



Mackie, I., Chatterjee, S. and Dalton, J. (2023) Generative Relevance Feedback with Large Language Models. In: SIGIR 2023: 46th International ACM SIGIR Conference on Research and Development in Information Retrieval, Taipei, Taiwan, 23-27 Jul 2023, pp. 2026-2031. ISBN 9781450394086 (doi: [10.1145/3539618.3591992](https://doi.org/10.1145/3539618.3591992))



Copyright © 2023 The Authors. Reproduced under a [Creative Commons Attribution 4.0 International License](https://creativecommons.org/licenses/by/4.0/).

For the purpose of open access, the author(s) has applied a Creative Commons Attribution license to any Accepted Manuscript version arising.

<https://eprints.gla.ac.uk/297215/>

Deposited on: 26 April 2023

# Generative Relevance Feedback with Large Language Models

Iain Mackie  
University of Glasgow  
i.mackie.1@research.gla.ac.uk

Shubham Chatterjee  
University of Glasgow  
shubham.chatterjee@glasgow.ac.uk

Jeffrey Dalton  
University of Glasgow  
jeff.dalton@glasgow.ac.uk

## ABSTRACT

Current query expansion models use pseudo-relevance feedback to improve first-pass retrieval effectiveness; however, this fails when the initial results are not relevant. Instead of building a language model from retrieved results, we propose Generative Relevance Feedback (GRF) that builds probabilistic feedback models from long-form text generated from Large Language Models. We study the effective methods for generating text by varying the zero-shot generation subtasks: queries, entities, facts, news articles, documents, and essays. We evaluate GRF on document retrieval benchmarks covering a diverse set of queries and document collections, and the results show that GRF methods significantly outperform previous PRF methods. Specifically, we improve MAP between 5-19% and NDCG@10 17-24% compared to RM3 expansion, and achieve the best R@1k effectiveness on all datasets compared to state-of-the-art sparse, dense, and expansion models.

## CCS CONCEPTS

• Information systems → Information retrieval.

## KEYWORDS

Pseudo-Relevance Feedback; Text Generation; Document Retrieval

### ACM Reference Format:

Iain Mackie, Shubham Chatterjee, and Jeffrey Dalton. 2023. Generative Relevance Feedback with Large Language Models. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '23)*, July 23–27, 2023, Taipei, Taiwan. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3539618.3591992>

## 1 INTRODUCTION

Recent advances in Large Language Models (LLMs) such as GPT-3 [4], PaLM [6], and ChatGPT demonstrate new capabilities to generate long-form fluent text. In addition, LLMs are being combined with search engines, including BingGPT or Bard, to create summaries of search results in interactive forms. In this work, we use these models not to generate end-user responses but as input to the core retrieval algorithm.

The classical approach to address the vocabulary mismatch problem [2] is query expansion using Pseudo-Relevance Feedback (PRF) [1, 30, 31, 51], where the query is expanded using terms from the top- $k$  documents in a feedback set. This feedback set is obtained using a first-pass retrieval, and the expanded query is then used for a second-pass retrieval. While query expansion with PRF often improves recall, its effectiveness hinges on the quality of the first-pass retrieval. Non-relevant results in the feedback set introduce noise and may pull the query off-topic.

To address this problem, we propose Generative Relevance Feedback (GRF) that uses LLMs to generate text independent of first-pass retrieval. Figure 1 shows how we use an LLM to generate diverse

types of query-specific text, before using these “generated documents” as input for proven query expansion models [1]. We experiment using the following types of generated text: keywords, entities, chain-of-thought reasoning, facts, news articles, documents, and essays. Furthermore, we find that combining text across all generation subtasks results in 2-7% higher MAP versus the best standalone generation.

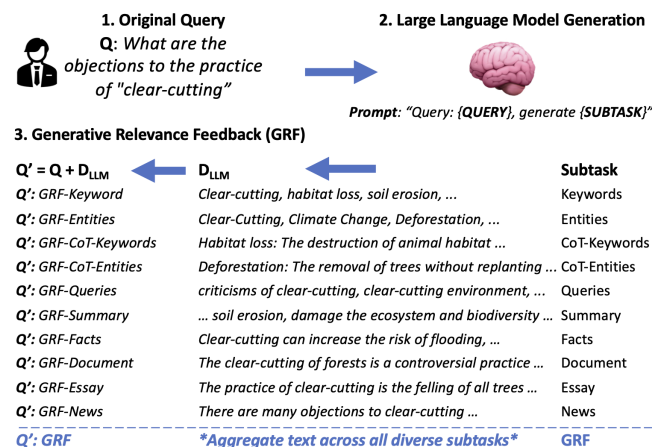


Figure 1: GRF uses diverse LLM-generate text content for relevance feedback to contextualise the query.

