

Synthetic Test Collections for Retrieval Evaluation

Hossein A. Rahmani University College London London, UK hossein.rahmani.22@ucl.ac.uk Nick Craswell Microsoft Bellevue, US nickcr@microsoft.com Emine Yilmaz University College London & Amazon London, UK emine.yilmaz@ucl.ac.uk

Bhaskar Mitra Microsoft Montréal, Canada bmitra@microsoft.com Daniel Campos Snowflake New York, US daniel.campos@snowflake.com

ABSTRACT

Test collections play a vital role in evaluation of information retrieval (IR) systems. Obtaining a diverse set of user queries for test collection construction can be challenging, and acquiring relevance judgments, which indicate the appropriateness of retrieved documents to a query, is often costly and resource-intensive. Generating synthetic datasets using Large Language Models (LLMs) has recently gained significant attention in various applications. In IR, while previous work exploited the capabilities of LLMs to generate synthetic queries or documents to augment training data and improve the performance of ranking models, using LLMs for constructing synthetic test collections is relatively unexplored. Previous studies demonstrate that LLMs have the potential to generate synthetic relevance judgments for use in the evaluation of IR systems. In this paper, we comprehensively investigate whether it is possible to use LLMs to construct fully synthetic test collections by generating not only synthetic judgments but also synthetic queries. In particular, we analyse whether it is possible to construct reliable synthetic test collections and the potential risks of bias such test collections may exhibit towards LLM-based models. Our experiments indicate that using LLMs it is possible to construct synthetic test collections that can reliably be used for retrieval evaluation.

 $\label{eq:constraint} \ref{eq:constraint} https://github.com/rahmanidashti/SyntheticTestCollections$

CCS CONCEPTS

• Information systems \rightarrow Information retrieval.

KEYWORDS

Synthetic Data Generation, Large Language Model, Test Collection

ACM Reference Format:

Hossein A. Rahmani, Nick Craswell, Emine Yilmaz, Bhaskar Mitra, and Daniel Campos. 2024. Synthetic Test Collections for Retrieval Evaluation. 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/10.1145/3626772.3657942



This work is licensed under a Creative Commons Attribution International 4.0 License.

SIGIR ¹24, July 14–18, 2024, Washington, DC, USA © 2024 Copyright held by the owner/author(s). ACM ISBN 979-8-4007-0431-4/24/07 https://doi.org/10.1145/3626772.3657942

1 INTRODUCTION

Test collection construction is a pivotal process to evaluate the effectiveness of information retrieval (IR) systems. The most widely used approach for constructing test collections is based on the Cranfield paradigm [4, 8], which involves creating a collection comprising queries and associated relevance judgments. Queries used in test collection construction are expected to come from real usage logs, representing real information needs. However, it is very difficult to get access to such logs outside of search engine companies. Hence, lots of existing test collections used in IR are based on manually created queries [9, 19]. This process demands time and expertise, making it costly in terms of labor and resources: furthermore, there are no guarantees that queries generated at the end of this process are representative of real information needs. Similarly, obtaining relevance judgments for a query is an expensive procedure requiring significant human effort. The highly demanding nature of the test construction process is a major bottleneck in constructing large test collections; hence, most existing publicly available test collections tend to consist of a small number of queries, which could degrade the reliability of these collections.

Recently, Large Language Models (LLMs) have demonstrated remarkable performance on unseen tasks by only considering the instructions (so-called 'prompts') provided to them [7, 12]. Synthetic datasets generated using LLMs have recently gained attention across a range of diverse tasks [5, 13, 20]. In IR, previous studies use LLMs to generate synthetic training data for augmentation to boost the quality of retrievers [3, 6]. LLMs have also been used to generate relevance labels [11, 18], as well as to generate query variants for evaluation and training of IR systems [2, 16]. However, to the best of our knowledge, no prior study has thoroughly explored the potential of LLMs to construct fully synthetic test collections where both queries and associated relevance judgments are automatically generated using LLMs.

Given the aforementioned challenges in constructing large-scale test collections, our goal in this paper is to investigate whether it is possible to create reliable synthetic test collections so that there is (i) no need for real usage query logs or manual creation of queries, and (ii) no need to obtain manual relevance judgments. We investigate different approaches to construct synthetic test collections using LLMs and show using synthetic test collections it is possible to obtain evaluation results that are similar to results obtained using real test collections.



Figure 1: Scatter plot of the effectiveness (i.e., NDCG@10) of TREC DL 2023 runs according to the real queries and synthetic queries with human judgments. A point represents a single run averaged over all queries.

