

Exploring the Nexus Between Retrievability and Query Generation Strategies

Aman Sinha[†][0009-0005-8227-3110], Priyanshu Raj Mall[†][0009-0005-4163-9649],
and Dwaipayan Roy^{*}[0000-0002-5962-5983]

Indian Institute of Science Education and Research Kolkata, India
{as18ms065, prm18ms118, dwaipayan.roy}@iiserkol.ac.in

[†]These authors contributed equally to this work.

^{*}Corresponding author

Abstract. Quantifying bias in retrieval functions through document retrievability scores is vital for assessing recall-oriented retrieval systems. However, many studies investigating retrieval model bias lack validation of their query generation methods as accurate representations of retrievability for real users and their queries. This limitation results from the absence of established criteria for query generation in retrievability assessments. Typically, researchers resort to using frequent collocations from document corpora when no query log is available. In this study, we address the issue of reproducibility and seek to validate query generation methods by comparing retrievability scores generated from artificially generated queries to those derived from query logs. Our findings demonstrate a minimal or negligible correlation between retrievability scores from artificial queries and those from query logs. This suggests that artificially generated queries may not accurately reflect retrievability scores as derived from query logs. We further explore alternative query generation techniques, uncovering a variation that exhibits the highest correlation. This alternative approach holds promise for improving reproducibility when query logs are unavailable.

Keywords: Reproducibility · Retrievability · Empirical Study.

1 Introduction

In the field of Information Retrieval (IR), the study of biases and fairness of retrieval models together with their effectiveness has been an active area of research [23,26]. Among the various metrics and methodologies developed, *retrievability* has emerged as a significant way of measuring how accessible a document is in a collection, given a retrieval system and a set of queries [8]. It reflects the degree to which the system favors certain documents while retrieval over others, regardless of their relevance. Retrievability can be used to evaluate the fairness, diversity, and coverage of retrieval systems, as well as to identify and analyze the factors that affect the accessibility of documents [7]. However, calculating the retrievability of a collection of documents is a complex task comprising several

interconnected steps. Initially, a set of queries is selected based on which the retrievals are performed. Researchers have used query logs as well as artificially generated queries in this step. Subsequently, a retrieval model is configured with optimal parameters and the actual retrieval process follows, producing ranked document lists. Retrievability scores are then computed based on the rank of the documents in the retrieved lists. Finally, the computed scores are used to estimate “retrievability bias” by metrics like the Gini coefficient[27], Atkinson index [4], Theli index [30] or Palma index [39] borrowed from economics, to quantify bias and inequality (mainly in wealth) within a population.

Various query generation techniques employed at the initial stage can yield distinct query sets, potentially resulting in diverse retrievability scores for the same collection and retrieval system. Previous research has indicated that retrievability scores calculated using simulated queries differ from those obtained using actual queries extracted from query logs [45]. However, a comprehensive and systematic comparison of retrievability scores computed using various query simulation methods has yet to be undertaken. The discrepancy, if any, poses a challenge to the reproducibility of retrievability experiments as different researchers may employ different query generation methods, leading to diverse conclusions on the same collection. To investigate the potential impact on result reproducibility and experimental validation, this paper conducts a comparative analysis of computed retrievability scores using the most commonly adopted query generation techniques in the research community. The main contributions of this paper include:

- Providing an overview of query generation techniques commonly used in retrievability experiments documented in the literature.
- Conducting a comprehensive empirical assessment of the reproducibility of retrievability experiments employing different query generation techniques and test collections.

Additionally, we address the following research questions in this paper:

- RQ1** - While keeping the other parameters constant, does varying the query simulation techniques lead to differences in retrievability scores, thus impacting the reproducibility of results?
- RQ2** - Do these differences in retrievability scores exhibit consistency across various collections, thereby influencing the reproducibility of results?

The rest of the paper is organized as follows. In the next section, we provide essential background information and review previous research related to retrievability, with a particular emphasis on the prevalent query simulation techniques in this field. Section 3 delves into the specifics of our experimental setups before presenting the empirical result on benchmark test collections in Section 4. We conclude the paper in Section 5 highlighting the findings considering the issues with reproducing results, and suggesting some directions to pursue further study. The codes and the simulated queries are available from the following repository: <https://mallpriyanshu.github.io/retrievability-survey/ecir2024>.

2 Background and Motivation

2.1 Retrievability: A Measure of Accessibility

Retrievability is a measure that evaluates the ease with which a document can be retrieved in a specific configuration of an IR system. Azzopardi and Vinay [8] introduced the concept and formally defined the retrievability score for a document d in a collection D as $r(d)$ with respect to a given IR system. Typically, the calculation of retrievability involves the execution of the following sequence of steps [36]: *i*) query set generation, *ii*) system configuration and retrieval, *iii*) computation of retrievability scores, and *iv*) summarising the scores.

