

Findability: A Novel Measure of Information Accessibility

Aman Sinha

Indian Institute of Science Education
and Research, Kolkata, India
as18ms065@iiserkol.ac.in

Priyanshu Raj Mall

Indian Institute of Science Education
and Research, Kolkata, India
prm18ms118@iiserkol.ac.in

Dwaipayan Roy

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

ABSTRACT

The overwhelming volume of data generated and indexed by search engines poses a significant challenge in retrieving documents from the index efficiently and effectively. Even with a well-crafted query, several relevant documents often get buried among a multitude of competing documents, resulting in reduced accessibility or “findability” of the desired document. Consequently, it is crucial to develop a robust methodology for assessing this dimension of Information Retrieval (IR) system performance. While previous studies have focused on measuring document accessibility disregarding user queries and document relevance, there exists no metric to quantify the findability of a document within a given IR system without resorting to manual labor. This paper aims to address this gap by defining and deriving a metric to evaluate the findability of documents as perceived by end-users. Through experiments, we demonstrate the varying impact of different retrieval models and collections on the findability of documents. Furthermore, we establish the findability measure as an independent metric distinct from retrievability, an accessibility measure introduced in prior literature.

CCS CONCEPTS

• Information systems → Retrieval effectiveness.

KEYWORDS

Accessibility, Bias, Evaluation, Empirical Study

ACM Reference Format:

Aman Sinha, Priyanshu Raj Mall, and Dwaipayan Roy. 2023. Findability: A Novel Measure of Information Accessibility. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management (CIKM '23)*, October 21–25, 2023, Birmingham, United Kingdom. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3583780.3615256>

1 INTRODUCTION

In the present digital era, where information is abundant and easily available, the challenge lies not in obtaining information but in finding the right information when we need it. Whether searching for a specific record in a vast database, locating a particular product on an e-commerce website, or identifying important content on the internet, finding relevant information has become a crucial aspect

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

To appear in the proceedings of *CIKM '23*, October 21–25, 2023, Birmingham, UK
© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.
<https://doi.org/10.1145/3583780.3615256>

of our daily lives. Active research in the domain of information retrieval continues to evolve and address the various challenges emerging from the complex information need of the users. One of the most explored areas of research is in exploratory search where the purpose is to discover new information, gain a deeper understanding of a subject, or explore different perspectives; in this scenario, the users are likely to be engaged in an open-ended search process to discover and learn about a particular topic or subject [19]. In these searches, users have a general idea of what they are interested in but may not have specific details or keywords in mind. However, various biases in the retrieval systems may also play an adverse role in the selection of documents among the top positions after a retrieval [16, 20], making a significant amount of documents non-accessible to the user [4].

On the other hand, known-item searches occur when a user is looking for a specific item or a particular piece of information. In this type of search, the user already has a clear idea of what they are searching for and often possesses some specific details about the item (text document in our context), such as a title, author, or specific keywords. The goal is to *find* an exact document quickly and accurately. However, finding a particular document can be challenging due to factors such as the large volume of available information, document variability, and biases in retrieval models.

In this paper, we propose a novel measurement scheme to quantify the overall findability of items in a document collection. Considering a collection of entities or items, if we assume that each entity will be relevant to a subset of all possible queries, the *findability* of an entity is the quantification of the ease with which the item can be found by the user by issuing queries from that subset only. To evaluate the practicality and effectiveness of the findability measure, we conduct an empirical study on benchmark retrieval datasets. By conducting experiments with various retrieval models and datasets of varying sizes, we observe the diverse effect and implications on the findability within an IR system.

2 BACKGROUND AND RELATED WORK

Research on the accessibility of documents in the domain of IR has focused on various aspects, such as crawlability [18], findability [13], discoverability [7], retrievability [4] and navigability [8]. The accessibility of information within a collection, as addressed in existing literature, encompasses two distinct perspectives: from the system’s standpoint [4], and from the user’s standpoint [12, 13]. Retrievability [4] measures gauge the ease with which a document can be retrieved utilizing a specific IR system. In contrast, the concept of findability [13] seeks to assess the ease with which a user, utilizing the IR system, can locate a particular document¹.

¹While the FAIR principle [22] distinguishes between accessibility and findability, in this study, we use the term findability to signify a method for gaining access to the documents within a collection.

