

Keyword Extraction for Improved Document Retrieval in Conversational Search

Oleg Borisov¹, Mohammad Aliannejadi² and Fabio Crestani¹

¹Università della Svizzera italiana (USI), Lugano, Switzerland

²University of Amsterdam, Amsterdam, Netherlands

Abstract

Recent research has shown that mixed-initiative conversational search, based on the interaction between users and computers to clarify and improve a query, provides enormous advantages. Nonetheless, incorporating additional information provided by the user from the conversation poses some challenges. In fact, further interactions could confuse the system as a user might use words irrelevant to the information need but crucial for correct sentence construction in the context of multi-turn conversations. To this aim, in this paper, we have collected two conversational keyword extraction datasets and propose an end-to-end document retrieval pipeline incorporating them. Furthermore, we study the performance of two neural keyword extraction models, namely, BERT and sequence to sequence, in terms of extraction accuracy and human annotation. Finally, we study the effect of keyword extraction on the end-to-end neural IR performance and show that our approach beats state-of-the-art IR models. We make the two datasets publicly available to foster research in this area.

Keywords

Conversational Search, Mixed-Initiative Conversations, Keyword Extraction

1. Introduction

Recent developments in speech recognition and deep learning have led to intelligent assistants, such as Google Assistant, Microsoft Cortana, and Apple Siri. Consequently, researchers and users are exploring novel means of communication and information access, such as spoken queries and conversations [1]. Research on information-seeking conversational systems has gained lots of attention recently. Various shared evaluation tasks have been raised in the community, focusing on single- [2] and mixed-initiative [3, 4] conversational search systems. The aim of research in mixed-initiative conversations is to enable a system to take the initiative of the conversation when necessary, aiming to provide a better experience to the user [5]. An example of mixed-initiative interaction is asking clarifying questions that has been recently studied in the context of information-seeking conversations [6, 7] and Web search [8, 9, 10, 11].

In Web search, where users usually type their queries, they take some time to formulate a query and often do not follow common sentence structures. For example, they only focus on using the most important words for their search. Consequently, a narrow focus is created for the search engine, making the inspection of documents for the most relevant query words easier.


IIR'21: Italian Information Retrieval Workshop (IIR)

✉ oleg.borisov@usi.ch (O. Borisov); m.aliannejadi@uva.nl (M. Aliannejadi); fabio.crestani@usi.ch (F. Crestani)

🌐 <http://aliannejadi.com> (M. Aliannejadi)



© 2021 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

 CEUR Workshop Proceedings (CEUR-WS.org)

In contrast to this, conversational IR faces challenges due to the inclination of users to follow their own speech patterns when formulating queries rhetorically. Here, users tend to include some unnecessary terms that appear crucial for a proper sentence construction but might derail the IR model in searching for relevant documents [12]. This could also be magnified when a conversation evolves into multiple turns [13, 14] and a new form of conversation is presented to the user, such as when the system asks clarifying questions. This happens mainly due to the context-dependence nature of multi-turn conversations and new types of responses that could emerge in a mixed-initiative conversation.

While the effectiveness of conversational systems has been studied before [6, 15, 7, 16], the main goal of this paper is to study if the identification of keywords retrieved from the human-computer interaction will help achieve better retrieval results. To this aim, we collect two datasets of keyword extraction and study the effectiveness of multiple generative models on them. Our first dataset is collected based on the performance of the retrieval model using different keywords, while the other is collected from news articles online. Every news article comes with a title and a set of keywords. Our intuition is that a neural model can learn to extract useful keywords from news titles and use this external knowledge for more effective keyword extraction in a conversation. We study the effect of various keyword extraction strategies on non-neural and neural document retrieval pipelines. To the best of our knowledge, keyword extraction in the context of mixed-initiative conversational IR has not been studied before.