For an ideal calculation of retrievability scores, a comprehensive universe of all possible queries is needed, although obtaining this is practically unattainable. In practice, researchers commonly select from two alternatives: either using a sufficiently large query log containing diverse user queries [42,45] or simulating a set of queries from the collection that adequately represents its overall topic distribution [8,9,14,16]. Once an appropriate set of queries has been chosen, the second step is selecting a retrieval model with an optimal parameter setting. Prior research has extensively examined the impact of altering system configurations [8,15,50] where the authors conclude that changing the retrieval model can significantly impact the retrievability scores of a document collection. The retrievals are then performed with the set of queries, and the retrievability of a document d is computed in the third step as *the summation of the number of times d is retrieved within the top c position*. Mathematically, the retrievability $r(d)$ of a document d ($d \in D$) with respect to an IR system is computed using the formula in Equation 1.

$$r(d) = \sum_{q \in Q} o_q \cdot f(k_{dq}, c) \quad (1)$$

Here, o_q indicates the likelihood of the query (usually considered uniform across query set) and k_{dq} is the rank of d for query q . Calculating $r(d)$ relies on the choice of a utility function $f(\cdot)$, either cumulative or gravity-based. Previous studies have shown a high correlation between these two measures [47,48], and consequent research predominantly reports retrievability analysis using the cumulative-based function [13,21,45,49]. Defined as a binary function, the cumulative utility function $f(k_{dq}, c)$ is equal to 1 if document d is retrieved within the top c documents for query q ; otherwise, it is considered 0. This utility function provides a straightforward interpretation of the retrievability score for each document. It simply corresponds to the number of times the document is retrieved within the top c ranks. Documents beyond the top c positions are disregarded, simulating a user who only examines the first c results. Therefore, the more queries that retrieve a document within rank c , the higher its retrievability score.

The final step involves summarizing the retrievability scores across the collection to assess existing bias. This quantification is commonly achieved using the Gini coefficient, a measure frequently employed in economics and social sciences

to assess inequality within a population [27]. Calculating the retrievability bias within a document collection, the Gini coefficient is computed as follows:

$$G = \frac{\sum_{i=1}^N (2i - N - 1) \cdot r(d_i)}{N \sum_{j=1}^N r(d_j)} \quad (2)$$

where N represents total number of documents in the collection ($|D|$). It provides a decimal value in the range $[0, 1]$. A Gini coefficient of zero denotes perfect equality, indicating that all documents in the collection have an equal retrievability score according to $r(d)$. Conversely, a Gini coefficient of one indicates total inequality, with only one document having $r(d) = |Q|$ while all other documents have $r(d) = 0$. Other inequality metrics, such as Atkinson index [4], Theli index [30], Palma index [39] etc., are also explored in [53] where the authors argue that various inequality metrics generally agree on which system and settings minimize the retrievability bias and the least inequality indicated by Gini coefficient is consistent with the other metrics discussed. Hence, the Gini coefficient G is commonly employed in subsequent works [8,16,18] to quantify the amount of bias in the distribution of the retrievability scores in a collection of documents.

To visually assess the retrievability distribution in a document collection, one can calculate retrievability scores for each document using Equation 1 and plot the sorted values to create a *Lorenz curve*. This curve represents the cumulative score distribution of documents sorted by retrievability scores in ascending order. If the retrievability scores are evenly distributed, the Lorenz Curve will be linear; a skewed curve in contrast indicates a greater level of inequality or bias. The Lorenz curve is a concept commonly used in economics to depict wealth distribution, where a straight line signifies equal wealth distribution across a population, while a skewed curve represents inequality, as shown in Figure 1.




Fig. 1: A sample Lorenz Curve. The values belonging to the distribution in which the disparity is to be observed are arranged in a sorted manner along the X-axis.

Retrievability, and the underlying theory of retrievability, have found applications in various domains. For instance, it has been used in the development of inverted indexes to enhance the efficiency and performance of retrieval systems by capitalizing on terms that contribute to a document’s retrievability [41]. Additionally, retrievability has been leveraged to investigate bias in search engines and retrieval systems on the web [6] and within patent collections [18], leading to improvements in system efficiency during pruning processes [54].

The initial and most critical step while calculating the retrievability of a collection of documents involves query generation, followed by retrieval using the generated queries, and finally, the computation of retrievability scores based on the appearance of documents within a rank cutoff. Prior research has extensively examined the impact of altering system configurations [8,9,15,45,50], modifying the nature of the function [48], and varying score summarization techniques [53]. However, there is limited empirical evidence concerning the influence of modifying query sets on the resultant retrievability scores. Filling the void, in this paper, we study the effect of various query-generation techniques on the retrievability scores of documents in a collection. We discuss some of the well-known techniques popular among researchers for query generation in the next section.

2.2 Query Generation

The initial step in a retrievability computation typically involves obtaining a set of all possible queries addressable by the collection (Q in Equation 1). As per Equation 1, the most ideal calculation of $r(d)$ would require an extensive set Q . In general, a query log would be the preferred option for Q ; however, obtaining such a log can be difficult and often not feasible. Hence, researchers have inclined towards artificial simulation of the set Q from the collection. Following the inception of the concept of retrievability in [8], a popular choice among researchers in simulating Q based on term statistics in the collection itself [14,15,16,17,18,19,21,38]. The approach by Azzopardi and Vinay [8] in generating the query set was influenced by Callan and Connel [24]. The authors generated a set of queries, which included single-term queries, formed by selecting each term in the vocabulary that appeared at least 5 times, and bi-term queries, formed by selecting each bigram in the collection that appeared at least 20 times. The list of bigrams was truncated at 20 million, and each query was issued to the system to estimate its retrievability. Azzopardi and Bache later followed a similar methodology to generate query sets for the AP and WSJ collections by ranking the most frequent 100,000 bigrams in each collection [5]. The same approach was also adapted in [9] where the authors applied retrievability to quantify the accessibility of patent documents. Although the queries used in patent search are often longer than two words, the authors choose two-word queries because this means there will be “just one Boolean operator and it thus affords a comparison between the use of AND and OR within the various IR models”. In the other studies on retrievability, the prevalent choice among researchers has been the adoption of a similar approach for query generation [47,48,49,50,51,52,53].

