

SF-QA: Simple and Fair Evaluation Library for Open-domain Question Answering

Xiaopeng Lu^{2*}, Kyusong Lee¹ and Tiancheng Zhao¹
SOCO Inc.

¹{kyusongl, tianchez}@soco.ai

Language Technologies Institute, Carnegie Mellon University

²xiaopen2@andrew.cmu.edu

Abstract

Although open-domain question answering (QA) draws great attention in recent years, it requires large amounts of resources for building the full system and it is often difficult to reproduce previous results due to complex configurations. In this paper, we introduce SF-QA: simple and fair evaluation framework for open-domain QA. SF-QA framework modularizes the pipeline open-domain QA system, which makes the task itself easily accessible and reproducible to research groups without enough computing resources. The proposed evaluation framework is publicly available and anyone can contribute to the code and evaluations.

1 Introduction

Open-domain Question Answering (QA) is the task of answering open-ended questions by utilizing knowledge from a large body of unstructured texts, such as Wikipedia, world-wide-web and etc. This task is challenging because researchers have to face issues in both scalability and accuracy. In the last few years, rapid progress has been made and the performance of open-domain QA systems has been improved significantly (Chen et al., 2017; Qi et al., 2019; Yang et al., 2019). Several different approaches were proposed, including two-stage ranker-reader systems (Chen et al., 2017), end-to-end models (Seo et al., 2019) and retrieval-free models (Raffel et al., 2019). Despite people’s increasing interest in open-domain QA research, there are still two main limitations in current open-domain QA research communities that makes research in this area not easily accessible:

The first issue is the high cost of ranking large knowledge sources. Most of the prior research used Wikipedia dumps as the knowledge source. For example, the English Wikipedia has more than

7 million articles and 100 million sentences. For many researchers, indexing data of this size with a classic search engine (e.g., Apache Lucene (McCandless et al., 2010)) is feasible but becomes impractical when indexing with a neural ranker that requires weeks to index with GPU acceleration and consumes very large memory space for vector search. Therefore, research that innovates in ranking mostly originates from the industry.

The second issue is about reproducibility. Open-domain QA datasets are collected at different time, making it depends on different versions of Wikipedia as the correct knowledge source. For example, SQuAD (Rajpurkar et al., 2016) uses the 2016 Wikipedia dump, and Natural Question (Kwiatkowski et al., 2019) uses 2018 Wikipedia dump. Our experiments found that a system’s performance can vary greatly when using the wrong version of Wikipedia. Moreover, indexing the entire Wikipedia with neural methods is expensive, so it is hard for researchers to utilize others’ new rankers in their future research. Lastly, the performance of an open-domain QA system depends on many hyperparameters, e.g. the number of passages passed to the reader, fusion strategy, etc., which is another confounding factor to reproduce a system’s results.

Thus, this work proposes SF-QA (Simple and Fair Question-Answering), a Python library to solve the above challenges for two-stage QA systems. The key idea of SF-QA is to provide pre-indexed large knowledge sources as public APIs or cached ranking results; a hub of reader models; and a configuration file that can be used to precisely reproduce an open-domain QA system for a task. The pre-indexed knowledge sources enable researchers to build on top of the previously proposed rankers without worrying about the tedious work needed to index the entire Wikipedia. Then the executable configuration file provides a complete

* This work was done during an internship at SOCO

snapshot that captures all of the hyperparameters in order to reproduce a result.

Experiments are conducted to validate the effectiveness of SF-QA. We show that one can easily reproduce previous state-of-the-art open-domain QA results on four QA datasets, namely Open SQuAD, Open Natural Questions, Open CMRC, and Open DRCD. More datasets will be included in the future. Also, we illustrate several use cases of SF-QA, such as efficient reader comparison, reproducible research, open-source community, and knowledge-empowered applications.

SF-QA is also completely open-sourced¹ and encourages the research community to contribute their rankers or readers into the repository, so that their methods can be used by the rest of the community.

In short, the contributions of this paper include:

1. The proposed open-source SF-QA project that provides a complete pipeline for simplifying open-domain QA research.
2. A hub of pre-indexed Wikipedia at different years with different ranking algorithms as public APIs or cached results.
3. Experiments and tutorials that explain use cases and scenarios of SF-QA and validate its effectiveness.