In our retrieval pipeline, after the conversational phase, where the system interacts with the user to clarify the query ambiguities, the conversational sequence is passed to the keyword extractor, identifying the most important terms from the sentences. In parallel to that, the document retrieval model performs the first relevance ranking of the documents based on the **original** conversation. Finally, the Neural IR performs re-ranking using top documents from the *IR phase 1* and **keywords** obtained by keyword extractor from *Keyword Extraction Phase* as inputs of the system.

2. Data Collection

As the topics addressed by this study have only recently surfaced in research, a substantial amount of work needed to be done to answer whether keywords could support the IR model with document retrieval tasks.

As was discussed earlier, keyword extraction from short-sized documents using Deep Learning is a relatively new topic. The previously created Inspec, SemEval-2010, SemEval-2017 datasets are not suitable for this research, as they are focused on keyword and keyphrase extraction from medium- and large-sized texts (e.g., abstracts or scientific articles) [17, 18, 19]. In contrast to this, the main focus of this research is keyword extraction from small-sized sentences of the length of no more than 20 words, which is the average English sentence length [20]. Therefore, we collect and release two types of datasets: (i) News-Keyword based and (ii) IR-Keyword based datasets.¹

¹Data available at: <https://github.com/aliannejadi/ConvKey>

2.1. News-Keywords Based dataset

Online newspaper websites and other social network Web pages tend to follow a content structure, where common articles are structured as title, main text, tags (sometimes hashtags). Content creators try not only to select an appealing and interesting title name but also to summarize the content in one sentence, thus selecting the most important words to portray the key message of an article. This can also be considered a reverse IR operation as the author, given the document's content, provides a title (considered the *query* in our case) that corresponds to the article in the best possible way.

Authors usually also choose some tags that either describe the article in the most general way or place the story in the context of other related articles that one could find on the website. From the user's point of view, tags provide an opportunity to navigate to other related material; however, having well-formulated tags is also crucial for Search Engine Optimization and could impact the website's visibility or the article [21].

Taking into consideration that writers pay very close attention to the title and the tags used, where it is not unusual for tag words to appear in the title, brings us to the first method of keyword dataset: considering *title* as the **input text**, and *tags* as the **target keywords**. If a tag does not appear in a corresponding title, we do not add it to the keywords list (as shown in Table 1).

Table 1

Example of scraped news article²

Title	Tags	Keywords List
Five German words you'll need to know this <i>summer</i>	[<i>summer</i> , holidays, members]	[<i>summer</i>]

To create the dataset, we scraped the following news websites: BBC³, The Local⁴ and Salon⁵. In total, over 104,000 title-tag pairs have been obtained using this method. After filtering the outliers and the items where the tags do not appear in the title, the dataset shrinks to 79,000 instances.

2.2. IR-Keyword-based dataset

Classical IR systems only focus on the basic preprocessing of the query, such as the removal of stopwords and punctuation. Having too many words could confuse the system and lead it to retrieve unwanted results. Therefore, a correct keyword identification could lead to better retrieval performance, while selecting less good keywords will inevitably worsen the output results. We developed the IR-Keyword-based dataset based on this assumption, applying the previously created Qulac dataset [6]. To create a dataset in this context, we used Qulac's first conversational round, containing three components: query *q*, question *t*, and answer *a*, which retrieve a set of relevant documents.

²Taken on 30 June 2020 from

<https://www.thelocal.ch/20200630/five-german-words-youll-need-to-know-this-summer>

³<https://www.bbc.com/>

⁴<https://www.thelocal.ch/>

⁵<https://www.salon.com/>

The main idea is to identify a set of words from q , t , and a , which will lead to the greatest relevance of retrieved documents. To evaluate the system’s performance, we used the Normalized Discounted Cumulative Gain at 20 (NDCG@20) metrics. Due to the complexity of the permutations of all potential keywords of the whole set q, t, a , we decided to focus on one component at a time. The algorithm that was used is presented in Algorithm 1. The main idea is to choose s_0 , which could be a query, question, or an answer. For example, let us consider s_0 a *query* and s_1, s_2 to represent the *question* and *answer*. Afterward, we would like to consider all possible subsets of words of s_0 (query in our example), which will form a set of *potential keywords*. In mathematical terms, such operation is known as **the powerset**. For instance, if s_0 is "How are you?", then a set of *potential keywords* would be: {"how", "are", "you", "how are", "how you", "you are", "how are you"}⁶. The cardinality of a powerset highly depends on the number of words that the input sentence contains. To address having a large powerset, we limit the maximum size of the subset to four words.

