



Delft University of Technology

## The Surprising Effectiveness of Rankers Trained on Expanded Queries

Anand, Abhijit; V, Venkatesh; Setty, Vinay; Anand, Avishek

**DOI**

[10.1145/3626772.3657938](https://doi.org/10.1145/3626772.3657938)

**Publication date**

2024

**Document Version**

Final published version

**Published in**

SIGIR '24

**Citation (APA)**

Anand, A., V, V., Setty, V., & Anand, A. (2024). The Surprising Effectiveness of Rankers Trained on Expanded Queries. In *SIGIR '24: Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval* (pp. 2652-2656). ACM. <https://doi.org/10.1145/3626772.3657938>

**Important note**

To cite this publication, please use the final published version (if applicable). Please check the document version above.

**Copyright**

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

**Takedown policy**

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.

***Green Open Access added to TU Delft Institutional Repository***

***'You share, we take care!' - Taverne project***

**<https://www.openaccess.nl/en/you-share-we-take-care>**

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.



# The Surprising Effectiveness of Rankers Trained on Expanded Queries

Abhijit Anand  
aanand@L3S.de  
L3S Research Center  
Hannover, Germany

Vinay Setty  
vsetty@acm.org  
University of Stavanger  
Stavanger, Norway

Venktesh V  
v.Viswanathan-1@tudelft.nl  
Delft University of Technology  
Delft, Netherlands

Avishek Anand  
avishek.anand@tudelft.nl  
Delft University of Technology  
Delft, Netherlands

## ABSTRACT

An significant challenge in text-ranking systems is handling hard queries that form the tail end of the query distribution. Difficulty may arise due to the presence of uncommon, underspecified, or incomplete queries. In this work, we improve the ranking performance of hard or difficult queries while maintaining the performance of other queries. Firstly, we do LLM-based query enrichment for training queries using relevant documents. Next, a specialized ranker is fine-tuned only on the enriched hard queries instead of the original queries. We combine the relevance scores from the specialized ranker and the base ranker, along with a query performance score estimated for each query. Our approach departs from existing methods that usually employ a single ranker for all queries, which is biased towards easy queries, which form the majority of the query distribution. In our extensive experiments on the DL-HARD dataset, we find that a principled query performance based scoring method using base and specialized ranker offers a significant improvement of up to **48.4%** on the document ranking task and up to **25%** on the passage ranking task compared to the baseline performance of using original queries, even outperforming SOTA model.

## CCS CONCEPTS

• Information systems → Retrieval models and ranking.

## KEYWORDS

Document Ranking, Hard Queries, Query Rewriting, QPP

### ACM Reference Format:

Abhijit Anand, Venkatesh V, Vinay Setty, and Avishek Anand. 2024. The Surprising Effectiveness of Rankers Trained on Expanded Queries. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '24)*, July 14–18, 2024, Washington, DC, USA. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3626772.3657938>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

*SIGIR '24, July 14–18, 2024, Washington, DC, USA*

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0431-4/24/07

<https://doi.org/10.1145/3626772.3657938>

## 1 INTRODUCTION

Hard queries are commonly characterized by incomplete, uncommon, incorrect, domain-specific and inherently complex queries [6, 25]. Though the leaderboards of widely-used IR benchmarks such as TREC-DL and MS MARCO indicate impressive performance gains by modern deep neural rankers, there's still significant room for improvement when addressing hard or obstinate queries [6, 25]. This paper aims to address the challenges posed by hard queries towards building effective and robust ranking models.

**The challenge of hard queries:** The presence of hard queries in training sets with large query workloads has two inherent problems. Firstly, the number of hard queries in training sets is small, with most of the queries being relatively easier to rank. Secondly, the relevance factors encoded in the query-document interactions in hard queries are more nuanced and are different from those of easy queries. As a consequence, training a single-ranking model that is agnostic to query hardness results in rankers that perform well on easy queries that dominate the training set. However, these rankers struggle with hard or “obstinate” queries [6, 25]. In this paper, we propose to train a specialized ranker (SR) for hard queries, in addition to the general ranker, to capture the subtle yet different query-document features induced by hard queries. For example, as shown in Table 1, given a query “anthropological definition of environment”, a base ranker retrieves a document that simply defines “anthropology”, while a SR can retrieve a more relevant document about “environmental anthropology”.

