



Gospodinov, M., [MacAvaney, S.](#) and [Macdonald, C.](#) (2023) Doc2Query--: When Less is More. In: 45th European Conference on Information Retrieval (ECIR'23), Dublin, Ireland, 02-06 Apr 2023, pp. 414-422. ISBN 9783031282379 (doi: [10.1007/978-3-031-28238-6_31](https://doi.org/10.1007/978-3-031-28238-6_31))

This is the author version of the work. You are advised to consult the publisher version if you wish to cite from it:

https://doi.org/10.1007/978-3-031-28238-6_31

<https://eprints.gla.ac.uk/287683/>

Deposited on: 11 January 2023

Enlighten – Research publications by members of the University of Glasgow
<http://eprints.gla.ac.uk>

Doc2Query--: When Less is More

Mitko Gospodinov¹, Sean MacAvaney², and Craig Macdonald²

University of Glasgow

¹2024810G@student.gla.ac.uk

²{first}.{last}@glasgow.ac.uk

Abstract. Doc2Query — the process of expanding the content of a document before indexing using a sequence-to-sequence model — has emerged as a prominent technique for improving the first-stage retrieval effectiveness of search engines. However, sequence-to-sequence models are known to be prone to “hallucinating” content that is not present in the source text. We argue that Doc2Query is indeed prone to hallucination, which ultimately harms retrieval effectiveness and inflates the index size. In this work, we explore techniques for filtering out these harmful queries prior to indexing. We find that using a relevance model to remove poor-quality queries can improve the retrieval effectiveness of Doc2Query by up to 16%, while simultaneously reducing mean query execution time by 30% and cutting the index size by 48%. We release the code, data, and a live demonstration to facilitate reproduction and further exploration.¹

1 Introduction

Neural network models, particularly those based on contextualised language models, have been shown to improve search effectiveness [3]. While some approaches focus on re-ranking document sets from a first-stage retrieval function to improve precision [27], others aim to improve the first stage itself [4]. In this work, we focus on one of these first-stage approaches: Doc2Query [29]. This approach trains a sequence-to-sequence model (e.g., T5 [33]) to predict queries that may be relevant to a particular text. Then, when indexing, this model is used to *expand* the document by generating a collection of queries and appending them to the document. Though computationally expensive at index time [34], this approach has been shown to be remarkably effective even when retrieving using simple lexical models like BM25 [28]. Numerous works have shown that the approach can produce a high-quality pool of results that are effective for subsequent stages in the ranking pipeline [19, 20, 23, 40].

However, sequence-to-sequence models are well-known to be prone to generate content that does not reflect the input text – a defect known in literature as “hallucination” [25]. We find that existing Doc2Query models are no exception. Figure 1 provides example generated queries from the state-of-the-art T5 Doc2Query model [28]. In this example, we see that many of the generated queries cannot actually be answered by the source passage (score ≤ 1).

¹ https://github.com/terrierteam/pyterrier_doc2query

Original Passage: Barley (<i>Hordeum vulgare</i> L.), a member of the grass family, is a major cereal grain. It was one of the first cultivated grains and is now grown widely. Barley grain is a staple in Tibetan cuisine and was eaten widely by peasants in Medieval Europe. Barley has also been used as animal fodder, as a source of fermentable material for beer and certain distilled beverages, and as a component of various health foods.	Generated Queries: (1) where does barley originate from · (2) what is the name of the cereal grain used in tibetan cooking? · (3) what is barley used for · (1) what is barley in food · (0) what is bare wheat · (3) what family of organisms is barley in · (1) why is barley important in tibetan diet · (3) what is barley · (2) where is barley grown · (1) where was barley first grown and eaten · (1) where was barley first used ...
--	--

Fig. 1. Example passage from MS MARCO and generated queries using the T5 Doc2Query model. The relevance of each query to the passage is scored by the authors on a scale of 0–3 using the TREC Deep Learning passage relevance criteria.

Based on this observation, we hypothesise that retrieval performance of Doc2Query would improve if hallucinated queries were removed. In this paper, we conduct experiments where we apply a new filtering phase that aims to remove poor queries prior to indexing. Given that this approach *removes* queries, we call the approach Doc2Query-- (Doc2Query-minus-minus). Rather than training a new model for this task, we identify that relevance models are already fit for this purpose: they estimate how relevant a passage is to a query. We therefore explore filtering strategies that make use of existing neural relevance models.