Next, we consider one instance k_i from the *potential keyword* set and retrieve the documents by supplying k_i, s_1 and s_2 to the document retrieval model. Consequently, it is important to evaluate the retrieved documents’ relevance and save the obtained score. In the end, we save the s_0 as the **input text** and k_{best} as the set of **keywords** that led to the retrieval of the most relevant documents. We repeat a similar operation by considering s_0 as the question, and s_1, s_2 as query and question, and later s_0 as the answer, and s_1, s_2 as query and answer, respectively. We apply the same process for all conversations from the Qulac dataset until we receive keywords from all queries, questions, and answers. Applying this approach, 15,320 data samples were obtained. The benefit which this approach suggests is that where the answer of a user in a computer interaction is uncertain or ambiguous and will not provide any important information, the system learns to ignore these. In this scenario, the system should ideally ask another question or base the search only on the initial query. Therefore, the proposed method of dataset generation will be able to mimic this behavior.

3. Proposed Methods

This section describes our conversational IR framework. We start with the neural models that we used for the keyword extraction task. Then we continue with the neural IR models and describe how keyword extraction fits into our pipeline.

3.1. Keyword Extraction Models

For the *Keyword Extraction Phase*, we experimented with two different types of neural models: *Sequence-to-Sequence* architecture and *BERT* model [22, 23]. Sequence-to-Sequence architecture uses Gated Recurrent Unit (GRU) as a recurrent neural network, the Attention mechanism to help the decoder, and pre-trained Word2Vec embeddings the performance on the words outside of the training set vocabulary [24, 25, 26]. We use Sequence-to-Sequence because it has been a state-of-the-art architecture for many different NLP tasks and established new benchmarks for the tasks of Neural Machine Translation [22]. In contrast to the previously described model, we

⁶We also keep the original order in which the words appear in the text

Algorithm 1: IR-Keyword Based Dataset Creation Method.

```
Method: find_keywords(query, question, answer)
for  $s_0$  in [query, question, answer] do
     $s_1, s_2 = [\text{query}, \text{question}, \text{answer}].\text{remove}(s_0)$ ;
    potential_keywords = PowerSet( $s_0$ , maxSubsetSize=4) ;

    scores_list = list() ;
    for  $k_i$  in potential_keywords do
        ranked_documents = IRmodel.retrieve( $k_i, s_1, s_2$ ) ;
        score = ranked_documents.evaluate(metrics="NDCG@20");
        scores_list.append(score) ;
    end
    max_score_index = argmax(scores_list) ;
     $k_{best} = \text{potential\_keywords}[\text{max\_score\_index}]$  ;
    save("input text" =  $s_0$ , "keywords"= $k_{best}$ )
end
```

Table 2Example of Data Supplied to Neural Networks.⁷

Original Sentence	"Conservatives and liberals drink different beer"
Tokenized Sentence	['conservatives', 'and', 'liberals', drink', 'different', 'beer']
Keywords	['conservatives', 'liberals', 'beer']
Named Entities	[1, 0, 1, 0, 0, 1]

also selected BERT as the most recently developed Transformer-based neural architecture in the field of NLP. One of its biggest advantages is that it has been pre-trained on a great amount of data using two main approaches: Masked Language Model, which is related to the prediction of masked/hidden tokens in the input sentence, and Next Sentence Prediction, which has the objective of predicting the next sentence from the input sequence. Therefore, by fine-tuning the model, it is possible to achieve great results in tasks, such as: Named Entity Recognition (NER), Sentence Classification, Answer Searching, and others [23].