The presence of two ranking models – a specialized ranker (SR) for hard queries, and a base ranker (BR) for the rest of the queries – poses two challenges: (1) *how to automatically determine which queries are hard during inference?* and (2) *how to train a specialized ranking model for hard queries?*

**Automated identification of hard queries:** We employ *query performance prediction (QPP)* approaches to automatically identify hard queries [5]. QPP has proved useful in determining the quality of the retrieval system on a range of queries [8, 40, 41]. Prior works have also demonstrated that estimates obtained using QPP [5, 14, 26] are good indicators of various query characteristics, including difficulty levels. In our approach, queries are first ranked by both the SR and BR. This is followed by a *fusion mechanism* to aggregate both scores based on a hardness estimate that is determined using QPP on the input query. This approach of combining the merits

<b>Query:</b>	anthropological definition of environment
<b>BR</b>	Anthropology is the scientific study of human beings as social organisms interacting with each other in their environment, and cultural aspects [...]
<b>SR</b>	Environmental anthropology is a sub-specialty within the field of anthropology that takes an active role in examining the relationships between humans and their environment across space and time[...]

**Table 1: Comparing top-1 document retrieved by Base Ranker (BR) and our Specialized Ranker (SR) for a hard query.**

of both rankers provides an automatic and robust way of ranking documents for any arbitrary queries.

**Context Aware Query enrichment:** Our second major challenge is training the SR for hard queries. A ranker, trained on hard queries in their original form, does not capture the query-document features for hard queries due to the lack of context. To help the SR identify the specific query-document features, our second novel contribution is to *contextually expand* the hard queries. Specifically, we hypothesize that if the hard queries are expanded or enriched with the active knowledge of the relevant documents, then the SR can extract relevant query-document features when training on the expanded queries. Towards this, we propose a *context-aware query enrichment approach*, to rewrite queries using LLMs. This method incorporates knowledge from the relevant documents as context during training, enhancing the rankers’ ability to learn the complex relationship between query-document pairs for hard queries.

We conduct extensive experiments on DL-HARD [25] (hard queries from MS MARCO) and TREC-DL [42] document datasets to showcase the performance benefits of using SR. We find that during inference, our SR are surprisingly effective for hard queries without needing any query expansions. Our approach shows significant performance gain of **20.2%** and **48.4%** in nDCG@10 and RR respectively compared to baseline and outperforming the SOTA model. We believe that this observation has an impact on the design of multiple ranking models for ranking to handle different query types.

## 1.1 Related Work

The field of IR has long explored the concept of query expansions, with both classical [4, 19, 20, 31, 38, 42] and recent Large Language Model (LLM)-based methodologies [18, 21, 35–37, 39, 43]. Distinct from these methods, which typically use expansion during the inference, our method uniquely applies expansion only during the training phase. Our technique is also different from document expansion methods [10, 29] and document augmentation strategies [2, 3, 7, 13, 32, 33], which target document modification rather than query rewriting. Furthermore, these generative expansion approaches are context-unaware, leading to topic drift. Existing strategies to mitigate such drift [1, 12] are again primarily implemented during inference, which is different from our approach.

## 2 METHODOLOGY

### 2.1 Query Enrichment

Given a training set, consider a set of queries  $Q = \{HQ \cup Q^+\}$ , consisting of a set of hard queries  $HQ$  and a set of easy queries  $Q^+$ . Given a hard query  $hq_i \in HQ$ , we generate a rewritten query  $q_i^*$  that enriches the hard query to result in a more clear and concise query for a ranker.  $q_i^*$  is generated by conditioning the LLM on a prompt using the query and top most relevant document  $d^+$  from a

document collection  $D$ , known as *context-aware query enrichment*. Formally,  $q_i^* = \text{LLM}(hq_i, d^+)$ . Note that there are no principled methods for automatically determining  $HQ$ . We use heuristics such as query length [34], presence of uncommon (acronyms and entities with varied semantics) or incorrect terms and lack of context (under-specified) which are defined according to characteristics of obstinate queries in DL-HARD and MS MARCO Chameleon datasets.

**Passage selector for documents:** Using the full document in the prompt for query enrichment could result in topic drift, since a document may consist of multiple aspects and different topics. To mitigate this, we employ supervised passage selection techniques (Attention, Linear) as proposed in [23]. This selects the most relevant passage aligned with the query.