2 Related Work

Existing deep learning open-domain QA approaches can be broadly divided into three categories.

2.1 Two-stage Approach

Recent open-domain QA systems mostly use a two-stage ranker-reader approach. Dr.QA (Chen et al., 2017) uses a modified TF-IDF bag-of-words method as the first-stage retriever. Selected documents are then fed into an RNN-based document reader to extract the final answer span. Wang et al. (2018) leverage reinforcement learning to update both ranker and reader components and shows improvement over Dr.QA in open-domain QA task. Lee et al. (2018) focuses on the ranker improvement and uses a learned reranker to boost first stage answer recall.

Some other works focus on second-stage reader improvement. Yang et al. (2019) adopts a BERT-based reader model (Devlin et al., 2018) instead

¹ <https://github.com/soco-ai/SF-QA.git>

of the previous RNN-based model and that significantly improved the end-to-end performance. To deal with span extraction in a multi-document setting, Wang et al. (2019) uses the global normalization approach (Clark and Gardner, 2017) to make the span scores comparable among candidate documents, which improved the performance by a large amount.

The graph-based ranker-reader approach has also been explored recently. Asai et al. (2019) proposes a graph-based retriever to retrieve supporting documents recursively based on entity link evidence, and then uses a BERT-based reader model to complete open-domain QA task.

2.2 End-to-End Approach

Open-domain QA using the end-to-end approach was not feasible for a long time, because this needs humongous memory to index the corpus and do the vector search. With the emergence of a large pre-trained language model (PLM), researchers revisit this idea and make the end-to-end open-domain QA feasible. Lee et al. (2019) proposed Open retrieval QA (ORQA) model, which updates the ranker and reader model in an end-to-end fashion by pre-training the model with an Inverse Cloze Task (ICT). Seo et al. (2019) experiments with considering open-domain QA task as a one-stage problem, and indexing corpus at phrase level directly. This approach shows promising inference speed with compromise in worse performance.

2.3 Retrieval-free Approach

Pre-trained language models have got rapid development in recent years. Querying a language model directly to get phrase-level answers becomes a possibility. The T5 model (11B version) (Rafael et al., 2019) can reach comparative scores on several open-domain QA datasets, compared with two-stage approaches with far less number of parameters ($\sim 330M$). However, as reported in Guu et al. (2020), decreasing the number of parameters hurts the model performance drastically. This leaves large room for future research on how to make retrieval-free open-domain QA feasible in the real-world setting.

3 The Proposed Method

3.1 Background

A typical ranker-reader-based open-domain QA system operates as follows: first, a large text

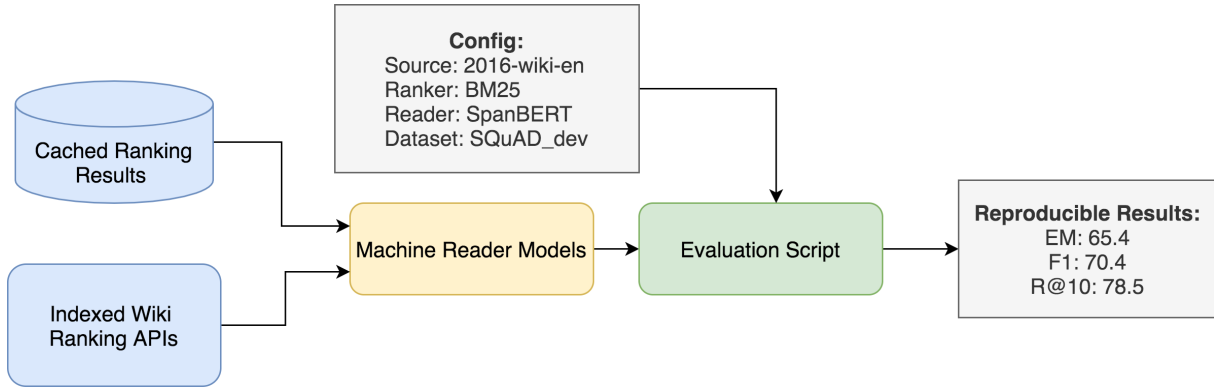


Figure 1: Overall pipeline for open-domain QA

knowledge-base is indexed by a ranker, e.g. a full-text search engine. Given a query, the ranker can return a list of relevant passages that may contain the correct answer. How to choose the size of a passage is still an open research question and many choices are available, e.g. paragraph, fixed-size chunks, and sentences. Note that it is not necessary that the ranker needs to return the final passages in one-shot: advanced ranker can iteratively refine the passage list to support multi-hop reasoning (Yang et al., 2018; Asai et al., 2019).