The retrievability measure focuses exclusively on the system-oriented aspect of document retrieval, neglecting user intent during a search. Further, the retrievability measure provides only a general approximation of the likelihood of a document being retrieved, regardless of which query is posed to the IR system. This approximation merely attempts to gauge the accessibility of documents within the collection facilitated by the IR system, disregarding the user’s perspective. When users interact with a search engine, they express their specific information need by submitting a query, reflecting their intent and purpose for utilizing the IR system. Therefore, to accurately estimate document accessibility within the collection, it is crucial to account for both the user and their query together with the IR system during the search process.

The ability to access relevant documents in response to a user’s query lies at the heart of document findability from the user’s standpoint. For instance, presenting a document related to "Java" in response to a user’s search for "Python" does not contribute to the findability of "Java" because the document lacks relevance to the query posed. In this context, the findability of a document refers to its capacity to be located solely for queries whose intent is satisfied or addressed by that particular document. In other words, a document is considered *found* when the user is satisfied to encounter it in the search results as per the query they entered into the IR system to find that document. A similar approach was followed in [12] where the authors introduced *Page Hunt*, a game specifically designed to collect web search log data. The game involves presenting participants with webpages and tasking them with finding them using the provided search interface.

Previous studies [3, 4] have introduced retrievability measures to estimate the ease of document retrieval using search engines. Additionally, other researchers [6, 14, 24] have explored something like findability from the perspective of browsing and navigability, evaluating how easily users can navigate websites. Authors of [5] say that successful validation of findability measures could enable the development of tools to assist Information Architects in analyzing websites, offering insights into the findability and utilization of content, as well as identifying features (such as terms, links, etc.) that contribute to the ease or difficulty of locating specific pages.

3 FINDABILITY - A MEASURE OF ACCESSIBILITY

Consider an IR system that employs model R to retrieve relevant documents in response to user queries from a document collection D . Let Q_d represent the set of all possible queries for which document d (where $d \in D$) is deemed relevant. We refer to these queries as “relevant queries” specifically for that particular document. The *findability* of a document is then defined as the expectation of the likelihood of a user finding that document for every query q ($q \in Q_d$). Mathematically, the findability measure is formulated as:

$$f(d) = \frac{1}{|Q_d|} \sum_{q \in Q_d} \xi(p_{dq}, c) \quad (1)$$

In Equation 1, p_{dq} is the rank of document d in the search result against query q . The function $\xi(p_{dq}, c)$ is a generalized convenience

function that captures users’ willingness to explore the search results up to rank p_{dq} , while c denotes the threshold rank at which users cease examining the ranked list.

The function $\xi(p_{dq}, c)$ is subject to two boundary constraints. The first constraint ensures that when document d appears at the top rank ($p_{dq} = 1$), it represents the most convenient and optimal scenario for the user; $\xi(1, c)$ is set to 1 to reflect this favorable situation. On the other hand, users typically do not continue indefinitely exploring the ranked list of search results; they stop investigating (denoted by c in Equation 1) at some point. In the worst-case scenario, if a document appears in the search results after the user has stopped investigating, it means the user does not find that document. Thus, for ranks p_{dq} greater than c , we set $\xi(p_{dq}, c) = 0$. With these two constraints, the function $\xi(\cdot)$ is bounded within the range of $[0, 1]$, making it a suitable measure for interpretation.

To define the convenience function $\xi(\cdot)$, we employ the concept of Click Through Rate (CTR): the net percentage of clicks that a document at a certain rank gets out of total clicks by users to open a search result document in the ranked list. The CTR on a search engine could be taken as a representation of the user tendency or user effort it takes to investigate a certain rank in the results. Notably, analyses conducted by Semrush Inc. and Backlinko² based on 4 million Google search results offer valuable insights into CTR for top ranks in the context of web searches. These findings indicate that a mere 0.63% users click on results beyond rank position 10, suggesting that a majority of users discontinue exploring the ranked list after a certain rank threshold. The observations from CTR of users lead us to propose the following two forms for the convenience function $\xi(p_{dq}, c)$ in the context of findability measure:

- **Exponential decay of Convenience:** Considering an exponential decay with a decay rate of approximately one-third, the convenience function can be defined as:

$$\xi(p_{dq}, c) = \begin{cases} e^{-(p_{dq}-1)/3} & \text{if } p_{dq} \leq c \\ 0 & \text{if } p_{dq} > c \end{cases} \quad (2)$$