To train selected architectures, we formulate the task of keyword extraction in the form of a NER task, as shown in Table 2. Where we say that a word is a **keyword** if its corresponding entity is labeled as "1", and **not a keyword** if it is marked as "0".

3.2. Neural IR models

We extend the solution available from previous research [6] by adding *Information Retrieval Phase 2*, represented by the Neural IR model. We study the effectiveness of the following commonly used two Neural IR models:

1. **Deep Relevance Matching Model** (DRMM): this model puts more emphasis on the

⁷Retrieved on 15th of October 2019, from:https://www.salon.com/2013/02/27/conservatives_and_lilberals_drink_different_beer_partner/

relevance (both semantic and lexical) matching of the query rather than on exclusively semantic matching. It considers three crucial factors of the "handling of the exact matching signals, query term importance, and diverse matching requirements" [27].

2. **Deep Semantic Similarity Model (DSSM)**: based on the Siamese network architecture, DSSM has the main focus in comparing cosine similarities of the vector representations of a query and the document, where vector representations are learned using Deep fully-connected layers [28]. Originally such a model was only used for short text matching tasks (for example, matching questions with the most relevant answers); however, later, DSSM proved to be useful for tasks involving documents containing long texts, thus being a perfect choice for IR related tasks [29].

4. Experimental Results

4.1. Experimental Setup

Data. We use the publicly available Qulac⁸ dataset, which is built based on the TREC Web Track 2009-2012, for our experiments. For keyword extraction experiments, we use the two datasets described in Section 2.

Metrics. We evaluate keyword extractions' performance in two ways, namely, the accuracy of extraction and end-to-end document retrieval. As for extraction accuracy, we use the following evaluation metrics: Precision, Recall, Average Tag Correct Identification (ATCI)⁹, and Correct per Response Fill (CpRF)¹⁰. Also, we perform a human evaluation on the extracted keywords, where we ask the human annotators to score each extracted keyword from 1 to 5. Our IR evaluation follows the standard IR metrics, namely, Normalized Discounted Cumulative Gain at k (nDCG@ k), Precision at k (P@ k), Mean Reciprocal Rank (MRR), and Mean Average Precision (MAP).

Statistical Significance. We perform two-tailed paired t-test $p - value < 0.05$ to determine significant improvements on the IR metrics.

4.2. Keyword Extractors

To evaluate the performance of Keyword Extracting Neural Networks, two methods have been used. The first one relies on *test dataset accuracy*, while the second one is a *human evaluation* method.

Test dataset accuracy. Table 3 shows the performance of the Keyword Extractors. We also created a simple "Non-Neural approach" to serve as a baseline. This method operates in a very elementary way: the word frequencies were calculated from the training dataset. Using a brute force approach, the optimal frequency threshold was found, which maximizes the

⁸<http://www.github.com/aliannejadi/qulac>


⁹tests the quality of the overall assigned tags, by checking if the model has correctly assigned *keyword* or *not keyword* tag

¹⁰It captures the ratio of fully correct and partially correct predictions to the total amount of sentences in the dataset (adopted from MUC-5)

Table 3

Performance of keyword extractors on the test set

	Precision	Recall	ATCI	CpRF	Human Evaluation
Non-Neural	0.4572	0.7186	0.6216	0.4873	-
seq2seq	0.6744	0.7795	0.8311	0.5863	4.268
BERT	0.8636	0.8792	0.9178	0.7612	3.964

**Figure 1:** Performance of Keyword Extractors according to the human evaluation for various keywords positions in sentences.

correct identification of tagged keywords (if the frequency of a certain word is below the threshold, the word is assigned to be a keyword). If the word has not been seen before, it is automatically assigned as the keyword as it is considered rare enough as it has not appeared in the training corpus. In Table 3, we see that BERT seems to be an ultimate solution to the keyword extraction problem, as the model achieves a much better test set performance than other keyword extractors.