Then given the returned passages, a machine reader model will process all passages jointly and extract potential phrase-level answers from them. A fusion strategy is needed to combine candidate answers and scores from each passage and to create a final list of N-best phrase-level answers by reading these passages. The reason to combine ranker with the reader is to solve the scalability challenge since the state-of-the-art readers are prohibitively slow to process very large corpus in real-time (Chen et al., 2017; Devlin et al., 2018).

3.2 The Proposed Library Overview

SF-QA is a library that is designed to make it easy to evaluate and reproduce open-domain systems that use ranker-reader architecture. SF-QA decreases the cost of indexing, hosting, and querying large unstructured text knowledge base, e.g. Wikipedia, and also provides a complete configuration snapshot that can be used to replicate a QA system’s performance. It is also a place for open-domain QA researchers to share their work, no matter it is innovating in better information retrieval or it is in stronger machine reading comprehension.

There are four main components in SF-QA: *ranker service*, *reader hub*, *evaluation*, and *pipeline configuration*.

3.3 Ranker Service

The goal of the ranker service is to reduce the cost and time to index and query large knowledge source for open-domain QA research using a variety of ranking technologies. Up to date, we have included the BM25 (Robertson et al., 2009) and SPARTA (Zhao et al., 2020) ranking methods with several configurations detailed below. More methods will be included and we also welcome community contributions.

Currently, SF-QA supports four ways of document splitting for indexing:

1. Sentence: sentence-level indexing
2. Paragraph: paragraph-level indexing
3. Chunk: fixed word size indexing
4. Context: context-level indexing, where the full sentence is always kept, with a maximum number of tokens

Also, Wikipedia dumps at different times are indexed separately so that users can choose to use the same dump as benchmark datasets used. The following versions are included:

1. English Wikipedia: 2016/2018/2020
2. Chinese Wikipedia: 2017/2018/2020

The returned passage is in the following JSON format: $\{\{question_id\}: [{"score": 42.86, "answer": "Super Bowl V, the fifth edition of the Super Bowl...", ...}]$, which contains all question ids as key, and top-k retrieved documents and scores as value.

There are two methods to use the ranking results: *cached ranking results* and *ranking API*.

3.3.1 Cached Ranking Results

The fastest way to use ranking service for experiments is via cached ranking results. SF-QA provides top-K ranked passages in JSON format for training, validation and test (if publicly available) set. One can directly use the cached results for training or for testing, saving time, and resources for processing the raw data. Another use case is one may use more computationally expensive re-ranking methods to re-rank the top-K passages and then feed them into the reader component.

3.3.2 Ranking APIs

The cached results are very useful for researchers who work on existing datasets and who do not need to have a live system. However, only cached results do not work for new datasets or live QA system that needs to handle user queries. Therefore, SF-QA also provides public API as a service to solve this need. The API is available as a RESTful API and can be reached via HTTPs. Detailed connection documentation can be found on the GitHub.

3.4 Reader Hub

Reader hub allows SF-QA’s user to specify which reader model to use to extract phrase-level answers. One can either uses their own models by implementing an abstract function or directly load any reader models that are compatible with the Hugging Face Transformer library (Wolf et al., 2019). SF-QA also includes its own reader model that is optimized for open-domain QA. For example, it offers a BERT reader that is globally normalized (Wang et al., 2019), which provides more reliable answer scores to compare multiple candidates’ answers from different passages.

Moreover, the reader hub allows the user to define the fusion mechanism that combines the ranking results with reading results. The current implementation supports a linear combination with two free variables, namely the type of score and the weight on reader score. Concretely, the final answer score is computed as follows:

$$y = (1 - \alpha)y_{reader} + \alpha y_{ranker} \quad (1)$$

where α is a coefficient between 0 and 1. y_{reader} is the reader score, which can either be logits or probability after the softmax layer. y_{ranker} is the ranker score, which depends on the ranking method. One may also specify different normalization strategies to normalize the score from ranker or reader. Normalization strategies include z-normalization, floor

normalization etc. Lastly, one may easily add their own strategy by overriding the fusion function.

