A Sensitivity Analysis of the MSMARCO Passage Collection[⋆]

Joel Mackenzie[1], Matthias Petri[2], and Alistair Moffat[1]

[1] The University of Melbourne, Australia
[2] Amazon Alexa, United States
{joel.mackenzie,ammoffat}@unimelb.edu.au, mkp@amazon.com

Abstract. The recent MSMARCO passage retrieval collection has allowed researchers to develop highly tuned retrieval systems. One aspect of this data set that makes it distinctive compared to traditional corpora is that most of the topics only have a single answer passage marked relevant. Here we carry out a "what if" sensitivity study, asking whether a set of systems would still have the same relative performance if more passages per topic were deemed to be "relevant", exploring several mechanisms for identifying sets of passages to be so categorized. Our results show that, in general, while run scores can vary markedly if additional plausible passages are presumed to be relevant, the derived system ordering is relatively *in*sensitive to additional relevance, providing support for the methodology that was used at the time the MSMARCO passage collection was created.

Keywords: Retrieval experimentation · Pooling · System comparison

1 Introduction

Offline retrieval evaluations make use of *test collections*, each of which includes a set of *documents*, a set of *topics* (or queries), and a set of relevance judgments (or *qrels*). A *run* is constructed for each combination of system and topic, and then those runs are scored using an *effectiveness metric*, making use of the qrels for the corresponding topic. Finally, the run scores are compared across the systems, usually via a paired (over topics) statistical test [17].

The recent MSMARCO passage test collection [5, 6, 15] differs from previous test collections, with the "documents" short *passages* extracted from larger entities, and with very sparse qrels. In particular, there is only a single passage marked as relevant for the majority of topics, and no passages are marked non-relevant. As a result, effectiveness metric values for most runs are drawn from a small set of distinct values; and systems might risk being deemed inaccurate if they present equally-attractive, but unjudged, answers in different orders.

[⋆] The work we report here was carried out in the period May-August 2021, and was conceived and executed independently of and concurrently with the complementary work of Arabzadeh et al. [1].

Query: how long is super bowl game

Passage 1: A traditional football game is approximately 3 hours long. However, the Super Bowl is approximately 4 hours long from start to finish. [*27 more words*]

Passage 2: However, the Super Bowl is approximately 4 hours long from start to finish. The game is longer due to the lengthened half time show and the focus on advertising and commercial breaks. [*82 more words*]

Passage 3: How long does the Super Bowl usually last? The Super Bowl is typically four hours long. The game itself takes about three and a half hours, with a 30 minute halftime show built in. [*63 more words*]

Fig. 1: One topic and (extracts of) three passages of MSMARCO. Only the third is marked as being relevant; the other two are neither relevant nor non-relevant.

To illustrate this risk, Figure 1 shows one of the MSMARCO topics, and the first three passages returned by a standard BM25 run. The passage ranked third is the (only) one that has been judged relevant for this topic, despite the apparent suitability of the first two. A system that, perhaps just by chance, had the third answer at rank one or rank two would have a notably different effectiveness score. There are many other instances of this effect.

Our goal here is to explore the extent to which unjudged, but arguably relevant, answers might affect system effectiveness scores, and also system versus system comparisons. To do that we develop a range of passage orderings based on "clairvoyant" knowledge of the qrel set, including ones that are a result of fusing multiple held-out systems' runs, and ask a critical question: if more passages taken from those lists of plausible candidates are deemed to be relevant, what happens to system scores and comparative orderings? Our results show that run scores vary markedly, but that the derived system ordering is relatively *in*sensitive to additional relevance, providing support for the methodology that was used at the time the MSMARCO passage collection was created.

2   Experimental Design

Our goal is to explore score consistency and system ordering stability as additional passages are assumed to be relevant, augmenting the set of passages labeled "relevant" in the original MSMARCO qrels. The next few paragraphs describe the process used for identifying plausible candidate passages.

