question answering. To the best of our knowledge, our

*work done while internship at Baidu.




Figure 1: **(a)** The OIA graph of "Plants supply the fungi with carbohydrates, in return, making it a symbiotic relationship." There are two types of nodes: constant and predicate. Constant nodes are simple nominal phrases while predicate nodes include simple verbal phrases and prepositional phrases. Edges in OIA graphs are labeled. *pred.arg.n* denotes the *n*-th arguments of a predicate node, *mod* indicates the modification, and *as:pred.arg.n* expresses an reversed relation of *pred.arg.n*. **(b)** The joint OIA graph consists of two factoid sentences that share the concepts "fungi" and "symbiotic relationship".

work is the first retrieve-and-reasoning approach for OpenCSR. (2) We present experimental results that show our model achieves great results on the benchmark OpenCSR dataset with an ablation study demonstrating the performance gain of integration structural information and semantic information. (3) The proposed method can potentially homogenize structured, i.e., knowledge base, and unstructured commonsense knowledge, i.e., textual corpus for answering open-ended commonsense questions since it can unify both knowledge formats into a graph-structured format.

2 Related Work

Commonsense Reasoning Traditional commonsense reasoning (CSR) techniques are mainly designed for multiple-choice QA. For instance, to independently score each decision, KagNet [8] and MHGRN [3] both leverage external commonsense knowledge graphs as structural priors. Although effective in selecting the best response for a multiple-choice question, these techniques are less useful for real-world situations because answer candidates are frequently unavailable. By fine-tuning a text-to-text transformer, UnifiedQA [7] generated answers to questions. However, a drawback of multiple-choice QA models is that they do not provide intermediate explanations for their answers, making them less suitable in many real-world scenarios. Lin et al. [9] introduced the open-ended commonsense reasoning and proposed DrFact to directly retrieve relevant facts, and then use the concepts mentioned in the top-ranked facts as answer predictions.

Subgraph Reasoning Many recent works learn representations of localized subgraphs. Alsentzer et al. [1] introduced a subgraph neural network to learn disentangled subgraph representations using a novel subgraph routing mechanism. Teru et al. [11] proposed a graph neural network that reasons over

local subgraph structures and performs inductive relation predictions. Han et al. [4] developed an explainable reasoning framework for forecasting future links on temporal knowledge graphs by employing a sequential reasoning process over local subgraphs.

3 Our Approach

Retrieving Relevant Facts Following the dense passage retrieval work [6], we use a bi-encoder transformer architecture that learns to maximize the inner product of the representation of a question and the relevant factual sentences from the knowledge corpus containing correct answers to the given question.

Constructing Question-dependent Joint OIA-Graph Following the steps in [10], we convert each retrieved factual sentence into an OIA-graph as shown in Figure 1a. For each node in an OIA graph, we link it with nodes in the OIA graphs of other sentences that include the same concept. We label this kind of link as *shared concepts*. As shown in Figure 1b, the factoid sentences "Plants supply the fungi with carbohydrates, in return, making it a symbiotic relationship." and "Fungi participate in symbiotic relationships to obtain their food." shares the same concepts "fungi" and "symbiotic relationship". Then, we construct a joint OIA-graph \mathcal{G}_{joint} by linking nodes that share the same concepts in different OIA graphs.

Subgraph Reasoning on the Joint OIA-Graph Inspired by [5], we conduct reasoning on a dynamically expanded inference graph \mathcal{G}_{inf} extracted from the joint OIA-graph. Given a commonsense question q , we build an initial inference graph via **entity linking** between the question q and the joint OIA-graph. We find all nodes of \mathcal{G}_{joint} that share the same concepts as q includes. We set such OIA-




Figure 2: Model Architecture.

nodes to be the initial nodes of the inference graph \mathcal{G}_{inf} . The inference graph expands by sampling one-hop neighbors of initial nodes in \mathcal{G}_{joint} . Besides, we propose a semantic-following operation to build skip connections between the initial nodes and their multi-hop neighbors. Taking a node v in \mathcal{G}_{inf} as an example, we compute the inner-product similarity between its representation and the representation of other nodes in \mathcal{G}_{joint} obtained by the retrieval and add the top K nodes into \mathcal{G}_{inf} by linking them with v . The contribution of the semantic-following has two folds: 1) It speeds up the reasoning process and broadens the receptive field of the subgraph reasoner by adding skip connections between multi-hop neighbors; 2) It allows the subgraph reasoner to take into account both semantic-relevant and symbolic-linked nodes regarding a given question. Next, we feed \mathcal{G}_{inf} into a relational graph attention layer that takes node embedding as the input, computes an attention score for each edge indicating the relevance to the given question, and produces a question-dependent representation for each node using message passing. Instead of treating all neighbors with equal importance in the message passing, we take the question information into account and assign varying importance levels to each neighbor by calculating the following question-dependent attention score:

$$e_{vu}^l(q, p_k) = \mathbf{W}_s^l(\mathbf{h}_v^{l-1} \parallel \mathbf{p}_k^{l-1} \parallel \mathbf{h}_q^{l-1})^T \mathbf{W}_t^l(\mathbf{h}_u^{l-1} \parallel \mathbf{p}_k^{l-1} \parallel \mathbf{h}_q^{l-1}), \quad (1)$$

where $e_{vu}^l(q, p_k)$ is the attention score of the edge (v, p_k, u) regarding the question q , p_k corresponds to the edge type between the source node v and the target node u , \mathbf{W}_s^l and \mathbf{W}_t^l are two weight matrices for capturing the dependencies between question representations and source node features specified for source node and target node, respectively. \mathbf{p}_k is the edge embedding indicating the relationship between u and v . \mathbf{h}_v^{l-1} denotes the hidden representation of the node v at the $(l-1)^{th}$ inference step.