In their work on patent retrieval [14], the authors created queries using terms from the *claim section* of the patent documents. Single-term queries are formed from each document with a term frequency lower bound of 2. To make longer queries, two and three-term queries are generated by the combination of the frequent terms in that document. In their subsequent works [15,16], the authors employed a controlled query generation technique proposed in [31] for patent retrieval. Subsequent works from the same research group introduced a different approach that involves generating query sets for each document in the corpus. They select terms that appear more than once within the document, resulting in four query subsets: single terms, combinations of two terms, combinations of three terms, and combinations of four terms [17,18,19]. Some further truncation of query set was adapted in [11,12,13,20,22] where queries were generated with combinations of terms that appear more than once in the document. Stopwords and terms having document frequency $> 25\%$ of collection size were removed.

Assessing document retrievability in a collection necessitates a comprehensive set of queries that cover all aspects of the collection. Ideally, this would involve an extensive query log spanning a significant time period. However, acquiring such a log is often costly and unavailable, prompting researchers to resort to simulated queries. Despite the challenges, some research has been conducted using real user query logs acquired from certain closed domains. The authors in [44,45] use a set of real queries submitted to the National Library of the Netherlands, between March 2015 to July 2015. Additionally, to make the results comparable to [8], the authors create a set of simulated queries: they select the top 2 million unigrams and bigrams respectively as single term and two term queries. All the queries were finalized after stemming and removal of stopwords and operators. On comparing the characteristics of simulated queries to those formed by real user queries, the authors report significant differences regarding the composition, the number of unique terms used, and the presence of named entities: real user queries exhibited a considerably higher fraction of named entities compared to the simulated query set. Additionally, due to the considerably low number of real queries in contrast to the vast size of the collection (exceeding 102 million), the authors acknowledged the presence of a varying document-query ratio compared to those reported in [8]. Authors in [2] employ a transformer-based generative model docT5query [37] for generating queries and conclude that BERT-based rankers demonstrate less retrievability bias than user-given queries. Recently, Roy et al. [42] studied the retrievability distribution in an integrated retrieval system containing documents in multiple categories [42,43]. However, this work does not conduct a direct comparison of the retrievability scores (and their distributions) when computed using real and simulated query sets.

The contrasting nature of simulated queries, when compared to real queries, is briefly reported in [45]. The various techniques used for generating simulated queries illustrate the diversity in approaches taken at the initial stage of a retrievability analysis by different groups of researchers. This also highlights the absence of a well-defined structure for such an analysis paving a potential issue with reproducing results. It prompts the question: can we anticipate a similar

Table 1: Basic statistics of datasets used for empirical verification of the variation in the computed retrievability scores when different query sets are employed.

Dataset	# documents	Collection Type	# terms
TREC Robust	528,155	News articles	1,502,031
WT10g	1,692,096	Web pages	9,674,707
Wikipedia	6,584,626	Encyclopedia	18,797,260

distribution of retrievability scores when the query set is altered? Answering that question, in this paper, we present a systematic comparison of the retrievability scores computed based on several of these techniques together with real-user query logs. According to the observation in [45], the real queries contain a significantly higher fraction of named entities compared to the simulated query set. Inspired by this observation, we additionally introduce a rule-based query simulation technique for retrievability assessment, further contributing to the examination of this diversity in approaches. To maintain the primary narrative of the paper, we have detailed the method in the Appendix A.

3 Experimental Setup

As presented in Section 2, the set of queries employed in computing the retrievability plays an important role in the computation and our focus of this paper is to foresee whether alteration in this set can affect the final score, potentially leading to challenges in reproducing results for different collections and retrieval models. In this section, we present the datasets and the query sets that are used in this paper to verify this empirically with real data.

3.1 Datasets

We perform the experiments of the retrievability score distributions on two benchmark retrieval datasets as well as on a diverse encyclopedia of documents. Particularly, the TREC ad-hoc Robust collection [46] containing news articles, WT10g collection [29] with web pages, together with the English Wikipedia article dump of February, 2023 are employed in this study¹. All the datasets are preprocessed with stopword removal following SMART stopword list² and stemmed using the Porter Stemmer before indexing them using Apache Lucene. Overall statistics of the datasets are presented in Table 1.

3.2 Simulated Queries

In order to foresee the variation, if any, in the retrievability distribution of documents in a collection, we use the following artificial query generation techniques:

¹ <https://dumps.wikimedia.org/enwiki/>

² <https://tinyurl.com/smart-stopword>

Table 2: Number of queries generated for each dataset using the individual query generation techniques. Note that, the simulated queries are generated from each dataset and applied in that particular dataset for the retrievability computation. AOL query set is filtered separately for each dataset as discussed in Section 3.3.