Notation. Let S be a *retrieval system*. When provided with a query q, S returns a ranked list of documents (here, passages) $S(q)$. Further, let M be an *effectiveness metric* which returns a score derived from a ranking $S(q)$ and a set of relevance judgments for that query, $J(q)$. That is, $M(S(q), J(q))$ is the score assigned by metric M to system S for query q, relative to the judgments $J(q)$. It also convenient to take q as being given, and use the shorthand $M(S, J) \equiv M(S(q), J(q))$. Finally, let $T_d(L)$ be the first d items in list L. For example, $T_d(S)$ is the first d elements of the ranking generated by S for some query q.
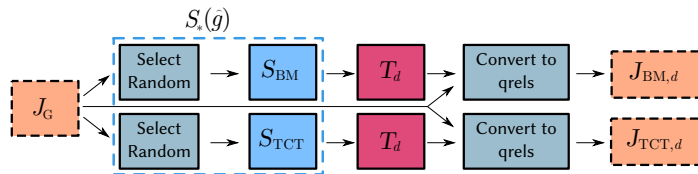
Fig. 2: Extrapolated judgments. A single gold answer from $J_G(q)$ is used as a query to generate BM25 and TCT runs. The d highest-ranked non-gold passages are taken to be relevant, and added to $J_G(q)$.

Gold Answers. The MSMARCO qrels establish at least one *gold answer* for each query q; we denote q's set of gold answers by $J_G(q)$. In the MSMARCO collection, $|J_G(q)| = 1$ for most q; and we suppose that $\hat{g}(q)$ is that passage. When $|J_G(q)| > 1$, we select $\hat{g}(q) \in J_G(q)$ as a random choice. Given a system, $S(\hat{g}(q))$ can be computed via a *query-by-document* [24] mechanism, with $\hat{g}(q)$ likely (but by no means guaranteed) to be the top-ranked passage. That is, in the majority of cases, $T_1(S(\hat{g}(q))) = \hat{g}(q)$; whereas there is no expectation that $T_1(S(q)) = \hat{g}(q)$.

Clairvoyant Rankings and Seed Systems. Our experiments are based on the hypothesis that if some answer d is "close" to $\hat{g} \in J_G(q)$, then d is also a plausible candidate for relevance to q [3]. To quantify closeness, we use the "query-by-passage" ordering $S(\hat{g})$, and determine the rank at which d occurs. We can think of $S(\hat{g})$ as being a *clairvoyant* ranking, since it is derived from knowledge of a relevant passage; that is, via a relevance feedback loop [18, 20]. We use two different *seed systems* to generate those rankings:

- S_{BM} is a bag-of-words BM25 run generated using the PISA search system [13] over an Anserini index [23] transferred via the Common Index File Format [9] following Mackenzie et al. [12].
- S_{TCT} is the neural TCT-ColBERT-V2-HN+ system described by Lin et al. [11]. We use Pyserini [10] to conduct brute force retrieval via FAISS [7].

Extrapolated Qrels. The experimental pipeline takes the S_{BM} and S_{TCT} runs, together with $J_G(q)$ and one gold passage $\hat{g} \in J_G(q)$, and generates three sets of variable-size *extrapolated qrels*:

- $J_{BM,\,d}$ contains $J_G(q)$ plus exactly d additional "deemed relevant" passages generated via the BM25-based query-by-passage process, see Figure 2:

$$J_{BM,\,d}(q) = J_G(q) \ \cup \ T_d(S_{BM}(\hat{g}(q)) \setminus J_G(q)).$$

- $J_{TCT,\,d}$ is derived from S_{TCT} in the same way.
- $J_{FUS,\,d}$ makes use of ten BM25 runs and ten TCT runs, with those query-by-passage runs in turn based on two original query-by-passage runs, one from
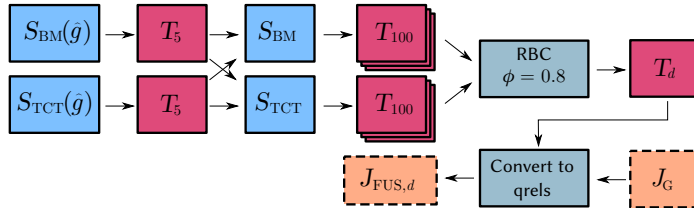
Fig. 3: Applying rank fusion. The top five passages from each of the BM25 and TCT query-by-passage runs (10 passages) are used as queries to both BM25 and TCT. The resultant 20 runs (to depth 100) are then fused using RBC [2], and the top-d passages of that final run are deemed relevant and joined with $J_{\mathrm{G}}(q)$.
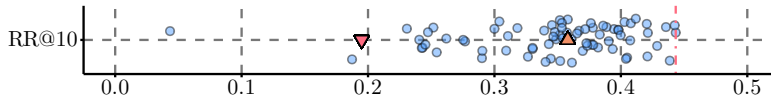

Fig. 4: Average system scores for RR@10 across the dev set, with triangles marking the BM25 (left) and TCT systems (right), and the red dot-dashed line representing the best dev run on the official leader board as of 24 August 2021.

each system. The fusion process applies the rank-biased centroid approach [2] to those twenty runs to obtain a single merged run; finally, d top passages are taken from it and added to $J_{\mathrm{G}}(q)$, see Figure 3.