Table 1: Statistics of queries per query type

	All	Real	T5 Generated	GPT-4 Generated
No. of Queries	82	51	13	18
Avg. Query Length	6.84	5.76	5.69	10.72
Min Query Length	2	2	4	6
Max Query Length	15	14	8	15

2 SYNTHETIC TEST COLLECTION

2.1 Synthetic Query Generation

The reliability of a test collection highly depends on the queries included in the collection. Since obtaining real user queries may not always be possible, we analyzed the feasibility of using synthetic queries in test collection construction, focusing on constructing test collections for passage retrieval task.

Our goal is to create meaningful queries, which can be characterized as a collection of terms that can be used to reach a document and make sense alone without any further context from the document from which the query is derived. The approach we followed is based on starting with a passage and generating a query to which this passage would be relevant.

To generate the synthetic queries, we first sampled 1000 passages at random from the MS MARCO v2 passage corpus¹ as anchor documents. The selected passages are then go through a filtering step whose goal is to get rid of "bad" passages, which can lead to low quality queries being generated. Examples of bad passages include passages that talk about a person without saying the person's name ("she then completed her first novel...") or passages that use terms that might not be clear once the passage is removed from its document context ("we eventually chose the second solution, since it gave a good balance of accuracy and efficiency"). To identify passages that might not be good stand-alone search results, we ran each of the 1000 passages through a GPT-4 prompt, which generates a *query independent passage quality score*. Since this is a rough one-off filter, we did not repeat it in the case of some error or malformed output, which eliminated 8.9% of passages. We Hossein A. Rahmani, Nick Craswell, Emine Yilmaz, Bhaskar Mitra, Daniel Campos

 Table 2: Average number of documents per query for each

 relevance grade for different query types.

Relevance Grade	All	Real	T5 Generated	GPT-4 Generated
Nonrelevant (0)	169.09	159.31	213.30	164.88
Related (1)	53.31	64.15	31.23	38.55
Highly relevant (2)	27.54	31.60	19.46	21.88
Perfectly relevant (3)	22.31	24.35	21.07	17.44

filtered low-quality passages with passage quality scores less than 50, removing 14.6% passages. The remaining passages were seeds for generating synthetic queries.

We then generated queries using two methods: The first method uses a small pre-trained model based on T5 [15], and the second method is based on zero-shot query generation using GPT-4 [1]. For query generation using T5, we applied the BeIR query generation model² that uses a T5-based model pre-trained on the MS MARCO Passage dataset [17]. We then generated one query per passage using both T5 and GPT-4 approaches. The details of the approaches can be found in our GitHub.

Then, we asked experts who are professional assessors with very good experience in relevance annotation to remove the queries that did not look reasonable, that contained too few or too many relevant documents as these queries tend to be noisy or not very informative for evaluation purposes. Amongst 48 T5-generated queries, 13 of them were selected and amongst 49 queries generated using GPT-4, 18 of them were selected to be included in the test collection. Queries generated as a result of this process were included in the test collections created as part of the TREC Deep Learning Track (TREC DL) 2023 [10], which consist of 51 real queries in addition to the synthetic queries. Systems submitted to the DL Track were run on both synthetic and real queries in the test collection. For all query types, depth-10 pooling was used to select the documents to be judged by the expert assessors who labelled the documents based on four relevance grades (irrelevant, related, highly relevant and perfectly relevant).

Table 1 shows the total number of queries included in the TREC DL 2023 test collection for each query type, together with the average, min and max query length for each category. We also included statistics about real queries for a better comparison of synthetic queries with human queries. It can be seen that queries generated using GPT-4 tend to be much longer than the queries generated using T5 and the real queries and in general, queries generated using T5 tend to be shorter than the other query types. For each query type, Table 2 shows the average number of documents per query for each relevance grade. It can be seen that synthetic queries contain much fewer relevant documents (documents with a relevance grade greater than zero) compared to the real queries (120.1 documents of relevance grade greater than zero for real queries vs. 71.76 documents for queries generated by T5 and 77.87 documents for queries generated by GPT-4.) Also, pools constructed using queries generated by T5 tend to contain significantly more non-relevant documents compared to the rest of the query types, suggesting that these queries may be more difficult than the other queries.

¹https://msmarco.blob.core.windows.net/msmarcoranking/msmarco_v2_passage.tar

 $^{^{2}} https://huggingface.co/BeIR/query-gen-msmarco-t5-large-v1$



Figure 2: Scatter plots of the effectiveness (i.e., NDCG@10) of TREC DL 2023 run according to the real queries with human judgments and our synthetic queries with (a) sparse judgments and (b) synthetic judgments.

