Commonsense Question Answerning with KB and Text

Peifeng Wang

1 Related Work

Both structured knowledge stored in a knowledge base and unstructured text from large web corpus proved to be valuable sources for question answering. Whether structured knowledge or unstructured text plays a dominant role in solving the QA problems mainly depends on the specific scenarios. For answering KB-based QA problems, i.e., questions concerning attributes of the entities, a large number of works dedicate efforts in the field of KBQA. Typically, this line of approaches require understanding the questions in natural language and translate them into KB queries either in formal SQL language (Berant et al., 2013; Cai and Yates, 2013; Kwiatkowski et al., 2013) or latent representation (Bordes et al., 2014; Dong et al., 2015; Yih et al., 2015) for searching plausible entities on KB as answers. For answering text-based QA problems, i.e., questions asking for text-span as answers, works in this field majorly fall in the reading comprehension-style (Seo et al., 2016; Shen et al., 2017; Yu et al., 2018). These approaches need to identify answers span either from some given articles or open-domain documents. Still, it could happen that in either scenario, relying on a single source might not be sufficient to solve the problem completely. On the one hand, KB is known to be extremely incomplete and usually facts in KB fail to cover the knowledge necessary for some QA datasets. On the other hand, while having high coverage, text could lead to more difficulties for locating the answers due to its unstructured nature. Therefore, a certain number of works began to look into methods which leverage both sources to conduct QA tasks. According to how the knowledge from both sizes cooperates, we list out three categories of these works which 1) supplement inference over KB with text, 2) supplement inference over text with KB, or 3) fusing knowledge from KB and text jointly.

1.1 Text to KB

In order to address the incompleteness of KB, several works augment their models with external evidence from text data. For some of them, text is only used as additional feature to enhance the inference over KB (Krishnamurthy and Mitchell, 2012; Reddy et al., 2014; Choi et al., 2015; Savenkov and Agichtein, 2016; Lin et al., 2019). They utilize external text to better understand the questions as well as enrich the features for candidate answers. Recent work (Fu et al., 2019) also make use of corpus for extracting new facts to complete KB during inference.

As methoned above, these methods are better at compositional reasoning over KB which unstructured text do not support (Das et al., 2017), and are greatly improved when enhanced by text evidence. However, when faced with more open-domain questions, evidence from text might be more useful and should serve as the main contribution instead of a complementing role. Moreover, they neglect the other side where structured knowledge could help inference over text.

1.2 KB to Text

There also exist some works investigating the reverse direction, i.e., leveraging KB to improve inference over text. For example, the work in Sun et al. 2015 links each candidate answer in search text to the entities in KB in order to get their semantic feature. Further, Xiong et al. 2019 employs gating mechanism to incorporate necessary strucutred knowledge to better encode questions and passages.

While these methods make up the shortage of KB-oriented counterparts, the obvious limitation is that the factoid knowledge in KB is not consulted to obtain answers directly. Likewise, they also omit the possible benefits brought by the text-to-KB line of approaches.

1.3 Fusing KB and Text

Limited attention is drawn to exploit evidence from KB and text jointly for integral reasoning. Early works utilizing both sources adopt a late fusion strategy. They either aggregate predictions which are grounded independently from each size (Ferrucci et al., 2010; Baudiš, 2015), or simply unify structured and unstructured knowledge with universal schema and feed them to memory network as input. As pointed out by Sun et al. 2018, this strategy is sub-optimal, as models have limited ability to aggregate evidence across the different sources and ignore the rich inter relations between both sizes. To bridge these gaps, Sun et al. 2018 adopts an early fusion strategy instead. They firstly construct a question subgraph to incorporate both KB and corpus via entity links. Then they propose heterogeneous update rules to fuse knowledge from different nodes. Lv et al. 2019 adopts a similar strategy to construct a graph from both sources but their method to fuse the heterogeneous knowledge is twisted. Firstly, nodes from both sizes are presorted as sequences and concatenated into one single input of a language model which generates a sequence representation. Then graph neural networks are used to generate representation for the whole graph. Finally, both the sequence and graph representations are used to compute the prediction score.

To some extent, these works step further to exploit both KB and text in a more unified way than the works introduced in the previous two subsections do. Therefore, evidence from both sizes could be considered jointly to better answer the question of any kind. Still, they emphasize more on relying question to select useful evidence from KB and text. The interaction of both sizes is fulfill only by knowledge fusion. Possible guidance from one size to encode the other size is not explicitly investigated.