Those three extrapolated qrels sets, parameterized by the augmentation parameter d, are employed in the experiments described in the next section, along with the original gold judgments J_{G}. Note again that for the majority of queries, $|J_{\mathrm{G}}(q)| = 1$. In contrast, $|J_{\mathrm{BM},d}(q)| = |J_{\mathrm{TCT},d}(q)| = |J_{\mathrm{FUS},d}(q)| = |J_{\mathrm{G}}(q)| + d$.

3   Experiments

Experimental Setup. We make use of the MSMARCO Passage Ranking Collection (version 1). The *dev set* contains qrels for 6,980 queries, with 6,590 (94.4%) having a single positive label ($|J_{\mathrm{G}}(q)| = 1$). Of the other 390 (5.6%) queries, 331 have two labels, 51 have three labels, and 8 have four labels. There are no negative (non-relevant) labels provided in the MSMARCO qrels.

A total of 75 system dev runs were used, truncated to 10 passages for each query, and with effectiveness computed "@10" in all cases. The runs were a mix of ones that we generated ourselves, and runs provided by the MSMARCO chairs. Figure 4 shows the distribution of system average (over queries) RR@10 scores. The two runs used to form the extrapolated qrels were not included in the 75.

Score Sensitivity. Figure 5 shows how metric scores are affected as additional "deemed relevant" passages are added into the qrels in a controlled manner. Unsurprisingly, all three metrics have upward trends, with the distinctive behavior
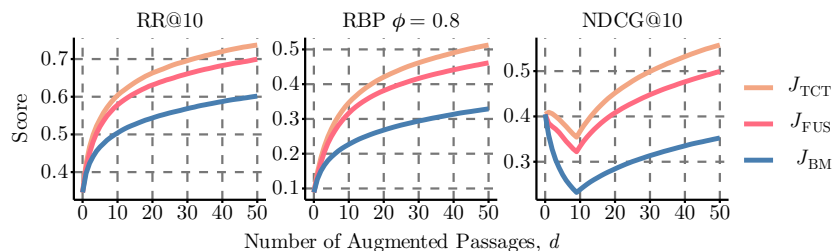
Fig. 5: Effectiveness scores, averaged across 75 system runs and the dev query set (that is, 75×6980 values) as a function of d, the number of further passages deemed relevant, for three metrics and three sets of extrapolated qrels. The original J_{G}-only metric scores correspond to $d = 0$.
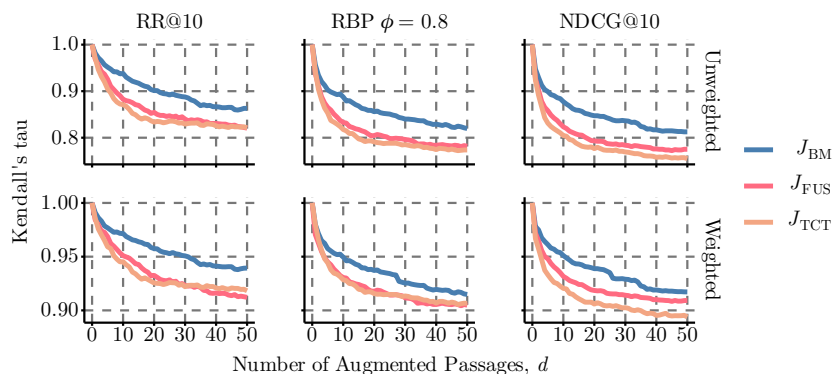

Fig. 6: Unweighted (top) and weighted (bottom) Kendall's τ correlations for 75 systems, all measured relative to the reference ordering computed using J_{G}. Three different sets of extrapolated judgments are used, and three metrics.

of NDCG@10 in the vicinity of $d = 10$ a consequence of the normalization process it employs. The J_{TCT} judgments result in the highest average system/query scores; while the J_{BM} judgments give the least score increase.