To compute the reliability of the synthetic queries, we compared the performance of systems that were submitted to TREC DL 2023 on real queries with the performance of systems solely on synthetic queries. For this purpose, we evaluate the quality of the 31 systems submitted to the full ranking task of the track using official judgments of the track obtained from expert human assessors from NIST, and compare the ranking of these systems on real queries and synthetic queries. Figure 1 shows how the performance of systems using synthetically generated queries (both T5 and GPT-4 based) compare with system performance on real queries. The figure also includes the line y = x for easy comparison. It can be seen that synthetic queries and real queries show similar patterns in terms of evaluation results and system ranking, with a system ordering agreement of Kendall's $\tau = 0.8151$.

2.2 Synthetic Relevance Judgment Generation

Once a sample of queries to be included in a test collection has been identified, the second step is to identify documents relevant to these queries, a process that is usually done by collecting manual annotations. While in the previous section we assumed that manual annotations are available for all query types (including synthetic queries), obtaining manual annotations is a very expensive process and it may not always be possible to get these many judgments from expert annotators. Hence, we next investigate the possibility of generating synthetic relevance judgments for test collections constructed using synthetic queries.

As the simplest method, one can assume that the passages that we used to generate the synthetic queries are the only passages that are relevant to those queries, referred to as *sparse judgments*. Figure 2a shows the system ranking and Kendall's τ correlation on real queries with human judgments and synthetic queries with sparse judgments. Compared to synthetic queries with human judgments ($\tau = 0.8151$), evaluation using synthetic queries with sparse judgments does not give similar results to human queries, with a very low system ordering agreement of $\tau = 0.157$.

Recent work [18] showed that it is possible to get very highquality relevance judgements using LLMs. Thus, the question we investigate further is can we generate a fully synthetic test collection by not only generating synthetic queries but also by generating

Table 3: Judgment agreement on TREC DL 2023 between TREC assessors and the LLM (i.e., GPT-4) on real queries and synthetic queries.

	GPT-4	Cahan'a K				
	Prediction	Perfect rel.	High. rel.	Related	Irrelevant	
Real	Perfect. rel.	597	469	496	324	0.24
	High. rel.	322	473	648	501	
	Related	298	548	1358	2736	
	Irrelevant	25	122	770	4564	
Synthetic	Perfect. rel.	166	92	40	150	0.26
	High. rel.	227	188	197	305	
	Related	170	301	695	1871	
	Irrelevant	25	66	168	3415	

synthetic relevance judgements? To answer this question, we used the GPT-4 language model as accessed through Open AI's API in order to automatically label the documents (that were originally annotated using human annotators) for the synthetic queries to generate synthetic relevance judgements. For this purpose, we used the prompt template introduced in the recent study [18] that shows the possibility of getting high-quality relevance judgements using LLMs. In our experiments the temperature was set at zero, so the model would select the single most likely output and other parameters of the model were top p = 1, frequency penalty 0.5, and presence penalty 0.

Table 3 shows how the judgments generated using GPT-4 compare with manual judgments. The table reports the "agreement" on the full 4-point relevance scale on real and synthetic queries, respectively. For both real and synthetic queries, we observe a fair level of agreement between synthetic judgements generated using GPT-4 and manual judgments: The Cohen's κ on real queries is 0.24 and on synthetic queries is 0.26. In earlier studies, Faggioli et al. [11] report a Cohen's κ of 0.26 between TREC DL 2021 assessors and GPT-3.5 with two types of human judgments, and Thomas et al. [18] show that Cohen's κ varies from 0.20 to 0.64 on two-level scale of TREC Robust data based on different versions of prompts.

The generated judgments on synthetic queries show that GPT-4 consistently underestimates 'Perfectly relevant' and 'Highly relevant' labels. On real queries, GPT-4 labels on 'Perfectly relevant' are more likely to be labelled as the same but on synthetic queries, according to GPT-4, a 'Perfectly relevant' document is more likely to be rated as a 'Highly relevant' document. For documents judged as 'Perfectly relevant' by human assessors, GPT-4 generates the same judgment in 28% of the cases. The results indicate that human assessors may use more detailed information to distinguish between 'Perfectly relevant' from 'Highly relevant' documents that are not fully captured by an LLM. We notice a similar pattern when we compare the 'Highly relevant' and 'Related' documents on both real and synthetic queries.