When $l = 1$, i.e., for the first layer, \mathbf{h}_v^0 is the aggregated token representation from Bert. Then, we compute the normalized attention score $\alpha_{vu}^l(q, p_k)$ using the *softmax function*. Once obtained, we aggregate the representations of the sampled neighbors of node v denoted as $\tilde{\mathcal{N}}_v$ and weight them using the normalized attention scores, which are written as

$$\mathbf{h}_v^l(q) = \sum_{u \in \tilde{\mathcal{N}}_v} \alpha_{vu}^l(q, p_k) \mathbf{h}_u^{l-1}(q). \quad (2)$$

Answer Prediction We compute the plausibility score $s_{v,q}^l$ of node v to be the answer of question q at the l^{th} inference step as follows:

$$s_{v,q}^l = s_{mips}(\mathbf{q}, \mathbf{f}_v) + \sum_{u \in \tilde{\mathcal{N}}_v} \sum_{p_k \in \mathcal{P}_{uv}} \alpha_{uv}^l(q, p_k) a_{u,q}^{l-1}, \quad (3)$$

where $s_{mips}(\mathbf{q}, \mathbf{f}_v)$ denotes the relevance score of the retrieved fact f_v , which mentions the node v , regarding the question q by the maximum inner product search. Since the same concept may appear in different nodes in the inference graph, we aggregate the plausibility score of nodes that share the same concept to assign each concept a unique attention score:

$$s_{c_i,q}^l = g(s_{v,q}^l | v(c) = c_i), \quad \text{for } v \in \mathcal{V}_{\mathcal{G}_{inf}}, \quad (4)$$

where $s_{c_i,q}^l$ denotes the plausibility score of concept c_i , $\mathcal{V}_{\mathcal{G}_{inf}}$ is the set of nodes in inference graph \mathcal{G}_{inf} . $v(c)$ represents the concept included in node v , and $g(\cdot)$ represents a score aggregation function. Here we use the maximum function.

Inference Graph Expansion and Pruning After several iterations of expansion, the inference graph \mathcal{G}_{inf} would grow rapidly and cover almost all nodes. To prevent the inference graph from exploding, we reduce the graph size by pruning the edges with a small plausibility score and keeping the edges with K largest contribution scores. After running

Datasets	ARC-Open				OBQA-Open			
	H@50	H@100	R@50	R@100	H@50	H@100	R@50	R@100
DPR	68.67	78.62	28.93	38.63	54.47	67.73	15.17	22.34
DrKIT	67.63	77.89	27.57	37.29	61.74	75.92	18.18	27.10
DrFact	71.60	80.38	31.48	40.93	69.01	80.03	21.27	30.32
Our model	72.76	80.38	31.09	40.24	62.30	73.80	18.11	26.83

Table 1: Results of the Hit@K and Rec@K (K=50/100) in % on OpenCSR.

L inference steps, the model selects the concept with the highest plausibility score in \mathcal{G}_{inf} as the answer to the given question, where the inference graph itself serves as a graphical explanation.

Loss Function We use the binary cross-entropy as the loss function, which is

$$\mathcal{L} = -\frac{1}{|\mathcal{Q}|} \sum_{q \in \mathcal{Q}} \frac{1}{|\mathcal{C}_q^{inf}|} \sum_{c_i \in \mathcal{C}_q^{inf}} (y_{c_i,q} \log(\frac{s_{c_i,q}^L}{\sum_{c_j \in \mathcal{C}_q^{inf}} s_{c_j,q}^L}) + (1 - y_{c_i,q}) \log((1 - \frac{s_{c_i,q}^L}{\sum_{c_j \in \mathcal{C}_q^{inf}} s_{c_j,q}^L}))),$$

where \mathcal{C}_q^{inf} represents the set of concepts in the inference graph of the question q , $y_{c_i,q}$ represents the binary label that indicates whether c_i is an answer for q , and \mathcal{Q} denotes the question set. $s_{c_i,q}^L$ denotes the plausibility score of concept c_i at the final inference step.

4 Experiments

Fact corpus and concept vocabulary Following settings in [9], GenericsKB-Best corpus serves as the main commonsense knowledge source that contains 1,025,413 unique facts. All sentences in the corpus are provided with concepts, which are frequent noun chunks, using the spaCy toolkit. There are 80,524 concepts in total.