We evaluate GRF<sup>1</sup> on four established document ranking benchmarks (see Section 4.1) and outperform several state-of-the-art sparse [1, 37], dense [19, 32, 41], and learned sparse [17] PRF models. We find that long-form text generation (i.e. news articles, documents, and essays) is 7-14% more effective as a feedback set compared to shorter texts (i.e. entities and keywords). Furthermore, the closer the generation subtask is to the style of the target dataset (i.e. news generation for newswire corpus or document generation for web document corpus), the more effective GRF is. Lastly, combining text across all generation subtasks results in 2-7% improvement in MAP over the best standalone generation subtask.

The contributions of this work are:

- We propose GRF, a generative relevance feedback approach which builds a relevance model using text generated from an LLM.
- We show LLM generated long-form text in the style of the target dataset is the most effective. Furthermore, combining text across multiple generation subtasks can further improve effectiveness.
- We demonstrate that GRF improves MAP between 5-19% and NDCG@10 between 17-24% compared to RM3 expansion, and achieves the best Recall@1000 compared to state-of-the-art sparse, dense and learned sparse PRF retrieval models.

<sup>1</sup>Prompts and generated data for reproducibility: [link](#)

## 2 RELATED WORK

**Query Expansion:** Lexical mismatch is a crucial issue in information retrieval, whereby a user query fails to capture their complete information need [2]. Query expansion methods [38] tackle this problem by incorporating terms closer to the user’s intended meaning. One popular technique for automatic query expansion is pseudo-relevance feedback (PRF), where the top  $k$  documents from the initial retrieval are assumed to be relevant. For example, Rocchio [38], KL expansion [51], relevance modelling [30], LCE [31], and RM3 expansion [1]. Additionally, we have seen approaches that expand queries with KG-based information [9, 29, 45, 47] or utilize query-focused LLM vectors for query expansion [32].

Recent advancements in dense retrieval [16, 21, 46] have led to the development of vector-based PRF models [19], such as ColBERT PRF [41], ColBERT-TCT PRF [49], and ANCE-PRF [49]. Furthermore, SPLADE [11] is a neural retrieval model that uses BERT and sparse regularization to learn query and document sparse expansions. Recent work has leveraged query expansion with PRF of learned sparse representations [17]. Unlike prior work, GRF does not rely on pseudo-relevance feedback, instead generating relevant text context for query expansion using LLMs.

**LLM Query Augmentation** The emergence of LLMs has shown progress across many different aspects of information retrieval [48]. This includes using LLMs to change the query representation, such as query generation and rewriting [15, 24, 34, 39, 44, 50], context generation [12, 23], and query-specific reasoning [10, 36]. For example, Nogueira et al. [34] fine-tune a T5 model to generate queries for document expansion for passage retrieval. More recent work by Bonifacio et al. [3] shows that GPT3 can be effectively leveraged for few-shot query generation for dataset generation. Furthermore, LLMs have been used for conversational query re-writing [44] and generating clarifying questions [50].

We have also seen facet generation using T5 [24] and GPT3 [39] to improve the relevance or diversity of search results. While in QA, Liu et al. [23] sample various contextual clues from LLMs and augment and fuse multiple queries. For passage ranking, HyDe [12] uses InstructGPT [35] to generate hypothetical document embeddings and use Contriever [14] for dense retrieval. Lastly, works have shown LLM generation used for query-specific reasoning [10, 36] to improve ranking effectiveness. Our approach differs from prior LLM augmentation approaches as we use LLMs to generate long-form text to produce a probabilistic expansion model to tackle query-document lexical mismatch.

## 3 GENERATIVE RELEVANCE FEEDBACK

Generative Relevance Feedback (GRF) tackles query-document lexical mismatch using text generation for zero-shot query expansion. Unlike traditional PRF approaches for query expansion [1, 30, 31], GRF is not reliant on first-pass retrieval effectiveness to find useful terms for expansion. Instead, we leverage LLMs [5] to generate zero-shot relevant text content.

We build upon prior work on Relevance Models [1] to incorporate the probability distribution of the terms generated by our LLM. This approach enriches the original query with useful terms from diverse generation subtasks, including keywords, entities, chain-of-thought reasoning, facts, news articles, documents, and essays.

We find that the most effective query expansions are: (1) long-form text generations and (2) text content closer in style to the target dataset. In essence, we show that LLMs can effectively generate zero-shot text context close to the target relevant documents.

Lastly, we propose our full GRF method that combines text content across all generation subtasks. The intuition behind this approach is that if terms are used consistently generated across subtasks (i.e. within the entity, fact, and news generations), then these terms are likely useful for expansion. Additionally, multiple diverse subtasks also help expose tail knowledge or uncommon synonyms helpful for retrieval. We find this approach is more effective than any standalone generation subtasks.