One of the primary uses of test collections is for system evaluation. Even though the synthetic judgments generated could be slightly different than the actual judgments, they may still evaluate systems in a similar way. Figure 2b shows how the ordering of system on real queries compare with system ordering on our fully synthetic test collection (synthetic queries + synthetic judgments). It can be seen that evaluation on the fully synthetic test



Figure 3: Scatter plots of the effectiveness of TREC DL 2023 runs based on synthetic vs. real test collections to analyse the bias towards systems using the same language model as the one used in synthetic test collection construction.

collection results in similar results to human queries with human judgments in terms of system ordering, with a Kendall's τ value of 0.8568. For comparison, Faggioli et al. [11] report Kendall's $\tau = 0.86$ for NDCG@10 on a similar experiment where only judgments were synthetically generated using GPT-3.5. Interestingly, compared with the system rankings using the human judgments on synthetic queries (see Figure 1), using the synthetic judgments generated with LLMs on synthetic queries result in higher correlations and Kendall's τ values as shown in Figure 2b.

3 ANALYSIS OF BIAS IN SYSTEM EVALUATION

One potential issue with using synthetic queries and judgements in test collection construction is the possible bias these collections may exhibit towards systems that are based on a similar approach (similar language model) to the one that was used in the synthetic test collection construction process (e.g., synthetic test collections constructed using T5 might favour systems that are based on T5).

In order to analyse the possible bias, we categorised the runs submitted to TREC DL 2023 based on the approach they use³ (i.e., language models used in their ranking or retrieval pipeline), resulting in four different system categories: systems based on GPT (×), T5 (|), GPT + T5 (+) (i.e., a combination of GPT and T5), and others (•) (i.e., traditional methods such as BM25, or any model that does not use either GPT or T5). Figure 3a shows that the synthetic test collection we have constructed that contains synthetic queries generated by LLMs (T5 and GPT-4) exhibits little to no bias towards LLM-based systems.

To further analyse possible bias that might arise from a system using a similar language model as the one used in test collection construction, Figure 3b shows how system performance computed on synthetic test collections generated using queries generated by GPT-4 compare with system performance on real test collections. It can be seen that synthetic test collections based on GPT-4 do not systematically overestimate the performance of systems based on GPT. Similarly, as can be seen in Figure 3c, synthetic test collections with queries generated using T5 exhibit almost no bias towards systems based on T5. Similar results were obtained on other evaluation metrics such as average precision and NDCG@100, results for which are omitted due to space limitations. However, all three plots in the figure show that for all system types, systems consistently achieve higher performance on synthetic test collections when compared to real queries, suggesting that synthetic test collections tend to be easier than real queries and they tend to overestimate system performance across all system types.

Although previous studies [14] on LLM evaluation discussed the potential bias towards LLM-generated text when we use LLMs for evaluation, in our experiments we did not observe a very clear evidence of systematic bias, where runs using GPT-4 were favored when evaluated using synthetic GPT-4 queries, or where runs using T5 were favored when evaluated on synthetic T5 queries. While our results look encouraging, we would like to emphasize that our results are based on one test collection we have constructed and further experiments are needed to analyse potential biases that might arise from using a fully synthetic test collection.

4 CONCLUSION AND FUTURE WORK

In this paper, we explored the construction of fully synthetic test collections using LLMs. Overall, our analysis suggests that by using fully synthetic test collections consisting of synthetically generated queries and judgments, it is possible to obtain evaluation results that are similar to evaluation results obtained using the traditional test collection approach. Future studies can examine more advanced prompting methods and different LLMs to compare with the test collection we have created.

Exploring the potential biases that may arise from generating a fully synthetic test collection is crucial for ensuring the quality, fairness, and reliability of the test collections. While our preliminary analysis showed that the synthetic test collections we have generated using a particular LLM do exhibit little to no bias towards systems based on the same LLM, further research is needed to provide a deeper understanding of the potential biases and to develop appropriate mitigation strategies.

ACKNOWLEDGMENTS

This research is supported by the Engineering and Physical Sciences Research Council [EP/S021566/1] and the EPSRC Fellowship titled "Task Based Information Retrieval" [EP/P024289/1].

³To this end, we carefully analysed the metadata file of submissions. GPT models are GPT-4 or GPT-3.5 and T5 models include MonoT5, FlanT5, and RankT5.