### 2.2 Training Ranker

Next, our goal is to train a model for *re-ranking* documents. We employ the pointwise loss to train the ranker. Given a query-document pair  $(q_i^*, d_i)$  as input, the ranker model outputs a relevance score. The ranking task is cast as a binary classification problem, where each training instance  $x_i = (q_i^*, d_i)$  is a query-document pair and  $y_i \in \{0, 1\}$  is a relevance label. We train two rankers, a Specialized Ranker (SR) trained on hard enriched query using *context-aware query enrichment* and a Base Ranker (BR) on original queries ( $Q$ ).

### 2.3 Specialized Ranking of Hard Queries

During inference, we use both SR and BR as below:

**2.3.1 Balanced Score Fusion (BSF):** To improve ranking performance for all queries, we propose to aggregate ranked lists from BR and SR during inference using the CombSUM technique [16]. Formally, for a given test query,  $q_i$  and corresponding document  $d_i$ ,  $r_i = BR(q_i, d_i)$ ;  $r_j = SR(q_i, d_i)$ . The final relevance score is calculated as the sum of document retrieval scores:  $r_i + r_j$ .

**2.3.2 Routing using Query Performance Prediction (R-QPP):** Using capabilities of SR and BR on specific queries should improve the ranking performance. To check if a query is hard during inference time, we use a Query Performance Prediction (QPP) model inspired by BERT-QPP [5]. We train a BERT-QPP model as per the original implementation optimized for nDCG@10. At inference time, we score queries using a trained BERT-QPP model and use the score to decide if a query is hard or easy. Hard queries are evaluated by SR and easy queries by BR and finally combine their ranklist.

**2.3.3 Weighted QPP Scoring (W-QPPS):** Finally, we propose a QPP based approach which combines the two ranking scores (BR and SR) weighted by the QPP score. Given a query ( $q_i$ ) at inference time, we first obtain a score which is the *hardness estimate* for  $q_i$  using BERT-QPP model ( $\psi$ ) defined above given the retrieved set.

$$\mathbb{P}(\text{hard}|(q_i, D_k^+)) = \psi(q_i, D_k^+)$$

Then the hardness estimate  $\psi(q_i, D_k^+)$  is used to compute final relevance score  $s_i$  by interpolating the relevance scores obtained by BR or  $s_{BR}(q_i, d_i)$  and SR or  $s_{SR}(q_i, d_i)$ .

$$s_i = \psi(q_i, D_k^+) * s_{SR}(q_i, d_i) + (1 - \psi(q_i, D_k^+)) * s_{BR}(q_i, d_i) \quad (1)$$

Intuitively, the higher the QPP score, more weight will be given to the specialized model score. The resulting score  $s_i$  is used to rank documents. We show in the experiments (Section 4.3) that QPP based scoring improves performance on hard queries without compromising performance on other queries.

### 3 EXPERIMENTAL SETUP

This section describes the experimental setup for three research questions – **RQ1**: How effective is the Specialized Ranker (SR) for hard queries? **RQ2**: What is the impact of QPP in identifying hard queries? **RQ3**: How effective is QPP based scoring for ranking?

#### 3.1 Datasets

**3.1.1 MS MARCO (Document dataset).** : We consider the document dataset from the TREC Deep Learning (DL) track (2019) [28]. Since, LLM based query enrichment is expensive, inspired by *selective query expansion* [12, 27], we randomly sample a small set of queries that satisfy several heuristics from the MS MARCO training set(described in Section 2.1). For training the specialized ranker, this approach results in the selection of 1200 hard queries from MS MARCO document training set, and then we rewrite these queries using our query rewrite models (Subsection 3.2). We evaluate our proposed model on TREC-DL-19 and TREC-DL-20, comprising 43 distinct queries each.

**3.1.2 DL-HARD dataset.** : Contains an annotated collection of 50 hard queries collected according to several criteria defined by Mackie et al. [25]. DL-HARD queries are primarily non-factoid.

#### 3.2 Query Rewriting Models

**3.2.1 QUERY2DOC.** We use a recently proposed document expansion approach, QUERY2DOC which generates pseudo documents to aid in document ranking, as a baseline. We use the gpt-3.5-turbo model with max tokens of 128 for generation and follow the original hyperparameters used by [36] for reproducibility. We further use the rewrites to train a BERT-base model for ranking.

**3.2.2 Chain-of-Thought (CoT).** We compare against the recent LLM based query expansion approach, which uses zero-shot chain of thought prompting to improve the query [21]. We use same hyperparameters as the paper to train ranker using these rewrites.