3.5 Evaluation

SF-QA evaluation is designed to offer a multilingual and comprehensive evaluation script that computes the performance of an open-domain QA system and also outputs useful intermediate metrics that are useful for analysis and visualization. For language support, SF-QA evaluation supports English and Chinese. For metrics, it has the most common EM (exact match) and F1 score for the final performance. It also provides other relevant metrics. The following is a list of metrics that are in the output:

- Exact match (EM)
- F-1 Score
- Ranking recall at K
- Oracle ranker score
- Mean reciprocal rank (MRR)

3.6 Pipeline Configurations

The pipeline configuration file is in YAML format, which defines all the hyperparameter for an open-domain QA system to do a forward inference. One can set the configuration for data, ranker ID, and reader ID, fusion strategy and etc. Therefore, the easiest way to share an open-domain QA system for results replication is via providing the right YAML configuration. The following is an example.

```
# config.yaml
data:
  lang: en
  name: squad
  split: dev-v1.1
ranker:
  use_cached: False
model:
  name: sparta
  es_index_name: en-wiki-2016
reader:
  model_id: squad-context-spanbert
param:
  n_gpu: 2
  score_weight: 0.8
  top_k: 10
```


4 Use Cases

SF-QA is designed to be modular and ready to use, with the hope that it can connect people from researchers interested in Question Answering (QA), Information Retrieval (IR), and developers from industries. In this section, we illustrate several use cases of SF-QA.

4.1 Efficient Reader Comparison

In open-domain QA, the first stage ranker consumes humongous resources in both time, memory, and storage. For researchers without enough computing power, it is not feasible to start open-domain QA research, even if they only want to improve the system on the reader stage. SF-QA provides solutions for researchers with this need. In SF-QA, existing publicly available open-domain QA datasets are already indexed with multiple rankers used in previous open-domain QA research work, currently including BM25 (Robertson et al., 2009), and SPARTA (Zhao et al., 2020), both with different granularity options. Researchers can call the RESTful API to get the cached ranking results directly if they want to focus on existing open-domain QA datasets, for example, Open SQuAD, Open CMRC, etc. Alternatively, they can call the backend live ranker to get the top retrieved results regarding the input query. We design SF-QA to be completely modular: the researcher is able to pick up a cached ranker and plug in their own reader model to evaluate the open-domain QA results.

4.2 Reproducible Research

Reproducibility is another problem that existed in current open-domain QA research. Since the first-stage retriever model needs researchers to collect large-scale data by themselves, it is hard to keep all the settings the same to make fair comparisons. In SF-QA, we collected data following the earliest works’ setting (Chen et al., 2017; Yang et al., 2019; Kwiatkowski et al., 2019). Therefore, researchers can check SF-QA to get data specifications for existing models. Moreover, parameter settings for different models are recorded and saved in another separate configuration file, as shown in the section 3.6. Therefore, any existing models in the current SF-QA project can be directly reproduced, which would greatly facilitate researchers in establishing benchmark scores and doing fair comparisons.

4.3 Knowledge-empowered Applications

SF-QA framework also considers the needs from an industry perspective. To show the potential of open-domain QA and to encourage more people to join the development of this task, we also provide a RESTful API (with a ready-to-use open-domain QA model in the backend) for users to ask questions and get the phrase-level answers directly as output. We also provide a tutorial to demonstrate that SF-QA can be seamlessly incorporated into RASA (Bocklisch et al., 2017), a popular open-source chatbot building platform, with only a few lines of code. We hope that this effort can attract people from different backgrounds to open-domain QA research.

5 Experiment Results

	Reported		Reproduced	
	EM	F1	EM	F1
Bertserini (Yang et al., 2019)	38.6	46.1	41.2	48.6
+DS (Xie et al., 2020)	51.2	59.4	51.6	59.2
Multipassage (Wang et al., 2019)	53.0	60.9	53.2	60.7
SpartaQA (Zhao et al., 2020)	59.3	66.5	59.3	66.5

Table 1: Comparison between reported performance and reproduced performance on Open SQuAD.