Human Evaluation. While the testing set evaluates the models’ performances in a similar training environment, it is also essential to test whether the extracted keywords would suit human judgments. To address this question, Google’s Query Wellformedness dataset has been used [30]. Judges were asked to select the least possible number of keywords given a sentence and rate the relevance of the keywords chosen by the keyword extractor on a scale from 1 to 5, where 5 is the best score. For the latter part, we asked judges to imagine themselves in a situation where they have to answer the question based on only the keywords provided. In the scenario that the Neural Network’s selected keywords were doubtful, we asked the judges to use Google’s Search Engine and plug the keywords to see if sufficiently good results were obtained.

As can be observed from Table 3, Sequence-to-Sequence (Seq2seq) models appear to retrieve better keywords than BERT as the scores given to the model by the judges are higher. Additionally, it is essential to focus on words that the judges selected, the Seq2seq and BERT models alike, to describe the reasoning of the classifiers and compare it to the experts’ judgment. We explore the locations in which the keywords tend to appear in the sentences more often. As clearly seen from Figure 1, keywords appear to be located closer to the end of the sentence. Neural Networks have correctly learned this trend. However, it can be noted that, in general,

Table 4

Performance of the IR models. Bold values denote the best performance for each metric. † denotes significant improvements compared to the *DRMM-orig* model using two-tailed paired t-test with $p < 0.05$.

		nDCG@3	nDCG@5	P@3	P@5	MRR	MAP
DRMM	orig	0.3039	0.3304	0.2502	0.2289	0.3925	0.3433
	s2s	0.2953	0.3117	0.2491	0.2200	0.3759	0.3272
	bert	0.2939	0.3126	0.2510	0.2215	0.3761	0.32586
	non-neural	0.2117	0.2328	0.1812	0.1718	0.3027	0.2529
DSSM	orig	0.2117	0.2328	0.1812	0.1718	0.3027	0.2529
	s2s	0.3194†	0.3415†	0.2710†	0.2383†	0.3914†	0.3420
	bert	0.3231†	0.3386†	0.2677†	0.2345†	0.4079†	0.3476
	non-neural	0.2117	0.2328	0.1812	0.1718	0.3027	0.2529

the models tend to underestimate the number of keywords in a sentence.

Both evaluation approaches (test set and human evaluation) give interesting insights as we can clearly see that BERT learned better keywords that lead to the best document retrieval. In contrast, according to the Human evaluation dataset, Sequence-to-Sequence was able to retrieve more relevant keywords. Therefore, we also study the impact of both models on end-to-end IR performance to see how they eventually affect IR performance.

4.3. Neural IR Models

The performance of the Neural IR Models is presented in Table 4. As can be observed, in the case of the DRMM Neural IR model usage, the models provided with keywords have achieved a similar performance and have outperformed the Non-Neural IR model. Interestingly, the DRMM supplied with original conversational sequences was able to show the best performance concerning other DRMM models.

Looking at the DSSM, we can observe that providing keywords using BERT or the Sequence-to-sequence architecture yields much better results than when using original conversational sequences or the Non-Neural model (the last two achieved relatively similar performances). DSSM models that have used keywords have achieved the overall best retrieval performance.

Keywords Extractor Influence on IR model. Another interesting insight is provided by considering how well the Neural IR model performs concerning the effectiveness of the Keyword Extractor. In this case, we are interested in the *precision* of keywords provided by the Sequence-to-sequence Keyword Extractor and how the produced keywords impact the performance of the DSSM model.

As we see in Table 5, the premise is that *the better the quality of the produced keywords, the better the IR model will perform*. It is also interesting to see how much the Neural IR model will benefit from high-quality keywords. First, we start with the test dataset ordered by the relevance scores assigned by the Neural IR models. The next step is to split the dataset into three sub-parts, based on the precision of the keywords obtained from the Keyword Extractor’s query-question-answer sequences.