	Query Type	Number of queries		
		Robust04	WT10g	Wikipedia
SQ ₁ [8]	Simulated	1061184	3902575	5320001
SQ ₂ [13]	Simulated	6903196	18724509	84920098
SQ ₃ [31]	Simulated	1644533	7346091	26272419
RSQ	Simulated	3900000	4500000	4550000
AOL [40]	Real Query	3885206	4302650	4550977

1. **Simulated Query set 1 (SQ₁):** Following the suggestion in [8], the first set of the queries with which the retrievals are performed is formed by a sampling method proposed in [24]. In this process, the terms undergo analysis and filtering, including stemming using the Porter stemmer and removing stopwords. Terms occurring more than 5 times in the collection are treated as single-term queries. Additionally, two-term queries are formed by pairing consecutively occurring terms that have a collection frequency of at least 20 times. These resulting bigrams are then ranked by occurrence, and the top two million are selected as the final set of two-term queries.
2. **Simulated Query set 2 (SQ₂):** The second technique is based on [13]. Specifically, queries are generated with the combinations of terms that appear more than once in the document. Stopwords and other terms with document frequency > 25% of collection size are removed. Queries up to 4 terms are formed using the boolean AND operator, with duplicate queries being removed.
3. **Simulated Query set 3 (SQ₃):** In this approach (based on [31]), a set of related documents is chosen as the source document set with clustering. Language models are defined for both the source documents and the corpus, with relative entropy used for comparison. Terms in the vocabulary are sorted based on their contribution to relative entropy, serving as a term discrimination scoring function. The most discriminating term initiates the query construction process, forming single-term and two-term queries for different query sets. The process continues by identifying next most informative term and generating subsequent queries until all terms have been used.
4. **Rule-based Simulated Query set (RSQ):** Building upon the insight in [45], that real queries contain a notably greater proportion of named entities compared to simulated queries, we introduce a rule-based query simulation technique for retrievability assessment. The approach is presented in Appendix A.

Table 3: Gini coefficient of Retrievability score distribution when computed using different query generation techniques. The highest and the lowest Gini values for each dataset are marked respectively with a superscript \dagger and $*$.

Dataset	SQ ₁	SQ ₂	SQ ₃	RSQ	AOL
Robust	0.3307	0.4556	0.4300	0.3052*	0.6032 \dagger
WT10g	0.5371	0.6391	0.6359	0.5009*	0.6541 \dagger
Wikipedia	0.5380	0.6210	0.6290	0.4820*	0.6798 \dagger

3.3 Real Queries

A majority of the works done in the domain of retrievability employ queries generated using simulation techniques. In this work, we use a general web search query log to make our study conclusive. Specifically, we utilize the AOL query log [40] that contains actual user query click data spanning from March 1st, 2006, to May 31st, 2006. It sparked significant controversy within the media, primarily due to privacy-related concerns [10,28,3]. Although the query log is more than a decade old, it is still widely used in recent literature [25,1,34,33,55,35].

To prepare the real query set for our experiments, the initial step involves extracting all unique queries from the AOL query log. Queries containing periods are eliminated to filter out website links. Given the distinct temporal scopes of our datasets and query sets, a conspicuous challenge arises in the form of topic and vocabulary mismatch. To mitigate this challenge as much as possible, we adopt an exclusionary approach: a query is selected if all of its terms are present in the corpus vocabulary. This strategic refinement serves to align the queries closely with the content of the target collection. Note that, the AOL queries are refined for each collection separately following the above exclusionary approach. Table 2 contains the individual count of queries finalized for each dataset.

3.4 Retrieval Model

In accordance with the insights and recommendations put forth in [8], we only use BM25 as the model for performing the retrieval. Specifically, we use the implementation available in Apache Lucene. To fine-tune the BM25 parameters, we utilize the TREC ad-hoc retrieval topic set [46], which is then applied on the Robust dataset for the retrievability study. Similarly, the TREC web topic sets [29] are employed for parameter tuning, and these adjusted parameters are subsequently applied to the WT10g and Wikipedia datasets.

4 Comparing Retrievability

4.1 Gini Coefficient and Lorenz Curve

We present the Gini coefficient for each dataset utilizing all five query sets in Table 3. The variation in Gini values is graphically presented in Figure 2. In

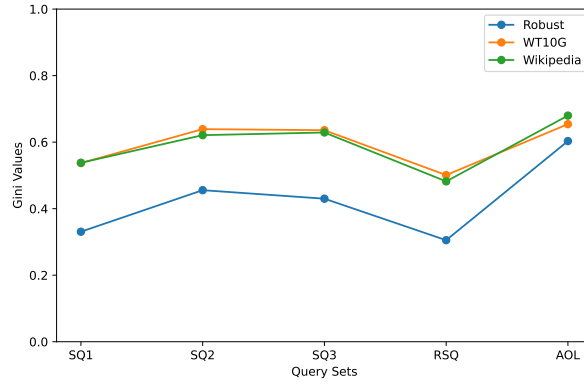


Fig. 2: Gini coefficient for each dataset when the retrievability scores are calculated using different query sets as discussed in Section 3.

the table, the least disparity is observed across all the datasets when utilizing the rule-based simulation for query generation (denoted as RSQ in Table 3). The distributions of retrievability values are presented using Lorenz curve in Figure 3. Furthermore, the original query simulation technique proposed by [8] (indicated by SQ₁ in Table 3) follows a similar trend. The other two simulation techniques (SQ₂ and SQ₃) exhibit higher disparity. In contrast, a significant disparity in retrievability scores across all the datasets is observed when the real-query log (AOL) is used. These results lead to the following conclusions:

- The web collection WT10g and the encyclopedia Wikipedia exhibit much more diverse retrievability scores as compared to the relatively smaller news collection TREC Robust across all the query sets.
- The approach in [8] selects terms from the collection based on their frequency, including non-informative queries (SQ₁). This method results in a more evenly distributed query generation covering a wide range of documents in the collection. As a result, it tends to produce lower Gini values compared to other query sets.
- Realistic queries generated using the rule-based approach (RSQ) result in even lower Gini coefficients across all datasets.
- AOL query logs mostly consist of known item searches rather than exploratory ones leading to a significant disparity in computed retrievability scores, with Gini coefficients consistently exceeding 0.6 across all datasets.