#### 3.3 Ranking Models

We use the cross-attention BERT-base [15] architecture(12 layer) for ranking. The input length is restricted to a maximum of 512 tokens. The BR is trained on the original set of queries using a pointwise ranking loss objective with learning rate of 1e-5 and 3e-3 for passage and document train set respectively. SR is trained with learning rate of 4e-4 for both sets. For all experiments, we use nDCG@10 and RR metrics for comparison.

MS HARD Document		
Ranking Models	RR	nDCG <sub>10</sub>
<b>Baselines</b>		
BERT-base (BR)	0.444	0.324
QUERY2Doc [36]	0.510(▲14.9%)	0.248(▼23.5%)*
CoT [21]	0.566(▲27.5%)	0.336(▲3.7%)#
BM25 <sup>†</sup>	0.368(▼17.1%)	0.272(▼16%)
BERT-MaxP (ZS) <sup>†</sup>	0.405(▼8.8%)	0.310(▼4.3%)
BERT-MaxP <sup>†</sup>	0.402(▼9.5%)	0.317(▼2.1%)
RM3+BERT-MaxP (ZS) <sup>†</sup> [22]	0.415(▼6.5%)	0.314(▼3%)
RM3+BERT-MaxP <sup>†</sup> [22]	0.443(▼0.2%)	0.295(▼8.9%)
T5-MaxP(ZS) <sup>†</sup> [30]	0.367(▼17.3%)	0.327(▲1.0%)
RM3+T5-MaxP(ZS) <sup>†</sup> [30]	0.359(▼19.1%)	0.307(▼5.2%)
Electra-MaxP <sup>†</sup> [11]	0.448(▲0.9%)	0.385(▲18.9%)
RM3+Electra-MaxP <sup>†</sup> [11]	0.461(▲3.8%)	0.380(▲17.4%)
PARADE-BERT <sup>†</sup> [24]	0.413(▼7.0%)	0.299(▼7.7%)
RM3+PARADE-BERT <sup>†</sup> [24]	0.419(▼5.6%)	0.313(▼3.3%)
PARADE-Electra <sup>†</sup> [24] [11]	0.498(▲12.2%)	0.356(▲9.9%)
RM3+PARADE-Electra <sup>†</sup> [24] [11]	0.489(▲10.1%)	0.357(▲10.3%)
<b>Our approach WITH query rewriting during inference</b>		
Specialised Ranker (SR)	0.345(▼22.4%)	0.257(▼20.6%)
Balanced Score Fusion (BSF)	0.568(▲28.0%)*	0.332(▲2.6%)*
Routing using QPP (R-QPP)	0.595(▲33.9%)*	0.334(▲3.2%)*
Weighted QPP Scoring (W-QPPS)	0.613(▲38.5%)*	0.335(▲3.5%)*
<b>Our approach W/O query rewriting during inference</b>		
Specialised Ranker (SR)	0.421(▼5.2%)	0.316(▼2.4%)
Balanced Score Fusion (BSF)	0.535(▲20.5%)*	0.368(▲13.7%)*
Routing using QPP (R-QPP)	<u>0.618(▲39.3%)*</u>	<u>0.382(▲17.9%)*</u>
Weighted QPP Scoring (W-QPPS)	<u>0.659(▲48.4%)*</u>	<u>0.389(▲20.2%)*</u>

**Table 2: Results on DL-HARD. Relative improvement are shown over BR. Statistically significant improvements at levels 95% (\*) and 90% (#) are reported [17]. <sup>†</sup> indicate values taken from DL-HARD leaderboard. Second best results are underlined.**

## 4 RESULTS

### 4.1 Effect of Specialized Ranker (SR)

Table 2, shows the comparison of our approaches with baseline models described in Section 3.2 and 3.3. Other baseline model performances are taken from the DL-HARD leaderboard. We also show the performance of our approach with both the original and the rewritten test queries. To answer **RQ1**, from Table 2 it is clear that just using SR does not outperform the baseline (BR). But when used in conjunction with BR using *BSF*, we see improvement of about **14%** and **21%** on nDCG@10 and RR respectively. This is because BR does better on easy queries and SR on hard queries as evident with the performance of *R-QPP*. Additionally, the performance of SR on hard queries can be attributed to its ability to learn relevance patterns specific to hard queries due to its training on contextually enriched queries. Further support is evident as we observe a performance improvement from *BSF* to *R-QPP*, where in *R-QPP* only hard queries are accessed by SR. While *BSF* has impressive gains, it still does not beat the performance of some baseline models (Electra-MaxP, RM3+Electra-MaxP). All our models outperform existing

query expansion approaches like QUERY2DOC and CoT. On further analysis, we observe that the drop in performance for QUERY2DOC is due to the topic drift in the rewrites obtained from QUERY2DOC. We observed the same in our experiments and mitigated topic drift by using a passage selector for documents (Section 2.1).

