

Learning to Transform, Combine, and Reason in Open Domain Question Answering*

Mostafa Dehghani^{1,2}, Hosein Azarbonyad^{1,3}, Jaap Kamps¹, and Maarten de Rijke¹

¹ University of Amsterdam

² Google AI

³ KLM Research

1 Extended Abstract

Users seek direct answers to complex questions from large open-domain knowledge sources like the Web. Open-domain question answering has become a critical task to be solved for building systems that help address users' complex information needs. Most open-domain question answering systems use a search engine to retrieve a set of candidate documents, select one or a few of them as context, and then apply reading comprehension models to extract answers. Some questions, however, require taking a broader context into account, e.g., by considering low-ranked documents that are not immediately relevant, combining information from multiple documents, and reasoning over multiple facts from these documents to infer the answer. In this paper, we propose a deep learning model based on the Transformer architecture that is able to efficiently operate over a larger set of candidate documents by effectively combining the evidence from these documents during multiple steps of reasoning, while it is robust against noise from low-ranked non-relevant documents included in the set.

For example, in Figure 1, in order to infer the correct answer to the question: “Who is the Spanish artist, sculptor and draughtsman famous for co-founding the Cubist movement?” given the top-ranked document, a reading comprehension system most likely will extract “Georges Braque” as the answer, which is not the correct answer. In this example, in order to infer the correct answer, one has to go down the ranked list, gather and encode facts, even those that are not immediately relevant to the question, like “Malaga is a city in Spain,” which can be inferred from a document at rank 66, and then in a multi-step reasoning process, infer some new facts, including “Picasso was a Spanish artist” given documents at ranks 12 and 66, and “Picasso, who was a Spanish artist, co-founded the Cubist” given the previously inferred fact and the document ranked third. In this example, and in general in many cases in open-domain question answering, a piece of information in a low-ranked document that is not immediately relevant to the question, may be useful to fill in the blanks and complete information extracted from the top relevant documents and eventually support inferring the correct answer.

In this paper, we propose TraCRNet (pronounced *Tracker Net*) to improve open-domain question answering by explicitly operating on a larger set of candidate documents during the whole question answering process and learning how to aggregate and

* This is an extended abstract of Dehghani et al. [1].

Copyright © 2019 for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).




Fig. 1: Example complex question answering that requires that information from multiple documents be combined and some amount of reasoning over the information extracted from those documents. (Best viewed in color.)

reason over information from these documents in an effective way while trying not to be distracted by noisy documents. Given the candidate documents and the question, to generate the answer, TraCRNet first **T**ransforms them into vectors by applying a stack of Transformer blocks with self-attention over words in each document in a layer called *Input Encoding*. Then, it updates the learned representations from the first stage by **C**ombining and enriching them through a multihop **R**easoning process by applying multiple steps of the Universal Transformer in a layer called *Multihop Reasoning*.

Returning to the example in Figure 1, after learning representations for each top-ranked document and the question, TraCRNet updates them by applying multiple steps of the Universal Transformer. Given the self-attention mechanism and inductive bias of the Universal Transformer, in the first step, TraCRNet can update the representation of document D#12 by attending to D#66 (as they are related by both mentioning Malaga) and augment the information in D#12 with the fact that “Malaga is city in Spain,” so the updated vector of D#12 has the fact that “Picasso is a Spanish artist” encoded in itself. Then, in the next step of reasoning, TraCRNet can update the representation of D#3 by attending over the vector representing D#12 estimated in the previous step, and enrich the information in D#3 with the fact that “Picasso is a Spanish artist,” and the updated vector of D#3 has the fact that “Picasso, who was a Spanish artist co-founded Cubism” encoded in it. After that, during answer generation, the decoder can attend to the final vector representing D#3 and give the correct answer.

We employ TraCRNet on two public open-domain question answering datasets, SearchQA and Quasar-T, and achieve results that meet or exceed the state-of-the-art.

Acknowledgments This research was supported in part by the Netherlands Organization for Scientific Research through the *Exploratory Political Search* project (ExPoSe, NWO CI # 314.99.108). All content represents the opinion of the authors, which is not necessarily shared or endorsed by their respective employers and/or sponsors.

References

- [1] Deghani, M., Azaronyad, H., Kamps, J., de Rijke, M.: Learning to transform, combine, and reason in open-domain question answering. In: WSDM’19: Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining. pp. 681–689 (2019), <https://doi.org/10.1145/3289600.3291012>