4.2 Comparing Correlation

Table 4 reports the degree of variation in retrievability scores when comparing the simulation-based query generation techniques (SQ₁, SQ₂, SQ₃, and RSQ) along with the AOL query log. We compute the Pearson’s r between the distribution of each pair of retrievability scores computed using the different query




Fig. 3: Lorenz curve with retrievability computed with various query sets on different collections.

sets and report in the upper half of Table 4. Kendall’s τ is presented in the lower half of the table which is computed using the document rank arranged following their retrievability scores. The table shows negligible correlations between scores computed using almost any query set pairs. Interestingly, the highest correlations, in terms of both Pearson’s r and Kendall’s τ occur when using the rule-based technique RSQ (Appendix A). Further, the massive difference in the highest and lowest correlations validates that distinct query sets lead to divergent retrievability scores, affirming both **RQ1** and **RQ2**.

5 Conclusion

The examination of retrievability has predominantly relied on the use of simulated queries since its inception with few works using real query logs. However, a comprehensive comparative analysis of the computed values has been absent in the existing literature. Bridging this gap, in this paper, we examine the impact of five distinct and diverse query sets, ranging from simulated queries to

Table 4: Correlation between retrievability scores when computed using the real query set (AOL) and the four simulated query sets (discussed in Section 3.2). Upper triangle contains Pearson’s r while lower triangle reports Kendall’s rank correlation τ . All differences statistically significant with $p < 0.05$. Highest (in bold) and lowest (superscripts *) correlations for each dataset are highlighted.

		Robust04				
	SQ₁	SQ₂	SQ₃	RSQ	AOL	
SQ₁		0.0923	-0.0025	0.7549	0.1250	
SQ₂	0.1758		-0.0004*	0.1461	0.0033	
SQ₃	0.0108	-0.0040*		0.0030	0.0165	
RSQ	0.5617	0.3023	0.0233		0.2076	
AOL	0.1413	-0.0174	0.0321	0.1811		
		WT10g				
	SQ₁	SQ₂	SQ₃	RSQ	AOL	
SQ₁		0.1206	0.2094	0.3103	0.1100	
SQ₂	0.3575		0.5503	0.1931	0.0900*	
SQ₃	0.3223	0.3047		0.2535	0.1795	
RSQ	0.5672	0.5312	0.3189		0.3135	
AOL	0.3731	0.2977*	0.3830	0.4126		
		Wikipedia				
	SQ₁	SQ₂	SQ₃	RSQ	AOL	
SQ₁		0.1586	0.2613	0.4768	0.2379	
SQ₂	0.3508		0.7019	0.1599	0.0598	
SQ₃	0.3897	0.5116		0.1048	0.0308*	
RSQ	0.5350	0.4867	0.3739		0.3065	
AOL	0.3603	0.3074	0.2338*	0.4194		

actual user query logs, on the computed retrievability scores when the retrieval is performed using BM25.

Our findings on diverse datasets, including news, web, and Wikipedia collections expose substantial variations in computed retrievability scores across query sets, providing an affirmative response to **RQ1** and underscoring the reproducibility challenges associated with retrievability scores. Among the popular query simulation techniques for retrievability studies, the query set generated following the method in [8] (**SQ₁**) exhibits the least disparity in computed retrievability scores. In contrast, the AOL query log generates the highest disparity. Furthermore, the utilization of a rule-based simulation technique (**RSQ**) results in the least inequality in the retrievability scores among all. These observations reveal a high sensitivity of these scores with respect to the query generation methods employed. Our research demonstrates that the selection of a query set directly impacts computed retrievability scores, regardless of the nature of the collection. This provides a conclusive response to **RQ2**. These findings emphasize the need for standardization of query set construction in retrievability studies and contribute to a deeper understanding of the nuances of accessibility within the field of information retrieval. In previous research on retrievability, various studies utilized multiple retrieval models for a single query set per collection. As part of our future work, we will explore the impact of diverse retrieval models on the retrievability scores with different query sets.

A Appendix - A POS-based Query Generation Technique

Identifying collocations in a text corpus typically involves counting co-occurring word pairs, revealing words that go beyond their individual meanings. Relying solely on the most frequent bigrams often yields uninteresting results, as many of them consist of function words, offering limited insights. To improve collocation quality, Justeson and Katz [32] introduced a simple yet effective heuristic. They apply a part-of-speech filter to candidate phrases, preserving patterns likely to represent genuine ‘phrases’ rather than random word combinations. This approach enhances the meaningfulness of the collocation identification process.

We apply this approach for query generation improving the relevance and effectiveness of the generated queries.