Datasets and evaluation metrics We evaluate our model on two benchmark open-ended commonsense reasoning datasets, i.e., ARC-Open and OBQA-Open [9], that contain 6600 and 5288 questions, separately. Every question could be answered using various concepts, where the average answer is 6.8 and 7.7 in ARC-Open and OBQA-Open. Each dataset provides the set of true answer concepts for each question. We use two metrics, Hits@K and Recall@K, where Hits@K denotes the percentage of times that at least one true concept appears in the top k of ranked concepts.

Experimental Results We compare our model with DPR [6], DrKIT [2], and DrFact [9]. Recall

that our model applies DPR as the retriever so it is a straightforward baseline. And DrFact is the strongest baseline in OpenCSR. As shown in Table 1, our model outperforms DPR and DrKIT on ARC-Open and achieves on-par performance as DrFact. All results are averaged over three trials. We provide implementation details in Appendix B and attach the source code in the supplementary material.

Datasets	ARC-Open			
	H@50	H@100	R@50	R@100
Model w/ SC	72.76	80.38	31.09	40.24
Model w/o SC	71.74	79.65	30.56	39.75

Table 2: Ablation Study on ARC-Open: we investigate the gain of adding skip connections (SC) to semantically relevant multi-hop neighbors.

Ablation Study Recall that the proposed sub-graph reasoner takes into account both the structurally linked one-hop neighbors and semantically relevant multi-hop neighbors by expanding the inference graph \mathcal{G}_{inf} . Table 2 shows an ablation study in that we disable the reasoner to add semantically relevant multi-hop neighbors while inference graph expansion called *Model w/o SC*, demonstrating the performance gain of integrating both structural and semantic information.

5 Conclusion

We present a novel graph-guided neural symbolic commonsense reasoning approach for the open-ended commonsense reasoning task. Specifically, The proposed method integrates both structural dependency information between facts and semantic information by constructing an open information annotation graph and employing a semantic-following operation. The model achieved state-of-the-art performance on two benchmark datasets while being more interpretable.

Limitations

The proposed model performs a sequential reasoning process, and thus, may cause long inference time when answering requires a quite long multi-hop reasoning chain. Besides, the "share-link", which connects nodes that share the same concept, would have a significant amount in some datasets and make the underlying graph much denser. It would make it difficult for the model to decide the expansion direction of the subgraph.

References

- [1] Emily Alsentzer, Samuel Finlayson, Michelle Li, and Marinka Zitnik. Subgraph neural networks. *Advances in Neural Information Processing Systems*, 33:8017–8029, 2020.
- [2] Bhuwan Dhingra, Manzil Zaheer, Vidhisha Balachandran, Graham Neubig, Ruslan Salakhutdinov, and William W Cohen. Differentiable reasoning over a virtual knowledge base. *arXiv preprint arXiv:2002.10640*, 2020.
- [3] Yanlin Feng, Xinyue Chen, Bill Yuchen Lin, Peifeng Wang, Jun Yan, and Xiang Ren. 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, Online, November 2020. Association for Computational Linguistics.
- [4] Zhen Han, Peng Chen, Yunpu Ma, and Volker Tresp. Explainable subgraph reasoning for forecasting on temporal knowledge graphs. In *International Conference on Learning Representations*, 2020.
- [5] Zhen Han, Peng Chen, Yunpu Ma, and Volker Tresp. xerte: Explainable reasoning on temporal knowledge graphs for forecasting future links. *arXiv preprint arXiv:2012.15537*, 2020.
- [6] Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. Dense passage retrieval for open-domain question answering. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6769–6781, Online, November 2020. Association for Computational Linguistics.
- [7] Daniel Khashabi, Sewon Min, Tushar Khot, Ashish Sabharwal, Oyvind Tafjord, Peter Clark, and Hananeh Hajishirzi. UNIFIEDQA: Crossing format boundaries with a single QA system. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1896–1907, Online, November 2020. Association for Computational Linguistics.
- [8] Bill Yuchen Lin, Xinyue Chen, Jamin Chen, and Xiang Ren. 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, Hong Kong, China, November 2019. Association for Computational Linguistics.
- [9] Bill Yuchen Lin, Haitian Sun, Bhuwan Dhingra, Manzil Zaheer, Xiang Ren, and William W Cohen. Differentiable open-ended commonsense reasoning. *arXiv preprint arXiv:2010.14439*, 2020.
- [10] Mingming Sun, Wenyue Hua, Zoey Liu, Xin Wang, Kangjie Zheng, and Ping Li. A predicate-function-argument annotation of natural language for open-domain information expression. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2140–2150, 2020.
- [11] Komal Teru, Etienne Denis, and Will Hamilton. Inductive relation prediction by subgraph reasoning. In *International Conference on Machine Learning*, pages 9448–9457. PMLR, 2020.

Appendices

A The license of the Artifacts

The datasets we used in this work are proposed by Lin et al. [9] which is licensed under the MIT License.

About the model we proposed in this work, we will release it and give the license, copyright information, and terms of use once the paper gets accepted.

B Implementation

We tune the hyperparameters of our models using the random search and report the best configuration in the source code in the supplementary material. The training costs 73 GPU hours on the ARC-Open dataset and 50 GPU hours on the OBQA-Open dataset with NVIDIA A40 instance.