5.1 Reproducing Prior Art

Results in Table 1 shows the performance comparison between several reported open-domain QA systems and our reproduced results. The first experiment conducted is to reproduce some prior results using SF-QA. We choose Bertserini (Yang et al., 2019), Bertserini with distant supervision (Xie et al., 2020), Multi-passage Bert (Wang et al., 2019), and SPARTA (Zhao et al., 2020) as three benchmark systems to reproduce.

To reproduce Bertserini (Yang et al., 2019), we follow the implementation described in the original paper and first index the 2016 English Wikipedia in paragraph level to get 29.5M documents in total. A BERT-base-cased model is trained with global normalization, following descriptions in the paper. We observe a slight improvement in the open-domain QA result, which may due to the usage of a newer version of the BM25 retriever. The same

	Indexing	Uploading	Retrieval	Reader	Total
Traditional Approach	16.2 h	5.2 h	6.1 h	4.4 h	21.9 h
SF-QA	-	-	-	4.4 h	4.4h

Table 2: Time elapsed to evaluate open-domain QA using Open SQuAD development set

open-domain QA setting	wiki 2016*			wiki 2018			wiki 2020		
	EM	F1	R@1	EM	F1	R@1	EM	F1	R@1
BM25 + SpanBERT	49.2	56.7	41.9	45.8	53.8	39.4	41.5	49.5	35.4
Sparta + SpanBERT	59.3	66.5	50.8	46.5	54.4	39.3	46.4	53.9	42.2

Table 3: Open SQuAD performance using Wikipedia dumps from different years. * represents the dump which SQuAD originally used for annotation.

phenomenon has also been reported in (Xie et al., 2020).

For Bertserini with distant supervision (Xie et al., 2020), we follow the two-stage distant supervision strategy proposed by the original author, where the model was first fine-tuned using the original SQuAD dataset, and then fine-tuned on the distantly supervised data retrieved from the full Wikipedia. The score we get matches the score reported by the original author.

To reproduce Multi-passage BERT (Wang et al., 2019), we first index the Wikipedia corpus using chunk size equals to 100, with a stride of 50 words. A BERT reranker is then trained to rerank the retrieved top 100 documents and the top 30 documents are passed to the reader. In the reader training stage, we train the model using BERT-large-cased model, also with global normalization to make the span score comparable. Our reproduced score matches the score reported in the original paper.

For SpartaQA, we follow the original author’s implementation on SPARTA retriever, and index the Wikipedia in the context level with a size of 150. During the reader stage, a SpanBERT (Joshi et al., 2020) model is used to train the model with distantly supervised data retrieved from Wikipedia with global normalization strategy. The score matches the reported score.

5.2 Time saved by SF-QA

This experiment shows results for elapsed time to evaluate open domain question answering with and without the SFQA evaluation framework (Table 2). Traditionally, we need to build the complete pipeline in order to evaluate the open-domain QA as following steps: (1) **Indexing**: converting full Wikipedia into sparse or dense representations;

(2) **Uploading**: inserting the text and representations to Elasticsearch (or similar database); 3) **Retriever**: retrieval n-best candidates from Elasticsearch; 4) **Reader**: span prediction using machine reading comprehension. We use GeForce RTX 2080 Ti GPU to index the entire Wikipedia dump of the total 89,544,689 sentences. The total amount of elapsed time for open-domain QA is 29 hours without using SF-QA for one experimental setting. In comparison to this, using cached retrieved results provided from SF-QA saves repetitive work in heavy indexing, and it only takes ~ 4 hours to get the final scores.

5.3 Model Accuracy v.s. Corpus release year

We conduct the last experiment to test the robustness of the state-of-the-art system against temporal shifting. Results are reported in Table 3. We observe that model accuracy is largely affected by the version of the Wikipedia dump, showing that it is essential to track the version of the input data and make sure that all open-domain QA researches are reproducible starting from the data input level.

6 Conclusion

In conclusion, this paper presents SF-QA, a novel evaluation framework to make open-domain QA research simple and fair. This framework fixes the gap among researchers from different fields, and make the open-domain QA more accessible. We show the robustness of this framework by successfully reproducing several existing models in open-domain QA research. We hope that SF-QA can make the open-domain QA research more accessible and make the evaluation easier. We expect to further improve our framework by including more models in both ranker and reader side, and encourage community contributions to the project as well.