1. **Perform Part-of-Speech (POS) tagging:** Initially, we employ POS tagging on all the documents within the collection. This step assigns appropriate grammatical tags to each word, facilitating the subsequent identification of N-grams.
2. **Extract N-grams:** N-grams, where N represents the desired length of the word sequences, are extracted from the POS-tagged documents. In our case, we consider N to range from 1 to 4, enabling the identification of unigrams, bigrams, trigrams, and quadgrams.
3. **Select N-grams with ‘query-like’ POS tag patterns:** From the pool of extracted N-grams, we apply Justeson and Katz’s [32] recommended POS tag patterns to filter and retain N-grams that exhibit patterns resembling queries. The specific POS tag patterns for each N-gram type are provided in Table 5. Tag patterns for Quadgrams are proposed by us heuristically from our observations.

Table 5: POS Tag rules for N-gram query generation

unigram	bigram	trigram	quadgram
noun	adj noun noun noun	adj adj noun adj noun noun noun adj noun noun noun noun noun adp noun	noun verb adp noun adj noun adj noun noun adp adj noun noun noun adp noun noun verb noun noun adv adj noun noun adj noun verb noun noun adj noun noun

Subsequently, the resulting list of N-grams is sorted in descending order based on their occurrence frequencies. To ensure a manageable and relevant set of queries, we truncate the list at specific thresholds. These thresholds are determined by drawing inspiration from the frequency distribution of queries found in the AOL query set [40]. We aim to maintain a proportional ratio between the selected N-grams and the query frequencies observed in the AOL real query set, preserving a close alignment with real-world query usage patterns.

References

1. Ahmad, W.U., Chang, K.W., Wang, H.: Context attentive document ranking and query suggestion. In: Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval. p. 385–394. SIGIR’19, Association for Computing Machinery, New York, NY, USA (2019). <https://doi.org/10.1145/3331184.3331246>, <https://doi.org/10.1145/3331184.3331246>
2. Amin Abolghasemi, Suzan Verberne, A.A., Azzopardi, L.: Retrieval bias estimation using synthetically generated queries. In: Proceedings of the First Workshop on Generative Information Retrieval - GenIR@SIGIR2023 held in conjunction with SIGIR 2023. GenIR@SIGIR2023 (2023), <https://coda.io/@sigir/gen-ir/accepted-papers-17>
3. Anderson, N.: The ethics of using aol search data <https://arstechnica.com/uncategorized/2006/08/7578/>
4. Atkinson, A.B.: On the measurement of inequality. *Journal of Economic Theory* **2**(3), 244–263 (1970). [https://doi.org/https://doi.org/10.1016/0022-0531\(70\)90039-6](https://doi.org/https://doi.org/10.1016/0022-0531(70)90039-6), <https://www.sciencedirect.com/science/article/pii/S0022053170900396>
5. Azzopardi, L., Bache, R.: On the relationship between effectiveness and accessibility. In: Proceedings of the 33rd international ACM SIGIR conference on Research and development in information retrieval. pp. 889–890 (2010)
6. Azzopardi, L., Owens, C.: Search engine predilection towards news media providers. In: Proceedings of the 32nd international ACM SIGIR conference on Research and development in information retrieval. pp. 774–775 (2009)
7. Azzopardi, L., Vinay, V.: Accessibility in information retrieval. In: Advances in Information Retrieval: 30th European Conference on IR Research, ECIR 2008, Glasgow, UK, March 30–April 3, 2008. Proceedings 30. pp. 482–489. Springer (2008)
8. Azzopardi, L., Vinay, V.: Retrieval: An evaluation measure for higher order information access tasks. In: Proceedings of the 17th ACM Conference on Information and Knowledge Management. p. 561–570. CIKM ’08, Association for Computing Machinery, New York, NY, USA (2008). <https://doi.org/10.1145/1458082.1458157>, <https://doi.org/10.1145/1458082.1458157>
9. Bache, R., Azzopardi, L.: Improving access to large patent corpora. *Trans. Large Scale Data Knowl. Centered Syst.* **2**, 103–121 (2010). https://doi.org/10.1007/978-3-642-16175-9_4, https://doi.org/10.1007/978-3-642-16175-9_4
10. Barbaro, Michael; Zeller Jr, T.: A face is exposed for aol searcher no. 4417749 <https://www.nytimes.com/2006/08/09/technology/09aol.html>
11. Bashir, S.: Improving retrievability with improved cluster-based pseudo-relevance feedback selection. *Expert Syst. Appl.* **39**(8), 7495–7502 (2012). <https://doi.org/10.1016/j.eswa.2012.01.041>, <https://doi.org/10.1016/j.eswa.2012.01.041>
12. Bashir, S.: Estimating retrievability ranks of documents using document features. *Neurocomputing* **123**, 216–232 (2014)
13. Bashir, S., Khattak, A.S.: Producing efficient retrievability ranks of documents using normalized retrievability scoring function. *Journal of Intelligent Information Systems* **42**, 457–484 (2014). <https://doi.org/10.1007/s10844-013-0274-3>, <https://doi.org/10.1007/s10844-013-0274-3>