Table 5

Performance of Neural IR models on dataset depending on the precision of Sequence-to-sequence Keyword Extractor

	NDCG@5	P@5	MAP
Low Precision	0.2445	0.1557	0.2476
Medium Precision	0.4399	0.3127	0.4324
High Precision	0.4503	0.3502	0.4595

$$P(k) = \begin{cases} \text{low} & 0 \leq P(k) < 0.5 \\ \text{medium} & 0.5 \leq P(k) < 0.75 \\ \text{high} & \text{otherwise} \end{cases} \quad (1)$$

Where $P(k)$ is the precision of the keyword extractor on the sentence k .

Table 5 shows how much the quality of the keywords influences the document retrieval of DSSM models. As it can be observed, the model has managed to rank the documents much better when the supplied keywords had a high relevance.

5. Conclusions

This research studied the application of keywords in the context of conversational IR is going to be advantageous. For this purpose, we created two keyword extraction datasets and studied two types of Keyword Extractor, one based on a seq2seq architecture and the other based on BERT. We tested the keyword extraction performance based on keyword extraction, as well as end-to-end document retrieval performance. To do so, we test the performance on two state-of-the-art neural IR models, namely, DRMM and DSSM. We showed that Neural IR models supplied with keywords from conversational communications with users improve the relevance of retrieved documents through experimental results. In addition, we showed that the higher the Keyword Extractor’s precision, the better is the performance of the DSSM IR model.

For further work, it would be interesting to train the Neural Networks on a newly created dataset manually labeled by humans. Likely, the keyword dataset creation approach, which we proposed in this paper, misses some important keywords that humans will identify easily.

Acknowledgement

The work constitutes part of the master thesis of Oleg Borisov at the Università della Svizzera italiana (USI) on conversational search.

References

- [1] F. Radlinski, N. Craswell, A theoretical framework for conversational search, in: CHIIR, ACM, 2017, pp. 117–126.
- [2] J. Dalton, C. Xiong, V. Kumar, J. Callan, Cast-19: A dataset for conversational information seeking, in: SIGIR, ACM, 2020, pp. 1985–1988.