### 3.1 GRF Query Expansion

For a given query  $Q$ , Equation 1 shows how  $P_{GRF}(w|R)$  is the probability of a term,  $w$ , being in a relevant document,  $R$ . Similar to RM3 [1], GRF expansion combines the probability of a term given the original query  $P(w|Q)$  with the probability of a term within our LLM-generated document,  $P(w|D_{LLM})$ , which we assume is relevant.  $\beta$  (original query weight) is a hyperparameter to weigh the relative importance of our generative expansion terms. Additionally,  $\theta$  (number of expansion terms) is a hyperparameter with  $W_\theta$  being the set of most probable LLM-generated terms.

$$P_{GRF}(w|R) = \beta P(w|Q) + \begin{cases} (1 - \beta)P(w|D_{LLM}), & \text{if } w \in W_\theta. \\ 0, & \text{otherwise.} \end{cases} \quad (1)$$

### 3.2 Generation Subtasks

We study how LLMs can generate relevant text,  $D_{LLM}$ , across diverse generation subtasks for GRF expansion. The 10 query-specific generation subtasks are:

- **Keywords (64 tokens):** Generates a list of the important words or phrases for the topic, similar to facet generation [24, 39].
- **Entities (64 tokens):** Generates a list of important concepts or named entities, similar to KG-based expansion approaches [9].
- **CoT-Keywords (256 tokens):** Generate chain-of-thought (CoT) [43] reasoning to explain “why” a list of keywords are relevant.
- **CoT-Entities (256 tokens):** Generate CoT reasoning to explain “why” a list of entities are relevant.
- **Queries (256 tokens):** Generate a list of queries based on the original query, similar to [3].
- **Summary (256 tokens):** Generate a concise summary (or answer) to satisfy the query.
- **Facts:** Generate a knowledge-intensive list of text-based facts on the topic, which is close to [23].
- **Document (512 tokens):** Generate a relevant document based on the query closest to a long-form web document.
- **Essay (512 tokens):** Generate a long-form essay-style response.
- **News (512 tokens):** Generate text in the style of a news article.

The full GRF expansion model concatenates text generated across all subtasks to produce  $D_{LLM}$ . We then calculate  $P(w|D_{LLM})$  using this aggregated text, as outlined above. Section 5 shows that the combination using the text across all types is most effective.

## 4 EXPERIMENTAL SETUP

### 4.1 Datasets

**4.1.1 Retrieval Corpora.** **TREC Robust04** [40] was created to investigate methods targeting poorly performing topics. This dataset comprises 249 topics, containing short keyword “titles” and longer natural-language “descriptions” queries. Relevance judgments are over a newswire collection of 528k long documents (TREC Disks 4 and 5), i.e. FT, Congressional Record, LA Times, etc.

**CODEC** [28] is a dataset that focuses on the complex information needs of social science researchers. Domain experts (economists, historians, and politicians) generate 42 challenging essay-style topics. CODEC has a focused web corpus of 750k long documents, which includes news (BBC, Reuters, CNBC etc.) and essay-based web content (Brookings, Forbes, eHistory, etc.).

**TREC Deep Learning (DL) 19/20** [7, 8] builds upon the MS MARCO web queries and documents [33]. The TREC DL dataset uses NIST annotators to provide judgments pooled to a greater depth, containing 43 topics for DL-19 and 45 topics for DL-20. Both query sets are predominately factoid-based [27].

**4.1.2 Indexing and Evaluation.** For indexing we use Pyserini version 0.16.0 [20], removing stopwords and using Porter stemming. We use cross-validation and optimise R@1k on standard folds for Robust04 [13] and CODEC [28]. On DL-19, we cross-validated on DL-20 and use the average parameters zero-shot on DL-19 (and vice versa for DL-20). We assess the system runs to a run depth of 1,000. With GRF being an initial retrieval model, recall-oriented evaluation is important, such as Recall@1000 and MAP to identify relevant documents. We also analyse NDCG@10 to show precision in the top ranks. We use ir-measures for all our evaluations [25] and a 95% confidence paired-t-test for significance.

### 4.2 GRF Implementation

**LLM Generation.** For our text generation we use the GPT3 API [5]. Specifically, we use the `text-davinci-002` model with parameters: *temperature* of 0.7, *top\_p* of 1.0, *frequency\_penalty* of 0.0, and *presence\_penalty* of 0.0. We release all code, generation subtask prompts, generated text content and runs for reproducibility.

**Retrieval and Expansion** To avoid query drift, all GRF runs in the paper use a tuned BM25 system for the input initial run [37]. We tune GRF hyperparameters: the number of feedback terms ( $\theta$ ) and the interpolation between the original terms and generative expansion terms ( $\beta$ ). The tuning methodology is the same as BM25 and BM25 with RM3 expansion to make the GRF directly comparable; see below for details.

### 4.3 Comparison Methods

**BM25** [37]: Sparse retrieval method, we tune  $k1$  parameter (0.1 to 5.0 with a step size of 0.2) and  $b$  (0.1 to 1.0 with a step size of 0.1).