14. Bashir, S., Rauber, A.: Analyzing document retrievability in patent retrieval settings. In: Database and Expert Systems Applications: 20th International Conference, DEXA 2009, Linz, Austria, August 31–September 4, 2009. Proceedings 20. pp. 753–760. Springer (2009)
15. Bashir, S., Rauber, A.: Identification of low/high retrievable patents using content-based features. In: Proceedings of the 2nd international workshop on Patent information retrieval. pp. 9–16 (2009)
16. Bashir, S., Rauber, A.: Improving retrievability of patents with cluster-based pseudo-relevance feedback documents selection. In: Proceedings of the 18th ACM conference on Information and knowledge management. pp. 1863–1866 (2009)
17. Bashir, S., Rauber, A.: Improving retrievability and recall by automatic corpus partitioning. Transactions on large-scale data-and knowledge-centered systems II pp. 122–140 (2010)
18. Bashir, S., Rauber, A.: Improving retrievability of patents in prior-art search. In: Advances in Information Retrieval: 32nd European Conference on IR Research, ECIR 2010, Milton Keynes, UK, March 28-31, 2010. Proceedings 32. pp. 457–470. Springer (2010)
19. Bashir, S., Rauber, A.: On the relationship between query characteristics and ir functions retrieval bias. Journal of the American Society for Information Science and Technology **62**(8), 1515–1532 (2011)
20. Bashir, S., Rauber, A.: Automatic ranking of retrieval models using retrievability measure. Knowledge and information systems **41**, 189–221 (2014)
21. Bashir, S., Rauber, A.: Retrieval models versus retrievability. Current Challenges in Patent Information Retrieval pp. 185–212 (2017)
22. Bashir, S., Rauber, A.: Retrieval Models Versus Retrievability, pp. 185–212. Springer Berlin Heidelberg, Berlin, Heidelberg (2017). https://doi.org/10.1007/978-3-662-53817-3_7, https://doi.org/10.1007/978-3-662-53817-3_7
23. Boratto, L., Faralli, S., Marras, M., Stilo, G. (eds.): Advances in Bias and Fairness in Information Retrieval. Springer Nature Switzerland (2023). <https://doi.org/10.1007/978-3-031-37249-0>, <https://doi.org/10.1007/978-3-031-37249-0>
24. Callan, J., Connell, M.: Query-based sampling of text databases. ACM Transactions on Information Systems (TOIS) **19**(2), 97–130 (2001)
25. Dehghani, M., Zamani, H., Severyn, A., Kamps, J., Croft, W.B.: Neural ranking models with weak supervision. In: Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval. p. 65–74. SIGIR '17, Association for Computing Machinery, New York, NY, USA (2017). <https://doi.org/10.1145/3077136.3080832>, <https://doi.org/10.1145/3077136.3080832>
26. Ekstrand, M.D., Das, A., Burke, R., Diaz, F.: Fairness in information access systems. Foundations and Trends® in Information Retrieval **16**(1-2), 1–177 (2022). <https://doi.org/10.1561/15000000079>, <http://dx.doi.org/10.1561/15000000079>
27. Gini, C.: On the measure of concentration with special reference to income and statistics. Colorado College Publication, General Series **208**(1), 73–79 (1936)
28. Hafner, K.: Tempting data, privacy concerns; researchers yearn to use aol logs, but they hesitate <https://www.nytimes.com/2006/08/23/technology/23search.html>
29. Hawking, D.: Overview of the TREC-9 web track. In: Voorhees, E.M., Harman, D.K. (eds.) Proceedings of The Ninth Text REtrieval Conference, TREC 2000,

- Gaithersburg, Maryland, USA, November 13-16, 2000. NIST Special Publication, vol. 500-249. National Institute of Standards and Technology (NIST) (2000), <http://trec.nist.gov/pubs/trec9/papers/web9.pdf>
30. Johnston, J.: H. Theil. Economics and Information Theory. *The Economic Journal* **79**(315), 601–602 (09 1969). <https://doi.org/10.2307/2230396>, <https://doi.org/10.2307/2230396>
 31. Jordan, C., Watters, C., Gao, Q.: Using controlled query generation to evaluate blind relevance feedback algorithms. In: Proceedings of the 6th ACM/IEEE-CS joint conference on Digital libraries. pp. 286–295 (2006)
 32. Justeson, J.S., Katz, S.M.: Co-occurrences of antonymous adjectives and their contexts. *Computational Linguistics* **17**(1), 1–20 (1991), <https://aclanthology.org/J91-1001>
 33. Kang, Y.M., Liu, W., Zhou, Y.: Queryblazer: Efficient query autocompletion framework. In: Proceedings of the 14th ACM International Conference on Web Search and Data Mining. p. 1020–1028. WSDM '21, Association for Computing Machinery (2021). <https://doi.org/10.1145/3437963.3441725>
 34. Ma, Z., Dou, Z., Bian, G., Wen, J.R.: Pstie: Time information enhanced personalized search. In: Proceedings of the 29th ACM International Conference on Information & Knowledge Management. p. 1075–1084. CIKM '20, Association for Computing Machinery (2020). <https://doi.org/10.1145/3340531.3411877>
 35. MacAvaney, S., Macdonald, C., Ounis, I.: Reproducing personalised session search over the aol query log. In: Advances in Information Retrieval: 44th European Conference on IR Research, ECIR 2022, Stavanger, Norway, April 10–14, 2022, Proceedings, Part I. p. 627–640. Springer-Verlag, Berlin, Heidelberg (2022). https://doi.org/10.1007/978-3-030-99736-6_42
 36. McLellan, C.: The relationship between retrievability bias and retrieval performance. Ph.D. thesis, University of Glasgow, UK (2019), <https://ethos.bl.uk/OrderDetails.do?uin=uk.bl.ethos.775857>
 37. Nogueira, R., Lin, J.: From doc2query to docttttquery. In: Online preprint 6 (2019), <https://github.com/castorini/docTTTTquery>
 38. Noor, S., Bashir, S.: Evaluating bias in retrieval systems for recall oriented documents retrieval. *International Arab Journal of Information Technology (IAJIT)* **12**(1) (2015)
 39. Palma, J.G.: Homogeneous middles vs. heterogeneous tails, and the end of the 'inverted-u': the share of the rich is what it's all about. Cambridge working papers in economics, Faculty of Economics, University of Cambridge (2011), <https://EconPapers.repec.org/RePEc:cam:camdae:1111>
 40. Pass, G., Chowdhury, A., Torgeson, C.: A picture of search. In: Proceedings of the 1st International Conference on Scalable Information Systems. p. 1–es. InfoScale '06, Association for Computing Machinery (2006). <https://doi.org/10.1145/1146847.1146848>
 41. Pickens, J., Cooper, M., Golovchinsky, G.: Reverted indexing for feedback and expansion. In: Proceedings of the 19th ACM international conference on Information and knowledge management. pp. 1049–1058 (2010)
 42. Roy, D., Carevic, Z., Mayr, P.: Studying retrievability of publications and datasets in an integrated retrieval system. In: Proceedings of the 22nd ACM/IEEE Joint Conference on Digital Libraries. JCDL '22, Association for Computing Machinery (2022). <https://doi.org/10.1145/3529372.3530931>
 43. Roy, D., Carevic, Z., Mayr, P.: Retrievability in an integrated retrieval system: an extended study. *International Journal on Digital Libraries* (Apr