Through experimentation on the MS MARCO dataset, we find that our filtering approach can improve the retrieval effectiveness of indexes built using Doc2Query-- by up to 16%; less can indeed be more. Meanwhile, filtering naturally reduces the index size, lowering storage and query-time computational costs. Finally, we conduct an exploration of the index-time overheads introduced by the filtering process and conclude that the gains from filtering more than make up for the additional time spent generating more queries. The approach also has a positive impact on the environmental costs of applying Doc2Query; the same retrieval effectiveness can be achieved with only about a third of the computational cost when indexing. To facilitate last-metre, last-mile, and complete reproduction efforts [36], we release the code, indices, and filtering scores.¹ In summary, we contribute a technique to improve the effectiveness and efficiency of Doc2Query by filtering out queries that do not reflect the original passage.

2 Related Work

The classical lexical mismatch problem is a key one in information retrieval - documents that do not contain the query terms may not be retrieved. In the literature, various approaches have addressed this: query reformulation – including stemming, query expansion models (e.g. Rocchio, Bo1 [1], RM3 [12]) – and document expansion [9, 30, 35]. Classically, query expansion models have been popular, as they avoid the costs associated with making additional processing for each document needed for document expansion. However, query expansion may result in reduced performance [11], as queries are typically short and the necessary evidence to understand the context of the user is limited.

The application of latent representations of queries and documents, such as using latent semantic indexing [8] allow retrieval using to not be driven directly by lexical signals. More recently, transformer-based language models (such as BERT [6]) have resulted in representations of text where the contextualised meaning of words are accounted for. In particular, in dense retrieval, queries and documents are represented in embeddings spaces [14, 37], often facilitated by Approximate Nearest Neighbour (ANN) data structures [13]. However, even when using ANN, retrieval can still be inefficient or insufficiently effective [15].

Others have explored approaches for augmenting lexical representations with additional terms that may be relevant. In this work, we explore Doc2Query [29], which uses a sequence-to-sequence model that maps a document to queries that it might be able to answer. By appending these *generated* queries to a document’s content before indexing, the document is more likely to be retrieved for user queries when using a model like BM25. An alternative style of document expansion, proposed by MacAvaney et al. [19] and since used by several other models (e.g., [10, 39, 40]), uses the built-in Masked Language Modelling (MLM) mechanism. MLM expansion generates individual tokens to append to the document as a bag of words (rather than as a sequence). Although MLM expansion is also prone to hallucination,² the bag-of-words nature of MLM expansion means that individual expansion tokens may not have sufficient context to apply filtering effectively. We therefore focus only on sequence-style expansion and leave the exploration of MLM expansion for future work.

3 Doc2Query--

Doc2Query-- consists of two phases: a *generation* phase and a *filtering* phase. In the generation phase, a Doc2Query model generates a set of n queries that each document might be able to answer. However, as shown in Figure 1, not all of the queries are necessarily relevant to the document. To mitigate this problem, Doc2Query-- then proceeds to a filtering phase, which is responsible for eliminating the generated queries that are least relevant to the source document. Because hallucinated queries contain details not present in the original text (by definition), we argue that hallucinated queries are less useful for retrieval than non-hallucinated ones. Filtering is accomplished by retaining only the most relevant p proportion of generated queries over the entire corpus. The retained queries are then concatenated to their corresponding documents prior to indexing, as per the existing Doc2Query approach.

More formally, consider an expansion function e that maps a document to n queries: $e : \mathbf{D} \mapsto \mathbf{Q}^n$. In Doc2Query, each document in corpus \mathcal{D} are concatenated with their expansion queries, forming a new corpus $\mathcal{D}' = \{\text{Concat}(d, e(d)) \mid d \in \mathcal{D}\}$, which is then indexed by a retrieval system. Doc2Query-- adds a filtering mechanism that uses a relevance model that maps a query and document to a real-valued relevance score $s : \mathbf{Q} \times \mathbf{D} \mapsto \mathbb{R}$ (with larger values indicating higher

² For instance, we find that SPLADE [10] generates the following seemingly-unrelated terms for the passage in Figure 1 in the top 20 expansion terms: *reed*, *herb*, and *troy*.

relevance). The relevance scoring function is used to filter down the queries to those that meet a certain score threshold t as follows:

$$\mathcal{D}' = \left\{ \text{Concat}(d, \{q \mid q \in e(d) \wedge s(q, d) \geq t\}) \mid d \in \mathcal{D} \right\} \quad (1)$$

The relevance threshold t is naturally dependent upon the relevance scoring function. It can be set empirically, chosen based on operational criteria (e.g., target index size), or (for a well-calibrated relevance scoring function) determined *a priori*. In this work, we combine the first two strategies: we pick t based on the distribution of relevance scores across all expansion queries. For instance, at $p = 0.3$ we only keep queries with relevance scores in the top 30%, which is $t = 3.215$ for the ELECTRA [31] scoring model on the MS MARCO dataset [26].