**4.1.1 Effect of Query Rewriting during inference.** We also try inference using rewritten queries, BR infers on original queries and SR on rewritten queries to keep it consistent with their respective training techniques. In Table 2, we see that rewriting test queries perform better than BR but do not outperform some other baseline models or our approach using original test queries. This is surprising enough but expected as rewriting hard queries does not always do well [9]. We theorize it does badly as the expanded/rewritten query does not have context, resulting in topic drift, contrary to the queries SR is trained on which have context.

**Insight 1:** *Our specialized ranker improves ranking performance on hard queries and also for general queries when used with BR.*

## 4.2 Effectiveness of document ranking using Query Performance Predictor (QPP)

In Table 2, we observe that *BSF* offers lesser gains compared to other approaches. This is primarily because *BSF* does not consider the nature of the query, and it applies equal weights to scores from BR and SR. However, *R-QPP* which considers the characteristics of individual queries and perform query routing accordingly provides significant gains. We observe an overall performance improvement of about **18%** and **39%** on *nDCG@10* and *RR*, respectively. Hence proving that using QPP along with SR and BR for ranking improves performance, answering **RQ2**.

On closer inspection, we observe that SR provides a significant gain in query performance for hard queries relative to the BR. For instance, the average *nDCG@10* for hard queries identified by QPP using BR is 0.278 and on the other hand, SR is 0.345 (+24% over BR).

**Insight 2:** *We find that query routing using QPP-based hardness estimation improves ranking performance by 18% on nDCG@10.*

## 4.3 Effectiveness of QPP based scoring approach

To answer **RQ3**, we evaluate *W-QPPS* using DL-HARD. The routing of the queries helps in choosing the ranker based on the nature of the queries. However, since the threshold to determine hardness decides the nature of the query, it could sometimes result in a misclassification of the queries. Using *W-QPPS*, we can leverage the estimated score directly to weigh the relevance scores from the ranking models that help combine the capabilities of both models for each query. We observe that our proposed approach *W-QPPS* provides the best performance with a gain of **20.2%** in *nDCG@10* compared to BR as shown in Table 2. This method even outperforms the SOTA Electra-MaxP [11]. It is interesting to note that the BERT model we use is BERT-base model whereas other baselines use BERT-Large and T5 models that have double the parameters of our model. This improvement of our approach over larger models also proves the effectiveness of QPP score based method(*W-QPPS*).

**Insight 3:** *Weighted QPP score-based method combines the best capabilities of both rankers, weighing them according to the hardness of the query and beating SOTA model with 20% improvement.*

Ranking Models	TREC-DL-Document	
	2019	2020
BERT-base (BR)	0.616	0.600
QUERY2Doc [36]	0.613(▲1.2%)	0.606(▲1.0%)
CoT [21]	0.590(▼4.3%)*	0.599(▲0%)
<b>Our Approach</b>		
Balanced Score Fusion( <i>BSF</i> )	0.622(▲1.0%)#	<b>0.618(▲3.1%)#</b>
Routing using QPP( <i>R-QPP</i> )	0.622(▲1.1%)#	0.610(▲1.7%)#
Weighted QPP Scoring( <i>W-QPPS</i> )	<b>0.625(▲1.5%)*</b>	0.590(▼1.7%)*

**Table 3: Results on TREC-DL-19 and TREC-DL-20. Relative improvement are shown over BR. Statistically significant improvements at levels 95% (\*) and 90% (#) are reported [17].**

## 4.4 Effectiveness on general queries

We also evaluate general ranking datasets like TREC-DL 19 and TREC-DL 20, document collections using our approaches. In Table 3 we observe a slight performance improvement of our model against baseline, this is because around 5% of queries in the dataset are hard, hence the contribution of SR in *R-QPP* or *W-QPPS* is very small to make a big difference in the overall ranking. We see *BSF* is the best approach. This illustrates that our approaches work on both general and hard-ranking datasets. We also experiment with rewritten queries and see similar results as in the case of DL-HARD.