References

- Akari Asai, Kazuma Hashimoto, Hannaneh Hajishirzi, Richard Socher, and Caiming Xiong. 2019. Learning to retrieve reasoning paths over wikipedia graph for question answering. *arXiv preprint arXiv:1911.10470*.
- Tom Bocklisch, Joey Faulkner, Nick Pawlowski, and Alan Nichol. 2017. Rasa: Open source language understanding and dialogue management. *arXiv preprint arXiv:1712.05181*.
- Danqi Chen, Adam Fisch, Jason Weston, and Antoine Bordes. 2017. Reading Wikipedia to answer open-domain questions. In *Association for Computational Linguistics (ACL)*.
- Christopher Clark and Matt Gardner. 2017. Simple and effective multi-paragraph reading comprehension. *arXiv preprint arXiv:1710.10723*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Ming-Wei Chang. 2020. Realm: Retrieval-augmented language model pre-training. *arXiv preprint arXiv:2002.08909*.
- Mandar Joshi, Danqi Chen, Yinhan Liu, Daniel S Weld, Luke Zettlemoyer, and Omer Levy. 2020. Spanbert: Improving pre-training by representing and predicting spans. *Transactions of the Association for Computational Linguistics*, 8:64–77.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, et al. 2019. Natural questions: a benchmark for question answering research. *Transactions of the Association for Computational Linguistics*, 7:453–466.
- Jinhyuk Lee, Seongjun Yun, Hyunjae Kim, Miyoung Ko, and Jaewoo Kang. 2018. Ranking paragraphs for improving answer recall in open-domain question answering. *arXiv preprint arXiv:1810.00494*.
- Kenton Lee, Ming-Wei Chang, and Kristina Toutanova. 2019. Latent retrieval for weakly supervised open domain question answering. *arXiv preprint arXiv:1906.00300*.
- Michael McCandless, Erik Hatcher, Otis Gospodnetić, and O Gospodnetić. 2010. *Lucene in action*, volume 2. Manning Greenwich.
- Peng Qi, Xiaowen Lin, Leo Mehr, Zijian Wang, and Christopher D Manning. 2019. Answering complex open-domain questions through iterative query generation. *arXiv preprint arXiv:1910.07000*.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2019. Exploring the limits of transfer learning with a unified text-to-text transformer. *arXiv preprint arXiv:1910.10683*.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. Squad: 100,000+ questions for machine comprehension of text. *arXiv preprint arXiv:1606.05250*.
- Stephen Robertson, Hugo Zaragoza, et al. 2009. The probabilistic relevance framework: Bm25 and beyond. *Foundations and Trends® in Information Retrieval*, 3(4):333–389.
- Minjoon Seo, Jinhyuk Lee, Tom Kwiatkowski, Ankur Parikh, Ali Farhadi, and Hannaneh Hajishirzi. 2019. Real-time open-domain question answering with dense-sparse phrase index. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4430–4441.
- Shuohang Wang, Mo Yu, Xiaoxiao Guo, Zhiguo Wang, Tim Klinger, Wei Zhang, Shiyu Chang, Gerry Tesauro, Bowen Zhou, and Jing Jiang. 2018. R 3: Reinforced ranker-reader for open-domain question answering. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Zhiguo Wang, Patrick Ng, Xiaofei Ma, Ramesh Nallapati, and Bing Xiang. 2019. Multi-passage bert: A globally normalized bert model for open-domain question answering. *arXiv preprint arXiv:1908.08167*.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. 2019. Huggingface’s transformers: State-of-the-art natural language processing. *ArXiv*, pages arXiv–1910.
- Yuqing Xie, Wei Yang, Luchen Tan, Kun Xiong, Nicholas Jing Yuan, Baoxing Huai, Ming Li, and Jimmy Lin. 2020. Distant supervision for multi-stage fine-tuning in retrieval-based question answering. In *Proceedings of The Web Conference 2020*, pages 2934–2940.
- Wei Yang, Yuqing Xie, Aileen Lin, Xingyu Li, Luchen Tan, Kun Xiong, Ming Li, and Jimmy Lin. 2019. End-to-end open-domain question answering with bertserini. *arXiv preprint arXiv:1902.01718*.
- 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. *arXiv preprint arXiv:1809.09600*.
- Tiancheng Zhao, Xiaopeng Lu, and Kyusong Lee. 2020. Sparta: Efficient open-domain question answering via sparse transformer matching retrieval. *arXiv preprint arXiv:2009.13013*.