- [3] M. Aliannejadi, J. Kiseleva, A. Chuklin, J. Dalton, M. S. Burtsev, Convai3: Generating clarifying questions for open-domain dialogue systems (clariq), CoRR abs/2009.11352 (2020).
- [4] M. Aliannejadi, L. Azzopardi, H. Zamani, E. Kanoulas, P. Thomas, N. Craswell, Analysing mixed initiatives and search strategies during conversational search, in: CIKM, ACM, 2021.
- [5] E. Horvitz, Principles of mixed-initiative user interfaces, in: M. G. Williams, M. W. Altom (Eds.), CHI, ACM, 1999, pp. 159–166.
- [6] M. Aliannejadi, H. Zamani, F. Crestani, W. B. Croft, Asking clarifying questions in open-domain information-seeking conversations, in: SIGIR, ACM, 2019, pp. 475–484.
- [7] H. Hashemi, H. Zamani, W. B. Croft, Guided transformer: Leveraging multiple external sources for representation learning in conversational search, in: SIGIR, ACM, 2020, pp. 1131–1140.
- [8] H. Zamani, S. T. Dumais, N. Craswell, P. N. Bennett, G. Lueck, Generating clarifying questions for information retrieval, in: WWW, ACM / IW3C2, 2020, pp. 418–428.
- [9] H. Zamani, B. Mitra, E. Chen, G. Lueck, F. Diaz, P. N. Bennett, N. Craswell, S. T. Dumais, Analyzing and learning from user interactions for search clarification, in: SIGIR, ACM, 2020, pp. 1181–1190.
- [10] I. Sekulic, M. Aliannejadi, F. Crestani, User engagement prediction for clarification in search, in: ECIR, Lecture Notes in Computer Science, Springer, 2021, pp. 619–633.
- [11] T. Lotze, S. Klut, M. Aliannejadi, E. Kanoulas, Ranking clarifying questions based on predicted user engagement, CoRR abs/2103.06192 (2021).
- [12] M. P. Kato, T. Yamamoto, H. Ohshima, K. Tanaka, Cognitive search intents hidden behind queries: a user study on query formulations, in: Proceedings of the 23rd International Conference on World Wide Web, 2014, pp. 313–314.
- [13] M. Aliannejadi, M. Chakraborty, E. A. Rissola, F. Crestani, Harnessing evolution of multi-turn conversations for effective answer retrieval, in: CHIIR, ACM, 2020, pp. 33–42.
- [14] N. Voskarides, D. Li, P. Ren, E. Kanoulas, M. de Rijke, Query resolution for conversational search with limited supervision, in: SIGIR, ACM, 2020, pp. 921–930.
- [15] Y. Zhang, X. Chen, Q. Ai, L. Yang, W. B. Croft, Towards conversational search and recommendation: System ask, user respond, in: Proceedings of the 27th ACM International Conference on Information and Knowledge Management, 2018, pp. 177–186.
- [16] A. M. Krasakis, M. Aliannejadi, N. Voskarides, E. Kanoulas, Analysing the effect of clarifying questions on document ranking in conversational search, in: ICTIR, ACM, 2020, pp. 129–132.
- [17] I. Augenstein, M. Das, S. Riedel, L. Vikraman, A. McCallum, Semeval 2017 task 10: Scienceie-extracting keyphrases and relations from scientific publications, arXiv preprint arXiv:1704.02853 (2017).
- [18] A. Hulth, Improved automatic keyword extraction given more linguistic knowledge, in: Proceedings of the 2003 conference on Empirical methods in natural language processing, 2003, pp. 216–223.
- [19] S. N. Kim, O. Medelyan, M.-Y. Kan, T. Baldwin, Automatic keyphrase extraction from scientific articles, Language resources and evaluation 47 (2013) 723–742.
- [20] M. Cutts, Oxford guide to plain English, Oxford University Press, USA, 2020.

- [21] N. Yalçın, U. Köse, What is search engine optimization: Seo?, *Procedia-Social and Behavioral Sciences* 9 (2010) 487–493.
- [22] I. Sutskever, O. Vinyals, Q. V. Le, Sequence to sequence learning with neural networks, in: *Advances in neural information processing systems*, 2014, pp. 3104–3112.
- [23] J. Devlin, M.-W. Chang, K. Lee, K. Toutanova, Bert: Pre-training of deep bidirectional transformers for language understanding, *arXiv preprint arXiv:1810.04805* (2018).
- [24] D. Bahdanau, K. Cho, Y. Bengio, Neural machine translation by jointly learning to align and translate, *arXiv preprint arXiv:1409.0473* (2014).
- [25] K. Cho, B. Van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, Y. Bengio, Learning phrase representations using rnn encoder-decoder for statistical machine translation, *arXiv preprint arXiv:1406.1078* (2014).
- [26] T. Mikolov, K. Chen, G. Corrado, J. Dean, Efficient estimation of word representations in vector space, *arXiv preprint arXiv:1301.3781* (2013).
- [27] J. Guo, Y. Fan, Q. Ai, W. B. Croft, A deep relevance matching model for ad-hoc retrieval, in: *Proceedings of the 25th ACM International on Conference on Information and Knowledge Management*, 2016, pp. 55–64.
- [28] P.-S. Huang, X. He, J. Gao, L. Deng, A. Acero, L. Heck, Learning deep structured semantic models for web search using clickthrough data, in: *Proceedings of the 22nd ACM international conference on Information & Knowledge Management*, 2013, pp. 2333–2338.
- [29] B. Mitra, F. Diaz, N. Craswell, Learning to match using local and distributed representations of text for web search, in: *Proceedings of the 26th International Conference on World Wide Web*, 2017, pp. 1291–1299.
- [30] M. Faruqui, D. Das, Identifying well-formed natural language questions, *arXiv preprint arXiv:1808.09419* (2018).