4 Experimental Setup

We conduct experiments to answer the following research questions:

- RQ1 Does Doc2Query-- improve the effectiveness of document expansion?
 RQ2 What are the trade-offs in terms of effectiveness, efficiency, and storage when using Doc2Query--?

Datasets and Measures. We conduct tests using the MS MARCO [26] v1 passage corpus. We use five test collections:³ (1) the MS MARCO Dev (small) collection, consisting of 6,980 queries (1.1 qrels/query); (2) the Dev2 collection, consisting of 4,281 (1.1 qrels/query); (3) the MS MARCO Eval set, consisting of 6,837 queries (held-out leaderboard set); (4/5) the TREC DL'19/'20 collections, consisting of 43/54 queries (215/211 qrels/query). We evaluate using the official task evaluation measures: Reciprocal Rank at 10 (RR@10) for Dev/Dev2/Eval, nDCG@10 for DL'19/'20. We tune systems⁴ on Dev, leaving the remaining collections as held-out test sets.

Models. We use the T5 Doc2Query model from Nogueira and Lin [28], making use of the inferred queries released by the authors (80 per passage). To the best of our knowledge, this is the highest-performing Doc2Query model available. We consider three neural relevance models for filtering: ELECTRA⁵ [31], MonoT5⁶ [32], and TCT-ColBERT⁷ [16], covering two strong cross-encoder models and one strong bi-encoder model. We also explored filters that use the probabilities from the generation process itself but found them to be ineffective and therefore omit these results due to space constraints.

Tools and Environment. We use the PyTerrier toolkit [22] with a PISA [24, 17] index to conduct our experiments. We deploy PISA's Block-Max WAND [7] implementation for BM25 retrieval. Inference was conducted on an NVIDIA 3090 GPU. Evaluation was conducted using the ir-measures package [18].

³ ir-datasets [21] IDs: `msmarco-passage/dev/small`, `msmarco-passage/dev/2`, `msmarco-passage/eval/small`, `msmarco-passage/trec-dl-2019/judged`, `msmarco-passage/trec-dl-2020/judged` ⁴ BM25's k_1 , b , and whether to remove stopwords were tuned for all systems; the filtering percentage (p) was also tuned for filtered systems. ⁵ `crystina-z/monoELECTRA_LCE_nneg31`
⁶ `castorini/monot5-base-msmarco` ⁷ `castorini/tct_colbert-v2-hnp-msmarco`

Table 1. Effectiveness and efficiency measurements for Doc2Query-- and baselines. Significant differences between Doc2Query and their corresponding filtered versions for Dev, Dev2, DL’19 and DL’20 are indicated with * (paired t-test, $p < 0.05$). Values marked with † are taken from the corresponding submissions to the public leaderboard.

System	RR@10			nDCG@10		ms/q	GB
	Dev	Dev2	Eval	DL’19	DL’20	MRT	Index
BM25	0.185	0.182	†0.186	0.499	0.479	5	0.71
Doc2Query ($n = 40$)	0.277	0.265	†0.272	0.626	0.607	30	1.17
w/ ELECTRA Filter (30%)	*0.316	*0.310	-	0.667	0.611	23	0.89
w/ MonoT5 Filter (40%)	*0.308	*0.298	0.306	0.650	0.611	29	0.93
w/ TCT Filter (50%)	*0.287	*0.280	-	0.640	0.599	30	0.94
Doc2Query ($n = 80$)	0.279	0.267	-	0.627	0.605	30	1.41
w/ ELECTRA Filter (30%)	*0.323	*0.316	0.325	0.670	0.614	23	0.95
w/ MonoT5 Filter (40%)	*0.311	*0.298	-	0.665	0.609	28	1.04
w/ TCT Filter (50%)	*0.293	*0.283	-	0.642	0.588	28	1.05

5 Results

We first explore RQ1: whether relevance filtering can improve the retrieval of Doc2Query models. Table 1 compares the effectiveness of Doc2Query with various filters. We observe that all the filters significantly improve the retrieval effectiveness on the Dev and Dev2 datasets at both $n = 40$ and $n = 80$. We also observe a large boost in performance on the Eval dataset.⁸ Though the differences in DL’19 and DL’20 appear to be considerable (e.g., 0.627 to 0.670), these differences are not statistically significant.