- 2023). <https://doi.org/10.1007/s00799-023-00363-4>, <https://doi.org/10.1007/s00799-023-00363-4>
44. Traub, M.C., Samar, T., van Ossenbruggen, J., Hardman, L.: Impact of crowd-sourcing ocr improvements on retrievability bias. In: Proceedings of the 18th ACM/IEEE on Joint Conference on Digital Libraries. p. 29–36. JCDL '18, Association for Computing Machinery (2018). <https://doi.org/10.1145/3197026.3197046>
 45. Traub, M.C., Samar, T., Van Ossenbruggen, J., He, J., de Vries, A., Hardman, L.: Querylog-based assessment of retrievability bias in a large newspaper corpus. In: 2016 IEEE/ACM Joint Conference on Digital Libraries (JCDL). pp. 7–16. IEEE (2016)
 46. Voorhees, E.M.: Overview of the TREC 2004 robust track. In: Voorhees, E.M., Buckland, L.P. (eds.) Proceedings of the Thirteenth Text REtrieval Conference, TREC 2004, Gaithersburg, Maryland, USA, November 16-19, 2004. NIST Special Publication, vol. 500-261. National Institute of Standards and Technology (NIST) (2004), <http://trec.nist.gov/pubs/trec13/papers/ROBUST.OVERVIEW.pdf>
 47. Wilkie, C., Azzopardi, L.: An initial investigation on the relationship between usage and findability. In: Advances in Information Retrieval: 35th European Conference on IR Research, ECIR 2013, Moscow, Russia, March 24-27, 2013. Proceedings 35. pp. 808–811. Springer (2013)
 48. Wilkie, C., Azzopardi, L.: Relating retrievability, performance and length. In: Proceedings of the 36th international ACM SIGIR conference on Research and development in information retrieval. pp. 937–940 (2013)
 49. Wilkie, C., Azzopardi, L.: Best and fairest: An empirical analysis of retrieval system bias. In: de Rijke, M., Kenter, T., de Vries, A.P., Zhai, C., de Jong, F., Radinsky, K., Hofmann, K. (eds.) Advances in Information Retrieval. pp. 13–25. Springer International Publishing, Cham (2014)
 50. Wilkie, C., Azzopardi, L.: Efficiently estimating retrievability bias. In: de Rijke, M., Kenter, T., de Vries, A.P., Zhai, C., de Jong, F., Radinsky, K., Hofmann, K. (eds.) Advances in Information Retrieval. pp. 720–726. Springer International Publishing, Cham (2014)
 51. Wilkie, C., Azzopardi, L.: A retrievability analysis: Exploring the relationship between retrieval bias and retrieval performance. In: Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management. pp. 81–90 (2014)
 52. Wilkie, C., Azzopardi, L.: Query length, retrievability bias and performance. In: Proceedings of the 24th ACM International on Conference on Information and Knowledge Management. p. 1787–1790. CIKM '15, Association for Computing Machinery, New York, NY, USA (2015). <https://doi.org/10.1145/2806416.2806604>, <https://doi.org/10.1145/2806416.2806604>
 53. Wilkie, C., Azzopardi, L.: Retrievability and retrieval bias: A comparison of inequality measures. In: Advances in Information Retrieval: 37th European Conference on IR Research, ECIR 2015, Vienna, Austria, March 29-April 2, 2015. Proceedings 37. pp. 209–214. Springer (2015)
 54. Zheng, L., Cox, I.J.: Document-oriented pruning of the inverted index in information retrieval systems. In: 2009 International Conference on Advanced Information Networking and Applications Workshops. pp. 697–702. IEEE (2009)
 55. Zhu, Y., Nie, J.Y., Dou, Z., Ma, Z., Zhang, X., Du, P., Zuo, X., Jiang, H.: Contrastive learning of user behavior sequence for context-aware document ranking. In: Proceedings of the 30th ACM International Conference on Information & Knowledge Management. p. 2780–2791. CIKM '21, Association for Computing Machinery (2021). <https://doi.org/10.1145/3459637.3482243>