**BM25+RM3** [1]: For BM25 with RM3 expansion, we tune  $fb\_terms$  (5 to 95 with a step of 5),  $fb\_docs$  (5 to 50 with a step of 5), and  $original\_query\_weight$  (0.2 to 0.8 with a step of 0.1).

**CEQE** [32]: Utilizes query-focused vectors for query expansion. We use the CEQE-MaxPool runs provided by the author.

**SPLADE+RM3**: We use RM3 [1] expansion with SPLADE [11]. We use `naver/splade-cocondenser-ensembledistil` checkpoint and Pyserini’s [20] “impact” searcher for max-passage aggregation. We tune  $fb\_docs$  (5,10,15,20,25,30),  $fb\_terms$  (20,40,60,80,100), and  $original\_query\_weight$  (0.1 to 0.9 with a step of 0.1).

**TCT+PRF**: [18] is a Roccio PRF approach using ColBERT-TCT [22]. We employ a max-passage approach with TCT-ColBERT-v2-HNP checkpoint. We tune Roccio PRF parameters:  $depth$  (2,3,5,7,10,17),  $\alpha$  (0.1 to 0.9 with a step of 0.1), and  $\beta$  (0.1 to 0.9 with a step of 0.1).

**ColBERT+PRF** [41]: We use the runs provided by Wang et al. [42], which use pyterrier framework [26] for ColBERT-PRF retrieval.

## 5 RESULTS & ANALYSIS

### 5.1 RQ1: What generative content is most effective for query expansion?

Table 1 shows the effectiveness of generative feedback with varying units of text (Keywords-News) and our full hybrid method that uses text from all subtasks. We test for significant improvements against BM25 with RM3 expansion, to ascertain whether our zero-shot generative feedback methods improve over RM3 expansion.

Generation subtasks that target short text span or lists (Keywords, Entities, Keywords-COT, Entities-COT, and Queries) do not

**Table 1: GRF with different generation subtasks. Significant improvements against BM25+RM3 (“+”) and best system (**bold**).**

	Robust04 -Title			CODEC			DL-19			DL-20		
	NDCG@10	MAP	R@1k	NDCG@10	MAP	R@1k	NDCG@10	MAP	R@1k	NDCG@10	MAP	R@1k
BM25	0.445	0.252	0.705	0.316	0.214	0.783	0.531	0.335	0.703	0.546	0.413	0.811
BM25+RM3	0.451	0.292	0.777	0.326	0.239	0.816	0.541	0.383	0.745	0.513	0.418	0.825
GRF-Keywords	0.435	0.252	0.717	0.327	0.218	0.748	0.565	0.377	0.749	0.554	0.435	0.822
GRF-Entities	0.452	0.252	0.698	0.341	0.216	0.750	0.531	0.363	0.741	0.544	0.414	0.824
GRF-CoT-Keywords	0.436	0.248	0.704	0.327	0.239	0.774	0.550	0.382	0.748	0.542	0.423	0.817
GRF-CoT-Entities	0.450	0.252	0.714	0.355	0.243	0.789	0.563	0.389	0.757	0.552	0.430	0.832
GRF-Queries	0.450	0.257	0.710	0.347	0.233	0.773	0.551	0.367	0.760	0.568	0.439	0.851
GRF-Summary	0.491 <sup>+</sup>	0.277	0.730	0.398 <sup>+</sup>	0.260	0.796	0.577	0.414	0.761	0.585 <sup>+</sup>	0.472 <sup>+</sup>	0.865
GRF-Facts	0.501 <sup>+</sup>	0.284	0.744	0.353	0.255	0.795	0.569	0.401	0.769	0.583 <sup>+</sup>	0.459 <sup>+</sup>	0.871
GRF-Document	0.480 <sup>+</sup>	0.276	0.728	0.376 <sup>+</sup>	0.265	0.795	0.618 <sup>+</sup>	0.428 <sup>+</sup>	0.787 <sup>+</sup>	0.589 <sup>+</sup>	0.476 <sup>+</sup>	0.872
GRF-Essay	0.494 <sup>+</sup>	0.284	0.736	<b>0.405<sup>+</sup></b>	0.270 <sup>+</sup>	0.803	0.609 <sup>+</sup>	0.421 <sup>+</sup>	0.779 <sup>+</sup>	0.551	0.440	0.859
GRF-News	0.501 <sup>+</sup>	0.287	0.745	0.398 <sup>+</sup>	0.270 <sup>+</sup>	0.828	0.609	0.409	0.777	0.578 <sup>+</sup>	0.457	0.853
GRF	<b>0.528<sup>+</sup></b>	<b>0.307</b>	<b>0.788</b>	<b>0.405<sup>+</sup></b>	<b>0.285<sup>+</sup></b>	<b>0.830</b>	<b>0.620<sup>+</sup></b>	<b>0.441<sup>+</sup></b>	<b>0.797<sup>+</sup></b>	<b>0.607<sup>+</sup></b>	<b>0.486<sup>+</sup></b>	<b>0.879<sup>+</sup></b>