References

Petr Baudiš. 2015. Yodaqa: a modular question answering system pipeline. In *POSTER 2015-19th International Student Conference on Electrical Engineering*, pages 1156–1165.

Jonathan Berant, Andrew Chou, Roy Frostig, and Percy Liang. 2013. Semantic parsing on freebase from question-answer pairs. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1533–1544.

Antoine Bordes, Sumit Chopra, and Jason Weston. 2014. Question answering with subgraph embeddings. *arXiv preprint arXiv:1406.3676*.

Qingqing Cai and Alexander Yates. 2013. Large-scale semantic parsing via schema matching and lexicon extension. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 423–433.

Eunsol Choi, Tom Kwiatkowski, and Luke Zettlemoyer. 2015. Scalable semantic parsing with partial ontologies. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1311–1320.

Rajarshi Das, Manzil Zaheer, Siva Reddy, and Andrew McCallum. 2017. Question answering on knowledge bases and text using universal schema and memory networks. *arXiv preprint arXiv:1704.08384*.

Li Dong, Furu Wei, Ming Zhou, and Ke Xu. 2015. Question answering over freebase with multi-column convolutional neural networks. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 260–269.

David Ferrucci, Eric Brown, Jennifer Chu-Carroll, James Fan, David Gondek, Aditya A Kalyanpur, Adam Lally, J William Murdock, Eric Nyberg, John Prager, et al. 2010. Building watson: An overview of the deepqa project. *AI magazine*, 31(3):59–79.

Cong Fu, Tong Chen, Meng Qu, Woojeong Jin, and Xiang Ren. 2019. Collaborative policy learning for open knowledge graph reasoning. *arXiv preprint arXiv:1909.00230*.

Jayant Krishnamurthy and Tom M Mitchell. 2012. Weakly supervised training of semantic parsers. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 754–765. Association for Computational Linguistics.

Tom Kwiatkowski, Eunsol Choi, Yoav Artzi, and Luke Zettlemoyer. 2013. Scaling semantic parsers with on-the-fly ontology matching. In *Proceedings of the 2013 conference on empirical methods in natural language processing*, pages 1545–1556.

Bill Yuchen Lin, Xinyue Chen, Jamin Chen, and Xiang Ren. 2019. Kagnet: Knowledge-aware graph networks for commonsense reasoning. *arXiv preprint arXiv:1909.02151*.

Shangwen Lv, Daya Guo, Jingjing Xu, Duyu Tang, Nan Duan, Ming Gong, Linjun Shou, Daxin Jiang, Guihong Cao, and Songlin Hu. 2019. Graph-based reasoning over heterogeneous external knowledge for commonsense question answering. *arXiv preprint arXiv:1909.05311*.

Siva Reddy, Mirella Lapata, and Mark Steedman. 2014. Large-scale semantic parsing without question-answer pairs. *Transactions of the Association for Computational Linguistics*, 2:377–392.

Denis Savenkov and Eugene Agichtein. 2016. When a knowledge base is not enough: Question answering over knowledge bases with external text data. In *Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval*, pages 235–244. ACM.

Minjoon Seo, Aniruddha Kembhavi, Ali Farhadi, and Hannaneh Hajishirzi. 2016. Bidirectional attention flow for machine comprehension. *arXiv preprint arXiv:1611.01603*.

Yelong Shen, Po-Sen Huang, Jianfeng Gao, and Weizhu Chen. 2017. Reasonet: Learning to stop reading in machine comprehension. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1047–1055. ACM.

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. *arXiv preprint arXiv:1809.00782*.

Huan Sun, Hao Ma, Wen-tau Yih, Chen-Tse Tsai, Jingjing Liu, and Ming-Wei Chang. 2015. Open domain question answering via semantic enrichment. In *Proceedings of the 24th International Conference on World Wide Web*, pages 1045–1055. International World Wide Web Conferences Steering Committee.

Wenhan Xiong, Mo Yu, Shiyu Chang, Xiaoxiao Guo, and William Yang Wang. 2019. Improving question answering over incomplete kbs with knowledge-aware reader. *arXiv preprint arXiv:1905.07098*.

Scott Wen-tau Yih, Ming-Wei Chang, Xiaodong He, and Jianfeng Gao. 2015. Semantic parsing via staged query graph generation: Question answering with knowledge base.

Adams Wei Yu, David Dohan, Minh-Thang Luong, Rui Zhao, Kai Chen, Mohammad Norouzi, and Quoc V Le. 2018. Qanet: Combining local convolution with global self-attention for reading comprehension. *arXiv preprint arXiv:1804.09541*.