Synthetic Test Collections for Retrieval Evaluation

REFERENCES

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. arXiv preprint arXiv:2303.08774 (2023).
- [2] Marwah Alaofi, Luke Gallagher, Mark Sanderson, Falk Scholer, and Paul Thomas. 2023. Can generative llms create query variants for test collections? an exploratory study. In Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval. 1869–1873.
- [3] Arian Askari, Mohammad Aliannejadi, Chuan Meng, Evangelos Kanoulas, and Suzan Verberne. 2023. Expand, Highlight, Generate: RL-driven Document Generation for Passage Reranking. In Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing. 10087–10099.
- [4] Javed A Aslam, Virgil Pavlu, and Emine Yilmaz. 2006. A statistical method for system evaluation using incomplete judgments. In Proceedings of the 29th annual international ACM SIGIR conference on Research and development in information retrieval. 541–548.
- [5] Jianzhu Bao, Rui Wang, Yasheng Wang, Aixin Sun, Yitong Li, Fei Mi, and Ruifeng Xu. 2023. A synthetic data generation framework for grounded dialogues. In Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers). 10866–10882.
- [6] 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.
- [7] Cheng-Han Chiang and Hung-yi Lee. 2023. Can Large Language Models Be an Alternative to Human Evaluations? arXiv preprint arXiv:2305.01937 (2023).
- [8] Cyril Cleverdon. 1967. The Cranfield tests on index language devices. In Aslib proceedings, Vol. 19. MCB UP Ltd, 173-194.
- [9] Nick Craswell, Bhaskar Mitra, Emine Yilmaz, Daniel Campos, Ellen M Voorhees, and Ian Soboroff. 2021. TREC deep learning track: Reusable test collections in the large data regime. In Proceedings of the 44th international ACM SIGIR conference on research and development in information retrieval. 2369–2375.
- [10] Nick Craswell, Bhaskar Mitra, Emine Yilmaz, Hossein A. Rahmani, Daniel Campos, Jimmy Lin, Ellen M. Voorhees, and Ian Soboroff. 2024. Overview of the TREC 2023 Deep Learning Track. In *Text REtrieval Conference (TREC)*. NIST,

 $TREC.\ https://www.microsoft.com/en-us/research/publication/overview-of-the-trec-2023-deep-learning-track/$

- [11] Guglielmo Faggioli, Laura Dietz, Charles LA Clarke, Gianluca Demartini, Matthias Hagen, Claudia Hauff, Noriko Kando, Evangelos Kanoulas, Martin Potthast, Benno Stein, et al. 2023. Perspectives on large language models for relevance judgment. In Proceedings of the 2023 ACM SIGIR International Conference on Theory of Information Retrieval. 39–50.
- [12] Ellen Jiang, Kristen Olson, Edwin Toh, Alejandra Molina, Aaron Donsbach, Michael Terry, and Carrie J Cai. 2022. Promptmaker: Prompt-based prototyping with large language models. In CHI Conference on Human Factors in Computing Systems Extended Abstracts. 1–8.
- [13] Zhuoyan Li, Hangxiao Zhu, Zhuoran Lu, and Ming Yin. 2023. Synthetic Data Generation with Large Language Models for Text Classification: Potential and Limitations. arXiv preprint arXiv:2310.07849 (2023).
- [14] Yang Liu, Dan Iter, Yichong Xu, Shuohang Wang, Ruochen Xu, and Chenguang Zhu. 2023. Gpteval: Nlg evaluation using gpt-4 with better human alignment. arXiv preprint arXiv:2303.16634 (2023).
- [15] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research* 21, 140 (2020), 1–67.
- [16] Thilina C Rajapakse and Maarten de Rijke. 2023. Improving the Generalizability of the Dense Passage Retriever Using Generated Datasets. In European Conference on Information Retrieval. Springer, 94–109.
- [17] Nandan Thakur, Nils Reimers, Andreas Rücklé, Abhishek Srivastava, and Iryna Gurevych. 2021. Beir: A heterogenous benchmark for zero-shot evaluation of information retrieval models. arXiv preprint arXiv:2104.08663 (2021).
- [18] Paul Thomas, Seth Spielman, Nick Craswell, and Bhaskar Mitra. 2023. Large language models can accurately predict searcher preferences. arXiv preprint arXiv:2309.10621 (2023).
- [19] Emine Yilmaz, Evangelos Kanoulas, and Javed A Aslam. 2008. A simple and efficient sampling method for estimating AP and NDCG. In Proceedings of the 31st annual international ACM SIGIR conference on Research and development in information retrieval. 603–610.
- [20] Lichao Zhang, Abel Gonzalez-Garcia, Joost Van De Weijer, Martin Danelljan, and Fahad Shahbaz Khan. 2018. Synthetic data generation for end-to-end thermal infrared tracking. *IEEE Transactions on Image Processing* 28, 4 (2018), 1837–1850.