**Table 2: GRF against state-of-the-art PRF models. Significant improvements against BM25+RM3 (“+”) and best system (*bold*).**

	Robust04 -Title			CODEC			DL-19			DL-20		
	nDCG@10	MAP	R@1k	nDCG@10	MAP	R@1k	nDCG@10	MAP	R@1k	nDCG@10	MAP	R@1k
BM25+RM3	0.451	0.292	0.777	0.326	0.239	0.816	0.541	0.383	0.745	0.513	0.418	0.825
CEQE-MaxPool	0.474	<b>0.310+</b>	0.764	-	-	-	0.518	0.378	0.746	0.473	0.396	0.841
SPLADE+RM3	0.418	0.248	0.703	0.311	0.216	0.770	0.566	0.328	0.651	0.533	0.379	0.784
TCT+PRF	0.493	0.274	0.684	0.358	0.239	0.757	<b>0.670+</b>	0.378	0.684	<b>0.618+</b>	0.442	0.784
ColBERT-PRF	0.467	0.272	0.648	-	-	-	0.668+	0.385	0.625	0.615+	<b>0.489+</b>	0.813
GRF (Ours)	<b>0.528+</b>	0.307	<b>0.788</b>	<b>0.405+</b>	<b>0.285+</b>	<b>0.830</b>	0.620+	<b>0.441+</b>	<b>0.797+</b>	0.607+	0.486+	<b>0.879+</b>

offer significantly improves over RM3 expansion. Conversely, sub-tasks targeting long text generation (Summary, Facts, Document, Essay, News) significantly improve at least two datasets over RM3 expansions. This indicates that more terms generated from the LLM provide a better relevance model, and increases MAP between 7-14% when we compare these two categories.

Furthermore, we find most effective generation subtasks are aligned with the style of the target dataset. For example, Facts and News are the best standalone generation methods across all measures on Robust04, where the dataset contains fact-heavy topics and a newswire corpus. Additionally, Essay and News are the best generation subtasks on CODEC across all measures, which aligns with its essay-style queries over a news (BBC, Reuters, CNBC, etc.) and essay-style (Brookings, Forbes, eHistory, etc.) corpus. Lastly, Document is the best generation subtask across DL-19 and DL-20, aligning with MS Marcos web document collection. Overall, this finding supports that LLM generative content in the styles of the target dataset is the most effective.

Although we see significant improvements from some standalone generation subtasks, particularly NDCG@10 (15/40 subtasks across the datasets), the full GRF method is consistently as good if not better than any standalone subtask. Specifically, GRF improves NDCG by 0.0-5.4%, MAP by 2.1-7.0% and R@1k by 0.2-5.8% across the datasets. This shows that combining LLM-generated text from various generation subtasks is a robust and effective method of relevance modelling.

Lastly, these results show that GRF expansion from generated text is consistently better, often significantly, than RM3 expansion that uses documents from the target corpus. Specifically, we find significant improvement on all measures across DL-19 and DL-20, NDCG@10 and MAP on CODEC, and NDCG@10 on Robust04 titles. Although not included for space limitations, on Robust04 description queries, GRF shows significant improvements with an NDCG of 0.550, MAP of 0.318, and R@1k of 0.776.

These results strongly support that LLM generation is an effective query expansion method without relying on first-pass retrieval effectiveness. For example, we look at the hardest 20% of Robust04 topics ordered by NDCG@10; we find that RM3 offers minimal uplift and only improves NDCG@10 by +0.006, MAP by +0.008, and R@1k by +0.052. In contrast, GRF is not reliant on first-pass retrieval effectiveness, and GRF improves NDCG@10 by +0.145, MAP by +0.068, and R@1k by +0.165 (a relative improvement of +100-200% on NDCG@10 and MAP).

## 5.2 RQ2: How does GRF compare to state-of-the-art PRF models?

Table 2 shows GRF against state-of-the-art sparse, dense, and learned sparse PRF models across target datasets. This allows us to directly compare GRF’s unsupervised term-based queries against PRF methods that use more complex LLM-based embeddings. We conduct significance testing against BM25 with RM3 expansion.

GRF has the best R@1k across all datasets and has comparable and often better effectiveness in the top ranks. Specifically, on more challenging datasets, such as CODEC and Robust04 titles, GRF is the best system across all measures, except Robust04 titles MAP, which is 0.003 less than CEQE. Although not included for space limitations, GRF is also the most effective system on Robust04 descriptions across all measures. GRF vastly outperforms dense retrieval and dense PRF on these challenging datasets, with a performance gap of 7-14% on NDCG@20, 13-21% MAP, and 10-22% R@1k.

Dense retrieval has been shown to be highly effective on the more factoid-focused datasets, such as DL-19 and DL-20. However, as well as the best R@1k, our unsupervised GRF queries have comparable NDCG@10 and MAP scores to dense PRF models. This is juxtaposed to other sparse methods (BM25 and BM25 with RM3 expansions) or LLM expansion (CEQE), which have much poorer precision in the top ranks. Overall, this supports that generative expansion is a highly effective initial retrieval method across various collections and query types.