Inter-System Sensitivity. The more important question is whether adding judgments – in this case, extrapolated ones – alters system relativities. In this experiment, the J_{G}-induced reference ordering of the 75 systems is compared with the orderings generated using the "plus d" extrapolated judgment sets. Unweighted [8] and top-weighted Kendall's τ coefficients were computed, in the latter case with a weight of $1/(k+1)$ assigned to the system at rank k [19, 21].

Figure 6 provides results. As increasing numbers of plausible passages are deemed to be relevant, the system orderings tend to slowly diverge from the reference ordering. But the top-weighted τ scores (the lower row) for all three effectiveness metrics remain above 0.9, even at $d = 20$, indicating high consistency in

> **Query:** idaho definition of signed
>
> **Passage 1:** of the state of Idaho and to the administrative jurisdiction of the Idaho real estate com - mission, and shall be subject to all penalties and remedies available under Idaho law for any violation of this chapter.
>
> **Passage 2:** July 2017 Idaho Real Estate License Law & Rules - i IDAHO REAL ESTATE COMMISSION STAFF (208) 334-3285 Administration MiChell M. Bird - Executive Director michell.bird@irec.idaho.gov.....ext. 105 Jessica Valerio - Administrative Assistant 2

Fig. 7: A query and two passages, the first from J_{G}, the second from J_{FUS}. After discussion we judged the second passage to "answer the query to approximately the same or greater accuracy" as the first. In absolute terms, neither is helpful.

Table 1: Limited-scale additional judgments as a demonstration of concept.

	$d = 1$	$d = 2$	$d = 10$
Fraction judged to be "as relevant as $\hat{g}(q)$"	20/20	16/20	14/20

the relative performance of the better-scoring systems. The J_{BM} approach gives the highest τ values, perhaps because it disrupts the metric scores the least.

Judgment Validation. To provide a limited-scale validation of the extrapolation method, three passages were extracted from J_{FUS} for each of twenty queries, those at ranks 1, 2, and 10. Those sixty passages were then judged by each of the authors, and discussed to reach consensus where we disagreed. The question considered in all cases was whether the added passage "answered the query to approximately the same or greater accuracy than the first passage", that is, we used the gold passage $\hat{g}(q)$ as an *anchor*. One such pair is shown in Figure 7; in this particular example, neither the anchor passage nor the extrapolated one are relevant, but nor is the second passage less relevant than the first.

The fraction at each depth d for which the consensus answer was "yes" is shown in Table 1. Of the twenty $d = 1$ passages, 19 were simply the gold passage, $\hat{g}(q)$, confirming that in most cases $T_1(S(\hat{g}(q))) = \hat{g}(q)$. On the other hand, the results for $d = 2$ and $d = 10$ provide compelling preliminary evidence that there are many more relevant passages in the MSMARCO collection than are captured by the reference qrels, supporting the claims made by Qu et al. [16].

4 Conclusion

We have explored the sensitivity of the MSMARCO collection, measuring the extent to which system scores and system orderings are stable if more than one passage per query is assumed to be relevant. Our results show that scores themselves increase as positive qrels are added, but that system orderings are comparatively resilient. These findings add credibility to the process used to construct the MSMARCO passage collection.

As a final comment, we observe that effective training of neural retrieval systems requires negative examples as well as positive ones [11, 16], a question that has been considered by a range of authors [4, 14, 22, 25]. Developing an effective mechanism for determining documents or passages that are plausible answers, but are non-relevant – as distinct from ones that are patently non-relevant – is thus another interesting challenge.

References

1. N. Arabzadeh, A. Vtyurina, X. Yan, and C. L. A. Clarke. Shallow pooling for sparse labels. arXiv:2109.00062, Aug. 2021.

2. P. Bailey, A. Moffat, F. Scholer, and P. Thomas. Retrieval consistency in the presence of query variations. In *Proc. SIGIR*, pages 395–404, 2017.

3. B. Carterette and J. Allan. Semiautomatic evaluation of retrieval systems using document similarities. In *Proc. CIKM*, pages 873–876, 2007.