We did additional experiments on DL-HARD, TREC-DL 19, and TREC-DL 20 passage datasets, and the results had similar trends as with their respective document collection approaches. For DL-HARD passage, the best performing model was *W-QPPS* with a gain of **10.4%** and **28.5%** over BR on *nDCG@10* and *RR*, respectively. Additionally, our approach showed a gain of **3.1%** and **3.4%** over respective BR models for TREC-DL 19 and TREC-DL 20 respectively. All results are not shown due to space constraints.

## 5 CONCLUSION

In this work, we propose a framework for improving the document ranking for hard queries without sacrificing the performance on other queries. We accomplish this by training a specialized document ranker on hard queries rewritten using context aware query enrichment using LLMs. We then perform query performance estimation using a neural scoring mechanism. Using the query performance scores as indicators of the degree of hardness of the queries, we propose a principled combination of the relevance scores from the base ranker and the specialized ranker. Through extensive experiments on diverse datasets, we demonstrate that the proposed approach offers performance gains for the ranking tasks on hard queries without sacrificing performance on other queries, even outperforming **SOTA** model on DL-HARD document test set. In the future, we plan to extend our approach to characterize diverse types of queries and also propose an end to end optimization of the query rewriter and ranker.

## ACKNOWLEDGMENTS

This work is partially supported by German Research Foundation (DFG), under the Project IREM with grant No. AN 996/1-1.