## 6 CONCLUSION

To our knowledge, this is the first work to study the use of long-form text generated from large-language models for query expansion. We show that generating long-form text in news-like and essay-like formats is effective input for probabilistic query expansion approaches. The results on document retrieval on multiple corpora show that the proposed GRF approach outperforms models that use retrieved documents (PRF). The results show GRF improves MAP between 5-19% and NDCG@10 between 17-24% when compared to RM3 expansion, and achieves the best Recall@1000 compared to state-of-the-art PRF retrieval models. We envision GRF as one of the many new emerging methods that use LLM-generated content to improve the effectiveness of core retrieval tasks.

## 7 ACKNOWLEDGEMENTS

This work is supported by the 2019 Bloomberg Data Science Research Grant and the Engineering and Physical Sciences Research Council grant EP/V025708/1.

## REFERENCES

- [1] Nasreen Abdul-Jaleel, James Allan, W Bruce Croft, Fernando Diaz, Leah Larkey, Xiaoyan Li, Mark D Smucker, and Courtney Wade. 2004. UMass at TREC 2004: Novelty and HARD. *Computer Science Department Faculty Publication Series* (2004), 189.
- [2] Nicholas J Belkin, Robert N Oddy, and Helen M Brooks. 1982. ASK for information retrieval: Part I. Background and theory. *Journal of documentation* (1982).
- [3] Luiz Bonifacio, Hugo Abonizio, Marzieh Fadaee, and Rodrigo Nogueira. 2022. Inpars: Unsupervised dataset generation for information retrieval. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 2387–2392.
- [4] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language Models are Few-Shot Learners. In *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (Eds.), Vol. 33. Curran Associates, Inc., 1877–1901. [https://proceedings.neurips.cc/paper\\_files/paper/2020/file/1457c0d6bfc4967418bfb8ac142f64a-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2020/file/1457c0d6bfc4967418bfb8ac142f64a-Paper.pdf)
- [5] Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165* (2020).
- [6] Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, Parker Schuh, Kensen Shi, Sasha Tsvyashchenko, Joshua Maynez, Abhishek Rao, Parker Barnes, Yi Tay, Noam Shazeer, Vinodkumar Prabhakaran, Emily Reif, Nan Du, Ben Hutchinson, Reiner Pope, James Bradbury, Jacob Austin, Michael Isard, Guy Gur-Ari, Pengcheng Yin, Toju Duke, Anselm Levskaya, Sanjay Ghemawat, Sunipa Dev, Henryk Michalewski, Xavier Garcia, Vedant Misra, Kevin Robinson, Liam Fedus, Denny Zhou, Daphne Ippolito, David Luan, Hyeontaek Lim, Barret Zoph, Alexander Spiridonov, Ryan Sepassi, David Dohan, Shivani Agrawal, Mark Omernick, Andrew M. Dai, Thanumalayan Sankaranarayanan Pillai, Marie Pellat, Aitor Lewkowycz, Erica Moreira, Rewon Child, Oleksandr Polozov, Katherine Lee, Zongwei Zhou, Xuezhi Wang, Brennan Saeta, Mark Diaz, Orhan Firat, Michele Catasta, Jason Wei, Kathy Meier-Hellstern, Douglas Eck, Jeff Dean, Slav Petrov, and Noah Fiedel. 2022. PaLM: Scaling Language Modeling with Pathways. *arXiv:cs.CL/2204.02311*
- [7] Nick Craswell, Bhaskar Mitra, Emine Yilmaz, and Daniel Campos. 2021. Overview of the TREC 2020 deep learning track. In *Text REtrieval Conference (TREC)*. TREC.
- [8] Nick Craswell, Bhaskar Mitra, Emine Yilmaz, Daniel Campos, and Ellen M Voorhees. 2020. Overview of the trec 2019 deep learning track. *arXiv preprint arXiv:2003.07820* (2020).
- [9] Jeffrey Dalton, Laura Dietz, and James Allan. 2014. Entity query feature expansion using knowledge base links. In *Proceedings of the 37th international ACM SIGIR conference on Research & development in information retrieval*. 365–374.
- [10] Fernando Ferraretto, Thiago Laitz, Roberto Lotufo, and Rodrigo Nogueira. 2023. ExaRanker: Explanation-Augmented Neural Ranker. *arXiv preprint arXiv:2301.10521* (2023).
- [11] Thibault Formal, Benjamin Piwowarski, and Stéphane Clinchant. 2021. SPLADE: Sparse lexical and expansion model for first stage ranking. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 2288–2292.
- [12] Luyu Gao, Xueguang Ma, Jimmy Lin, and Jamie Callan. 2022. Precise Zero-Shot Dense Retrieval without Relevance Labels. *arXiv preprint arXiv:2212.10496* (2022).
- [13] Samuel Huston and W Bruce Croft. 2014. Parameters learned in the comparison of retrieval models using term dependencies. *Ir, University of Massachusetts* (2014).
- [14] Gautier Izacard, Mathilde Caron, Lucas Hosseini, Sebastian Riedel, Piotr Bojanowski, Armand Joulin, and Edouard Grave. 2021. Unsupervised Dense Information Retrieval with Contrastive Learning. <https://doi.org/10.48550/ARXIV.2112.09118>
- [15] Vitor Jeronymo, Luiz Bonifacio, Hugo Abonizio, Marzieh Fadaee, Roberto Lotufo, Jakub Zavrel, and Rodrigo Nogueira. 2023. InPars-v2: Large Language Models as Efficient Dataset Generators for Information Retrieval. *arXiv preprint arXiv:2301.01820* (2023).
- [16] Omar Khattab and Matei Zaharia. 2020. ColBERT: Efficient and Effective Passage Search via Contextualized Late Interaction over BERT. *arXiv preprint arXiv:2004.12832* (2020).
- [17] Carlos Lassance and Stéphane Clinchant. 2023. Naver Labs Europe (SPLADE)@ TREC Deep Learning 2022. *arXiv preprint arXiv:2302.12574* (2023).