4. D. Cohen, S. M. Jordan, and W. B. Croft. Learning a better negative sampling policy with deep neural networks for search. In *Proc. ICTIR*, pages 19–26, 2019.

5. N. Craswell, B. Mitra, E. Yilmaz, D. Campos, and J. Lin. MS MARCO: Benchmarking ranking models in the large-data regime. In *Proc. SIGIR*, pages 1566–1576, 2021.

6. N. Craswell, B. Mitra, E. Yilmaz, D. Campos, E. M. Voorhees, and I. Soboroff. TREC deep learning track: Reusable test collections in the large data regime. In *Proc. SIGIR*, pages 2369–2375, 2021.

7. J. Johnson, M. Douze, and H. Jégou. Billion-scale similarity search with GPUs. *IEEE Trans. Big Data*, 7(3):535–547, 2021.

8. M. G. Kendall. A new measure of rank correlation. *Biometrika*, 30(1/2):81–93, 1938.

9. J. Lin, J. Mackenzie, C. Kamphuis, C. Macdonald, A. Mallia, M. Siedlaczek, A. Trotman, and A. de Vries. Supporting interoperability between open-source search engines with the common index file format. In *Proc. SIGIR*, pages 2149–2152, 2020.

10. J. Lin, X. Ma, S.-C. Lin, J.-H. Yang, R. Pradeep, and R. Nogueira. Pyserini: A Python toolkit for reproducible information retrieval research with sparse and dense representations. In *Proc. SIGIR*, pages 2356–2362, 2021.

11. S.-C. Lin, J.-H. Yang, and J. Lin. In-batch negatives for knowledge distillation with tightly-coupled teachers for dense retrieval. In *Proc. RepL4NLP at ACL 2021*, pages 163–173, 2021.

12. J. Mackenzie, Z. Dai, L. Gallagher, and J. Callan. Efficiency implications of term weighting for passage retrieval. In *Proc. SIGIR*, pages 1821–1824, 2020.

13. A. Mallia, M. Siedlaczek, J. Mackenzie, and T. Suel. PISA: Performant indexes and search for academia. In *Proc. OSIRRC at SIGIR 2019*, pages 50–56, 2019.

14. B. Mitra, F. Diaz, and N. Craswell. Learning to match using local and distributed representations of text for web search. In *Proc. WWW*, pages 1291–1299, 2017.

15. 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 *Proc. NeurIPS Workshop on Cognitive Computation*, pages 96–105, 2016.

16. Y. Qu, Y. Ding, J. Liu, K. Liu, R. Ren, W.-X. Zhao, D. Dong, H. Wu, and H. Wang. RocketQA: An optimized training approach to dense passage retrieval for open-domain question answering. In *Proc. NAACL*, pages 5835–5847, 2021.

17. M. Sanderson. Test collection based evaluation of information retrieval systems. *Found. Trnd. Inf. Retr.*, 4(4):247–375, 2010.

18. M. Sanderson and H. Joho. Forming test collections with no system pooling. In *Proc. SIGIR*, pages 33–40, 2004.

19. G. S. Shieh. A weighted Kendall's tau statistic. *Statistics & Probability Letters*, 39(1):17–24, 1998.

20. I. Soboroff and S. Robertson. Building a filtering test collection for TREC 2002. In *Proc. SIGIR*, pages 243–250, 2003.

21. S. Vigna. A weighted correlation index for rankings with ties. In *Proc. WWW*, pages 1166–1176, 2015.

22. L. Xiong, C. Xiong, Y. Li, K.-F. Tang, J. Liu, P. N. Bennett, J. Ahmed, and A. Overwijk. Approximate nearest neighbor negative contrastive learning for dense text retrieval. In *Proc. ICLR*, 2021.

23. P. Yang, H. Fang, and J. Lin. Anserini: Reproducible ranking baselines using Lucene. *J. Data Inf. Qual.*, 10(4):1–20, 2018.

24. Y. Yang, N. Bansal, W. Dakka, P. Ipeirotis, N. Koudas, and D. Papadias. Query by document. In *Proc. WSDM*, pages 34–43, 2009.

25. J. Zhan, J. Mao, Y. Liu, J. Guo, M. Zhang, and S. Ma. Optimizing dense retrieval model training with hard negatives. In *Proc. SIGIR*, pages 1503–1512, 2021.