Digging a little deeper, Figure 2 shows the retrieval effectiveness of Doc2Query with various numbers of generated queries (in dotted black) and the corresponding performance when filtering using the top-performing ELECTRA scorer (in solid blue). We observe that performing relevance filtering at each value of n improves the retrieval effectiveness. For instance, keeping only 30% of expansion queries at $n = 80$, performance is increased from 0.279 to 0.323 – a 16% improvement.

In aggregate, results from Table 1 and Figure 2 answer RQ1: Doc2Query-- filtering can significantly improve the retrieval effectiveness of Doc2Query across various scoring models, numbers of generated queries (n) and thresholds (p).

Next, we explore the trade-offs in terms of effectiveness, efficiency, and storage when using Doc2Query--. Table 1 includes the mean response time and index sizes for each of the settings. As expected, filtering reduces the index size since fewer terms are stored. For the best-performing setting ($n = 80$ with ELECTRA

⁸ Significance cannot be determined due to the held-out nature of the dataset. Further, due to restrictions on the number of submissions to the leaderboard, we only are able to submit two runs. The first aims to be a fair comparison with the existing Doc2Query Eval result, using the same number of generated queries and same base T5 model for scoring. The second is our overall best-performing setting, using the ELECTRA filter at $n = 80$ generated queries.




Fig. 2. Effectiveness (RR@10) on the Dev set, compared with the total number of indexed tokens. The generation phase is shown in dotted black (at various values of n), and the ELECTRA filtering phase is shown in solid blue (at various values of p).

filter), this amounts to a 48% reduction in index size (1.41 GB down to 0.95 GB). Naturally, such a reduction has an impact on query processing time as well; it yields a 30% reduction in mean response time (30ms down to 23ms).

Doc2Query-- filtering adds substantial cost an indexing time, mostly due to scoring each of the generated queries. Table 2 reports the cost (in hours of GPU time) of the generation and filtering phases. We observe that ELECTRA filtering can yield up to a 78% increase in GPU time ($n = 10$). However, we find that the improved effectiveness makes up for this cost. To demonstrate this, we allocate the time spent filtering to generating additional queries for each passage. For instance, the 15 hours spent scoring $n = 5$ queries could instead be spent generating 6 more queries per passage (for a total of $n = 11$). We find that when comparing against an unfiltered n that closely approximates the total time when

Table 2. Retrieval effectiveness comparison for comparable indexing computational budgets (in hours of GPU time). Values of n without a filter are chosen to best approximate the total compute hours or the Dev effectiveness of the corresponding filtered version. Significant differences between in RR@10 performance are indicated with * (paired t-test, $p < 0.05$).

n	Filter	GPU Hours	RR@10		Comment
		Gen+Filt=Tot	Dev	Dev2	
5	ELECTRA	20 + 15 = 34	0.273	0.270	
11	None	34 + 0 = 34	*0.261	*0.256	-4% Dev RR for sim. GPU hrs
31	None	99 + 0 = 99	0.273	0.265	$\times 2.9$ GPU hrs to match Dev RR
10	ELECTRA	32 + 25 = 57	0.292	0.292	
18	None	59 + 0 = 59	*0.270	*0.260	-8% Dev RR for sim. GPU hrs
20	ELECTRA	66 + 47 = 113	0.307	0.303	
36	None	113 + 0 = 113	*0.275	*0.265	-10% Dev RR for sim. GPU hrs
40	ELECTRA	128 + 86 = 214	0.316	0.310	
68	None	216 + 0 = 216	*0.279	*0.267	-12% Dev RR for sim. GPU hrs

filtering, the filtered results consistently yield significantly higher retrieval effectiveness. As the computational budget increases, so does the margin between Doc2Query and Doc2Query--, from 4% at 34 hours up to 12% at 216 hours.