- [18] Hang Li, Ahmed Mourad, Shengyao Zhuang, Bevan Koopman, and G. Zuccon. 2021. Pseudo Relevance Feedback with Deep Language Models and Dense Retrievers: Successes and Pitfalls. *ArXiv abs/2108.11044* (2021).
- [19] Hang Li, Shengyao Zhuang, Ahmed Mourad, Xueguang Ma, Jimmy Lin, and Guido Zuccon. 2022. Improving Query Representations for Dense Retrieval with Pseudo Relevance Feedback: A Reproducibility Study. In *European Conference on Information Retrieval*. Springer, 599–612.
- [20] Jimmy Lin, Xueguang Ma, Sheng-Chieh Lin, Jheng-Hong Yang, Ronak Pradeep, and Rodrigo Nogueira. 2021. Pyserini: A Python toolkit for reproducible information retrieval research with sparse and dense representations. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 2356–2362.
- [21] Sheng-Chieh Lin, Jheng-Hong Yang, and Jimmy Lin. 2020. Distilling dense representations for ranking using tightly-coupled teachers. *arXiv preprint arXiv:2010.11386* (2020).
- [22] Sheng-Chieh Lin, Jheng-Hong Yang, and Jimmy Lin. 2021. In-batch negatives for knowledge distillation with tightly-coupled teachers for dense retrieval. In *Proceedings of the 6th Workshop on Representation Learning for NLP (ReplANLP-2021)*. 163–173.
- [23] Linqing Liu, Minghan Li, Jimmy Lin, Sebastian Riedel, and Pontus Stenetorp. 2022. Query Expansion Using Contextual Clue Sampling with Language Models. *arXiv preprint arXiv:2210.07093* (2022).
- [24] Sean MacAvaney, Craig Macdonald, Roderick Murray-Smith, and Iadh Ounis. 2021. IntenT5: Search Result Diversification using Causal Language Models. *arXiv preprint arXiv:2108.04026* (2021).
- [25] Sean MacAvaney, Craig Macdonald, and Iadh Ounis. 2022. Streamlining Evaluation with ir-measures. In *European Conference on Information Retrieval*. Springer, 305–310.
- [26] Craig Macdonald, Nicola Tonello, Sean MacAvaney, and Iadh Ounis. 2021. PyTerrier: Declarative experimentation in Python from BM25 to dense retrieval. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*. 4526–4533.
- [27] Iain Mackie, Jeffrey Dalton, and Andrew Yates. 2021. How deep is your learning: The DL-HARD annotated deep learning dataset. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 2335–2341.
- [28] Iain Mackie, Paul Owoicho, Carlos Gemmell, Sophie Fischer, Sean MacAvaney, and Jeffery Dalton. 2022. CODEC: Complex Document and Entity Collection. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*.
- [29] Edgar Meij, Dolf Trieschnigg, Maarten De Rijke, and Wessel Kraaij. 2010. Conceptual language models for domain-specific retrieval. *Information Processing & Management* 46, 4 (2010), 448–469.
- [30] Donald Metzler and W Bruce Croft. 2005. A markov random field model for term dependencies. In *Proceedings of the 28th annual international ACM SIGIR conference on Research and development in information retrieval*. 472–479.
- [31] Donald Metzler and W Bruce Croft. 2007. Latent concept expansion using markov random fields. In *Proceedings of the 30th annual international ACM SIGIR conference on Research and development in information retrieval*. 311–318.
- [32] Shahrazad Naseri, Jeffrey Dalton, Andrew Yates, and James Allan. 2021. Ceqe: Contextualized embeddings for query expansion. In *Advances in Information Retrieval: 43rd European Conference on IR Research, ECIR 2021, Virtual Event, March 28–April 1, 2021, Proceedings, Part I* 43. Springer, 467–482.
- [33] Tri Nguyen, Mir Rosenberg, Xia Song, Jianfeng Gao, Saurabh Tiwary, Rangan Majumder, and Li Deng. 2016. Ms marco: A human-generated machine reading comprehension dataset. (2016).
- [34] Rodrigo Nogueira, Jimmy Lin, and AI Epistemic. 2019. From doc2query to docTTTTTquery. *Online preprint* 6 (2019).
- [35] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems* 35 (2022), 27730–27744.
- [36] Jayr Pereira, Robson Fidalgo, Roberto Lotufo, and Rodrigo Nogueira. 2023. Visconde: Multi-document QA with GPT-3 and Neural Reranking. In *Advances in Information Retrieval: 45th European Conference on Information Retrieval, ECIR 2023, Dublin, Ireland, April 2–6, 2023, Proceedings, Part II*. Springer, 534–543.
- [37] Stephen E Robertson and Steve Walker. 1994. Some simple effective approximations to the 2-poisson model for probabilistic weighted retrieval. In *SIGIR'94*. Springer, 232–241.
- [38] Joseph Rocchio. 1971. Relevance feedback in information retrieval. *The Smart retrieval system-experiments in automatic document processing* (1971), 313–323.
- [39] Chris Samarin, Arkin Dharawat, and Hamed Zamani. 2022. Revisiting Open Domain Query Facet Extraction and Generation. In *Proceedings of the 2022 ACM SIGIR International Conference on Theory of Information Retrieval*. 43–50.
- [40] Ellen M. Voorhees. 2004. Overview of the TREC 2004 Robust Track. In *Proceedings of the Thirteenth Text REtrieval Conference (TREC 2004)*. Gaithersburg, Maryland, 52–69.
- [41] Xiao Wang, Craig Macdonald, Nicola Tonello, and Iadh Ounis. 2022. ColBERT-PRF: Semantic Pseudo-Relevance Feedback for Dense Passage and Document Retrieval. *ACM Transactions on the Web* (2022).
- [42] Xiao Wang, Craig Macdonald, Nicola Tonello, and Iadh Ounis. 2023. ColBERT-PRF: Semantic Pseudo-Relevance Feedback for Dense Passage and Document