## REFERENCES

- [1] G. Amati, C. Carpineto, and G. Romano. Query difficulty, robustness, and selective application of query expansion. In S. McDonald and J. Tait, editors, *Advances in Information Retrieval*, pages 127–137, Berlin, Heidelberg, 2004. Springer Berlin Heidelberg. ISBN 978-3-540-24752-4.
- [2] A. Anand, J. Leonhardt, K. Rudra, and A. Anand. Supervised contrastive learning approach for contextual ranking. In *Proceedings of the 2022 ACM SIGIR International Conference on Theory of Information Retrieval*, pages 61–71, 2022.
- [3] A. Anand, J. Leonhardt, J. Singh, K. Rudra, and A. Anand. Data augmentation for sample efficient and robust document ranking. *ACM Transactions on Information Systems*, 2023.
- [4] N. Arabzadeh, A. Bigdeli, S. Seyedsalehi, M. Zihayat, and E. Bagheri. Matches made in heaven: Toolkit and large-scale datasets for supervised query reformulation. In *Proceedings of the 30th ACM International Conference on Information and Knowledge Management*, pages 4417–4425, 2021.
- [5] N. Arabzadeh, M. Khodabakhsh, and E. Bagheri. Bert-qp: Contextualized pre-trained transformers for query performance prediction. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*, CIKM '21, page 2857–2861, New York, NY, USA, 2021. Association for Computing Machinery. ISBN 9781450384469. doi: 10.1145/3459637.3482063. URL <https://doi.org/10.1145/3459637.3482063>.
- [6] N. Arabzadeh, B. Mitra, and E. Bagheri. Ms marco chameleons: Challenging the ms marco leaderboard with extremely obstinate queries. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*, CIKM '21, page 4426–4435, New York, NY, USA, 2021. Association for Computing Machinery. ISBN 9781450384469. doi: 10.1145/3459637.3482011. URL <https://doi.org/10.1145/3459637.3482011>.
- [7] L. Bonifacio, H. Abonizio, M. Fadaee, and R. Nogueira. Inpars: Data augmentation for information retrieval using large language models. *arXiv preprint arXiv:2202.05144*, 2022.
- [8] D. Carmel, E. Yom-Tov, A. Darlow, and D. Pelleg. What makes a query difficult? In *Proceedings of the 29th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '06, page 390–397, New York, NY, USA, 2006. Association for Computing Machinery. ISBN 1595933697. doi: 10.1145/1148170.1148238. URL <https://doi.org/10.1145/1148170.1148238>.
- [9] C. Carpineto and G. Romano. A survey of automatic query expansion in information retrieval. *ACM Comput. Surv.*, 44(1), jan 2012. ISSN 0360-0300. doi: 10.1145/2071389.2071390. URL <https://doi.org/10.1145/2071389.2071390>.
- [10] D. R. Cheriton. From doc2query to doc2ttquery. 2019.
- [11] K. Clark, M.-T. Luong, Q. V. Le, and C. D. Manning. Electra: Pre-training text encoders as discriminators rather than generators. *arXiv preprint arXiv:2003.10555*, 2020.
- [12] S. Cronen-Townsend, Y. Zhou, and W. B. Croft. A framework for selective query expansion. In *Proceedings of the Thirteenth ACM International Conference on Information and Knowledge Management*. CIKM '04, page 236–237, New York, NY, USA, 2004. Association for Computing Machinery. ISBN 1581138741. doi: 10.1145/1031171.1031220. URL <https://doi.org/10.1145/1031171.1031220>.
- [13] Z. Dai, V. Y. Zhao, J. Ma, Y. Luan, J. Ni, J. Lu, A. Bakalov, K. Guo, K. B. Hall, and M.-W. Chang. Promptagator: Few-shot dense retrieval from 8 examples. *arXiv preprint arXiv:2209.11755*, 2022.
- [14] S. Datta, D. Ganguly, D. Greene, and M. Mitra. Deep-qp: A pairwise interaction-based deep learning model for supervised query performance prediction. In *Proceedings of the Fifteenth ACM International Conference on Web Search and Data Mining*, WSDM '22, page 201–209, New York, NY, USA, 2022. Association for Computing Machinery. ISBN 9781450391320. doi: 10.1145/3488560.3498491. URL <https://doi.org/10.1145/3488560.3498491>.
- [15] J. Devlin, M. Chang, K. Lee, and K. Toutanova. BERT: pre-training of deep bidirectional transformers for language understanding. *CoRR*, abs/1810.04805, 2018. URL <http://arxiv.org/abs/1810.04805>.
- [16] E. A. Fox and J. A. Shaw. Combination of multiple searches. In *Text Retrieval Conference*, 1993. URL <https://api.semanticscholar.org/CorpusID:1309301>.
- [17] L. Gallagher. Pairwise t-test on TREC Run Files. <https://github.com/lgrz/pairwise-ttest/>, 2019.
- [18] M. Gospodinov, S. MacAvaney, and C. Macdonald. DocQuery-: When less is more, 2023.
- [19] M. Grbovic, N. Djuric, V. Radosavljevic, F. Silvestri, and N. Bhamidipati. Context-and content-aware embeddings for query rewriting in sponsored search. In *Proceedings of the 38th international ACM SIGIR conference on research and development in information retrieval*, pages 383–392, 2015.
- [20] Y. He, J. Tang, H. Ouyang, C. Kang, D. Yin, and Y. Chang. Learning to rewrite queries. In *Proceedings of the 25th ACM International Conference on Information and Knowledge Management*, pages 1443–1452, 2016.
- [21] R. Jagerman, H. Zhuang, Z. Qin, X. Wang, and M. Bendersky. Query expansion by prompting large language models. *arXiv preprint arXiv:2305.03653*, 2023.
- [22] V. Lavrenko and W. B. Croft. Relevance-based language models. In *ACM SIGIR Forum*, volume 51, pages 260–267. ACM New York, NY, USA, 2017.