From the opposite perspective, Doc2Query consumes $2.9\times$ or more GPU time than Doc2Query-- to achieve similar effectiveness ($n = 13$ with no filter vs. $n = 5$ with ELECTRA filter). Since the effectiveness of Doc2Query flattens out between $n = 40$ and $n = 80$ (as seen in Figure 2), it likely requires a massive amount of additional compute to reach the effectiveness of Doc2Query-- at $n \geq 10$, if that effectiveness is achievable at all. These comparisons show that if a deployment is targeting a certain level of effectiveness (rather than a target compute budget), Doc2Query-- is also preferable to Doc2Query.

These results collectively answer RQ2: Doc2Query-- provides higher effectiveness at lower query-time costs, even when controlling for the additional compute required at index time.

6 Conclusions

This work demonstrated that there are untapped advantages in generating natural-language for document expansion. Specifically, we presented Doc2Query--, which is a new approach for improving the effectiveness and efficiency of the Doc2Query model by filtering out the least relevant queries. We observed that a 16% improvement in retrieval effectiveness can be achieved, while reducing the index size by 48% and mean query execution time by 30%.

The technique of filtering text generated from language models using relevance scoring is ripe for future work. For instance, relevance filtering could potentially apply to approaches that generate alternative forms of queries [38], training data [2], or natural language responses to queries [5] — all of which are potentially affected by hallucinated content. Furthermore, future work could explore approaches for relevance filtering over masked language modelling expansion [19], rather than sequence-to-sequence expansion.

Acknowledgements

Sean MacAvaney and Craig Macdonald acknowledge EPSRC grant EP/R018634/1: Closed-Loop Data Science for Complex, Computationally- & Data-Intensive Analytics.

References

1. Amati, G., Van Rijsbergen, C.J.: Probabilistic models of information retrieval based on measuring the divergence from randomness. *ACM Trans. Inf. Syst.* **20**(4) (2002)
2. Bonifacio, L., Abonizio, H., Fadaee, M., Nogueira, R.: InPars: Unsupervised dataset generation for information retrieval. In: *Proceedings of SIGIR* (2022)

3. Dai, Z., Callan, J.: Deeper text understanding for IR with contextual neural language modeling. In: *Proceedings of SIGIR (2019)*
4. Dai, Z., Callan, J.: Context-aware document term weighting for ad-hoc search. In: *Proceedings of The Web Conference (2020)*
5. Das, R., Dhuliawala, S., Zaheer, M., McCallum, A.: Multi-step retriever-reader interaction for scalable open-domain question answering. In: *Proceedings of ICLR (2019)*
6. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: BERT: Pre-training of deep bidirectional transformers for language understanding. In: *Proceedings of NAACL-HLT (2019)*
7. Ding, S., Suel, T.: Faster top-k document retrieval using block-max indexes. In: *Proceedings of SIGIR (2011)*
8. Dumais, S.T., Furnas, G.W., Landauer, T.K., Deerwester, S., Harshman, R.: Using latent semantic analysis to improve access to textual information. In: *Proceedings of SIGCHI CHI (1988)*
9. Efron, M., Organisciak, P., Fenlon, K.: Improving retrieval of short texts through document expansion. In: *Proceedings of SIGIR (2012)*
10. Formal, T., Piwowarski, B., Clinchant, S.: SPLADE: Sparse lexical and expansion model for first stage ranking. In: *Proceedings of SIGIR (2021)*
11. He, B., Ounis, I.: Studying query expansion effectiveness. In: *Proceedings of ECIR (2009)*
12. Jaleel, N.A., Allan, J., Croft, W.B., Diaz, F., Larkey, L.S., Li, X., Smucker, M.D., Wade, C.: Umass at TREC 2004: Novelty and HARD. In: *TREC (2004)*
13. Johnson, J., Douze, M., Jegou, H.: Billion-scale similarity search with GPUs. *IEEE Transactions on Big Data* **7**(03) (2021)
14. Khattab, O., Zaharia, M.: ColBERT: Efficient and effective passage search via contextualized late interaction over BERT. In: *Proceedings of SIGIR (2020)*
15. Lin, J., Ma, X., Mackenzie, J., Mallia, A.: On the separation of logical and physical ranking models for text retrieval applications. In: *Proceedings of DESIRES (2021)*
16. Lin, S.C., Yang, J.H., Lin, J.: In-batch negatives for knowledge distillation with tightly-coupled teachers for dense retrieval. In: *Proceedings of RepL4NLP (2021)*
17. MacAvaney, S., Macdonald, C.: A Python interface to PISA! In: *Proceedings of SIGIR (2022)*
18. MacAvaney, S., Macdonald, C., Ounis, I.: Streamlining evaluation with ir-measures. In: *Proceedings of ECIR (2022)*
19. MacAvaney, S., Nardini, F.M., Perego, R., Tonellotto, N., Goharian, N., Frieder, O.: Expansion via prediction of importance with contextualization. In: *Proceedings of SIGIR (2020)*
20. MacAvaney, S., Tonellotto, N., Macdonald, C.: Adaptive re-ranking with a corpus graph. In: *Proceedings of CIKM (2022)*
21. MacAvaney, S., Yates, A., Feldman, S., Downey, D., Cohan, A., Goharian, N.: Simplified data wrangling with `ir_datasets`. In: *Proceedings of SIGIR (2021)*
22. Macdonald, C., Tonellotto, N.: Declarative experimentation in information retrieval using PyTerrier. In: *Proceedings of ICTIR (2020)*
23. Mallia, A., Khattab, O., Suel, T., Tonellotto, N.: Learning passage impacts for inverted indexes. In: *Proceedings of SIGIR (2021)*
24. Mallia, A., Siedlaczek, M., Mackenzie, J., Suel, T.: PISA: performant indexes and search for academia. In: *Proceedings of OSIRRC@SIGIR (2019)*
25. Maynez, J., Narayan, S., Bohnet, B., McDonald, R.: On faithfulness and factuality in abstractive summarization. In: *Proceedings of ACL (2020)*