- Retrieval. *ACM Transactions on the Web* 17, 1 (2023), 1–39.
- [43] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed H Chi, Quoc V Le, Denny Zhou, et al. [n. d.]. Chain-of-Thought Prompting Elicits Reasoning in Large Language Models. In *Advances in Neural Information Processing Systems*.
- [44] Zeqiu Wu, Yi Luan, Hannah Rashkin, David Reitter, Hannaneh Hajishirzi, Mari Ostendorf, and Gaurav Singh Tomar. 2022. CONQRR: Conversational Query Rewriting for Retrieval with Reinforcement Learning. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Abu Dhabi, United Arab Emirates, 10000–10014. <https://aclanthology.org/2022.emnlp-main.679>
- [45] Chenyan Xiong and Jamie Callan. 2015. Query Expansion with Freebase. In *Proceedings of the 2015 International Conference on The Theory of Information Retrieval (ICTIR '15)*. Association for Computing Machinery, New York, NY, USA, 111–120. <https://doi.org/10.1145/2808194.2809446>
- [46] Lee Xiong, Chenyan Xiong, Ye Li, Kwok-Fung Tang, Jialin Liu, Paul N Bennett, Junaid Ahmed, and Arnold Overwijk. [n. d.]. Approximate Nearest Neighbor Negative Contrastive Learning for Dense Text Retrieval. In *International Conference on Learning Representations*.
- [47] Yang Xu, Gareth J.F. Jones, and Bin Wang. 2009. Query Dependent Pseudo-Relevance Feedback Based on Wikipedia. In *Proceedings of the 32nd International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '09)*. Association for Computing Machinery, New York, NY, USA, 59–66. <https://doi.org/10.1145/1571941.1571954>
- [48] Andrew Yates, Rodrigo Nogueira, and Jimmy Lin. 2021. Pretrained Transformers for Text Ranking: BERT and Beyond. In *WSDM*. 1154–1156.
- [49] HongChien Yu, Chenyan Xiong, and Jamie Callan. 2021. Improving query representations for dense retrieval with pseudo relevance feedback. *arXiv preprint arXiv:2108.13454* (2021).
- [50] Hamed Zamani, Susan Dumais, Nick Craswell, Paul Bennett, and Gord Lueck. 2020. Generating clarifying questions for information retrieval. In *Proceedings of the web conference 2020*. 418–428.
- [51] Chengxiang Zhai and John Lafferty. 2001. Model-based feedback in the language modeling approach to information retrieval. In *Proceedings of the tenth international conference on Information and knowledge management*. 403–410.