- [23] J. Leonhardt, K. Rudra, and A. Anand. Extractive explanations for interpretable text ranking. *ACM Trans. Inf. Syst.*, dec 2022. ISSN 1046-8188. doi: 10.1145/3576924. URL <https://doi.org/10.1145/3576924>.
- [24] C. Li, A. Yates, S. MacAvaney, B. He, and Y. Sun. Parade: Passage representation aggregation for document reranking. *ACM Transactions on Information Systems*, 2020.
- [25] I. Mackie, J. Dalton, and A. Yates. How deep is your learning: the dl-hard annotated deep learning dataset, 2021.
- [26] C. Meng, N. Arabzadeh, M. Aliannejadi, and M. de Rijke. Query performance prediction: From ad-hoc to conversational search. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM, jul 2023. doi: 10.1145/3539618.3591919. URL <https://doi.org/10.1145/3539618.3591919>.
- [27] J. Mothe and M. Z. Ullah. Selective query processing: A risk-sensitive selection of search configurations. *ACM Trans. Inf. Syst.*, 42(1), aug 2023. ISSN 1046-8188. doi: 10.1145/3608474. URL <https://doi.org/10.1145/3608474>.
- [28] T. Nguyen, M. Rosenberg, X. Song, J. Gao, S. Tiwary, R. Majumder, and L. Deng. MS MARCO: A human generated machine reading comprehension dataset. In T. R. Besold, A. Bordes, A. S. d'Ávila Garcez, and G. Wayne, editors, *Proceedings of the Workshop on Cognitive Computation: Integrating neural and symbolic approaches 2016 co-located with the 30th Annual Conference on Neural Information Processing Systems (NIPS 2016), Barcelona, Spain, December 9, 2016*, volume 1773 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2016. URL [http://ceur-ws.org/Vol-1773/CoCoNIPS\\_2016\\_paper9.pdf](http://ceur-ws.org/Vol-1773/CoCoNIPS_2016_paper9.pdf).
- [29] R. Nogueira, W. Yang, J. Lin, and K. Cho. Document expansion by query prediction, 2019. URL <https://arxiv.org/abs/1904.08375>.
- [30] C. Raffel, N. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, and P. J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *The Journal of Machine Learning Research*, 21(1):5485–5551, 2020.
- [31] S. Rao and H. Daumé III. Learning to ask good questions: Ranking clarification questions using neural expected value of perfect information. *arXiv preprint arXiv:1805.04655*, 2018.
- [32] K. Rudra and A. Anand. Distant supervision in bert-based adhoc document retrieval. In *Proceedings of the 29th ACM International Conference on Information and Knowledge Management*, pages 2197–2200, 2020.
- [33] K. Rudra, Z. T. Fernando, and A. Anand. An in-depth analysis of passage-level label transfer for contextual document ranking. *Information Retrieval Journal*, 26(1):13, 2023.
- [34] M. Sanderson. Ambiguous queries: Test collections need more sense. In *Proceedings of the 31st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '08, page 499–506, New York, NY, USA, 2008. Association for Computing Machinery. ISBN 9781605581644. doi: 10.1145/1390334.1390420. URL <https://doi.org/10.1145/1390334.1390420>.
- [35] J. Trienes and K. Balog. Identifying unclear questions in community question answering websites. In *Advances in Information Retrieval: 41st European Conference on IR Research, EDIR 2019, Cologne, Germany, April 14–18, 2019, Proceedings, Part 1 41*, pages 276–289. Springer, 2019.
- [36] L. Wang, N. Yang, and F. Wei. Query2doc: Query expansion with large language models, 2023.
- [37] X. Wang, S. MacAvaney, C. Macdonald, and I. Ounis. Generative query reformulation for effective adhoc search. *arXiv preprint arXiv:2308.00415*, 2023.
- [38] H. Zamani and W. B. Croft. Relevance-based word embedding. In N. Kando, T. Sakai, H. Joho, H. Li, A. P. de Vries, and R. W. White, editors, *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval, Shinjuku, Tokyo, Japan, August 7–11, 2017*, pages 505–514. ACM, 2017. doi: 10.1145/3077136.3080831. URL <https://doi.org/10.1145/3077136.3080831>.
- [39] H. Zamani, S. Dumais, N. Craswell, P. Bennett, and G. Lueck. Generating clarifying questions for information retrieval. In *Proceedings of the web conference 2020*, pages 418–428, 2020.
- [40] O. Zendel, A. Shtok, F. Raiber, O. Kurland, and J. S. Culpepper. Information needs, queries, and query performance prediction. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR'19, page 395–404, New York, NY, USA, 2019. Association for Computing Machinery. ISBN 9781450361729. doi: 10.1145/3331184.3331253. URL <https://doi.org/10.1145/3331184.3331253>.
- [41] O. Zendel, A. Shtok, F. Raiber, O. Kurland, and J. S. Culpepper. Information needs, queries, and query performance prediction. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR'19, page 395–404, New York, NY, USA, 2019. Association for Computing Machinery. ISBN 9781450361729. doi: 10.1145/3331184.3331253. URL <https://doi.org/10.1145/3331184.3331253>.
- [42] G. Zerveas, R. Zhang, L. Kim, and C. Eickhoff. Brown university at trec deep learning 2019, 2020.
- [43] S. Zuo, Q. Yin, H. Jiang, S. Xi, B. Yin, C. Zhang, and T. Zhao. Context-aware query rewriting for improving users' search experience on e-commerce websites, 2022.