26. Nguyen, T., Rosenberg, M., Song, X., Gao, J., Tiwary, S., Majumder, R., Deng, L.: MS MARCO: A human generated machine reading comprehension dataset. In: Proceedings of CoCo@NIPS (2016)
27. Nogueira, R., Cho, K.: Passage re-ranking with BERT. ArXiv [abs/1901.04085](#) (2019)
28. Nogueira, R., Lin, J.: From doc2query to docttttquery (2019)
29. Nogueira, R., Yang, W., Lin, J.J., Cho, K.: Document expansion by query prediction. ArXiv [abs/1904.08375](#) (2019)
30. Pickens, J., Cooper, M., Golovchinsky, G.: Reverted indexing for feedback and expansion. In: Proceedings of CIKM (2010)
31. Pradeep, R., Liu, Y., Zhang, X., Li, Y., Yates, A., Lin, J.: Squeezing water from a stone: A bag of tricks for further improving cross-encoder effectiveness for reranking. In: Proceedings of ECIR (2022)
32. Pradeep, R., Nogueira, R., Lin, J.: The expando-mono-duo design pattern for text ranking with pretrained sequence-to-sequence models. ArXiv [abs/2101.05667](#) (2021)
33. Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., Zhou, Y., Li, W., Liu, P.J., et al.: Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.* **21**(140) (2020)
34. Scells, H., Zhuang, S., Zuccon, G.: Reduce, reuse, recycle: Green information retrieval research. In: Proceedings of SIGIR (2022)
35. Tao, T., Wang, X., Mei, Q., Zhai, C.: Language model information retrieval with document expansion. In: Proceedings of HLT-NAACL (2006)
36. Wang, X., MacAvaney, S., Macdonald, C., Ounis, I.: An inspection of the reproducibility and replicability of TCT-ColBERT. In: Proceedings of SIGIR (2022)
37. Xiong, L., Xiong, C., Li, Y., Tang, K.F., Liu, J., Bennett, P.N., Ahmed, J., Overwijk, A.: Approximate nearest neighbor negative contrastive learning for dense text retrieval. In: Proceedings of ICLR (2021)
38. Yu, S.Y., Liu, J., Yang, J., Xiong, C., Bennett, P.N., Gao, J., Liu, Z.: Few-shot generative conversational query rewriting. In: Proceedings of SIGIR (2020)
39. Zhao, T., Lu, X., Lee, K.: SPARTA: Efficient open-domain question answering via sparse transformer matching retrieval. arXiv [abs/2009.13013](#) (2020)
40. Zhuang, S., Zuccon, G.: TILDE: Term independent likelihood model for passage re-ranking. In: Proceedings of SIGIR (2021)