- **Inverse law of Convenience:** An alternative approach to incorporate the decaying effect is by considering the inverse of the document rank, which can be expressed as follows:

$$\xi(p_{dq}, c) = \begin{cases} \frac{1}{p_{dq}} & \text{if } p_{dq} \leq c \\ 0 & \text{if } p_{dq} > c \end{cases} \quad (3)$$

3.1 Estimating Document Findability

In order to estimate *findability* scores of documents in an operational setting, one crucial requirement is the creation of a relevant query set for every document. Ideally, we would require a comprehensive list of all the queries for which a document could be considered relevant. However, generating such an exhaustive list manually would be an impractically labor-intensive task, even for a moderately-sized collection of documents: human experts need to read the documents and submit search queries that are deemed relevant to the respective documents. Avoiding the involvement of human efforts, a known-item query generation strategy can be employed as a proxy for automatically generating a smaller but representative sample of relevant queries for each document.

²<https://backlinko.com/google-ctr-stats>

Table 1: Statistics of datasets used for experimentation.

Dataset	# documents	Collection type	# terms	# queries
TREC Robust	528,155	News	1,502,031	10,230,070
WT10g	1,692,096	Web	9,674,707	26,041,327
MS MARCO passage	8,841,823	Web excerpts	1,410,558	19,839,452

3.2 Mean Findability and Findability Bias

When evaluating document findability within a fixed collection, two aspects of access provided by different retrieval models can be assessed. Firstly, the *mean of findability* scores for all documents in the collection reflects the overall effectiveness of the retrieval model in retrieving the correct document at the top. This provides a measure of the aggregate performance of the retrieval model in delivering relevant documents. Secondly, the Gini coefficient [10] can be employed to quantify the disparity of access imposed by the retrieval model across the collection. In this context, the Gini coefficient represents the imbalance of access among the documents within the collection itself. It captures the extent to which certain documents are favored over others in terms of findability. By considering both the mean findability, denoted as $\langle f \rangle$, and the findability bias, represented by the Gini coefficient G , a comprehensive assessment of document findability offered by a retrieval model can be obtained. This dual approach provides a holistic understanding of how effectively and fairly the retrieval model enables access to the documents in the collection.

The applicability of Gini coefficient as a measure of inequality in the context of accessibility in IR has been employed in earlier studies [4, 21]. This measure, borrowed from economics and social sciences, provides a quantitative way to assess the level of inequality in access to information. Lorenz curve [9, 11], the graphical representation of Gini coefficient is used to visualize the deviation of wealth distribution from equality. In the context of findability, the Gini coefficient can be computed as follows:

$$G = \frac{\sum_{i=1}^N (2i - N - 1) \cdot f(d_i)}{N \sum_{j=1}^N f(d_j)} \quad (4)$$

where $f(d_i)$ and $f(d_j)$ are i^{th} and j^{th} documents when documents are ordered in ascending order by their findability scores; N is the total number of documents in the collection. Gini coefficient G ranges from 0 to 1, where $G = 0$ means no bias (i.e., an equal distribution which implies $f(d)$ is equal for all documents) and $G = 1$ indicates maximum bias (i.e., all documents have $f(d) = 0$ except one document) - implying that only one document consistently appears at the top ranks for its relevant queries, while the rest of the documents are never found within the top ranks and remain hidden among other documents (that take up the top c positions).

4 EMPIRICAL ANALYSIS

In this section, we present an experiment that showcases a practical use-case scenario for the findability measure. We evaluate the Mean Findability and Findability Bias of three standard retrieval models across three different benchmark collections to determine which retrieval model offers the best accessibility in diverse collection types. This evaluation allows us to assess the effectiveness of each retrieval model in retrieving the correct documents at the top ranks

Table 2: Gini coefficient G and Mean Findability $\langle f \rangle$ for Findability $f(d)$

		Robust04	WT10g	MS MARCO
LM-Dir	G	0.1587	0.2847	0.3774
	$\langle f \rangle$	0.6327	0.5209	0.5173
BM25	G	0.1456	0.2503	0.3116
	$\langle f \rangle$	0.6640	0.5985	0.5895
DFR-PL2	G	0.1424	0.2497	0.3007
	$\langle f \rangle$	0.6672	0.6133	0.5888

Table 3: Correlation between Findability $f(d)$ and Retrievability $r(d)$ when both $f(d)$ and $r(d)$ are estimated using query sets generated from their respective datasets (Self query set) as well as the known-item search query set (Known-item query), where all are statistically significant correlation with $p < 0.05$.

	Self query set		Known-item query set	
	Pearson's r	Kendall's τ	Pearson's r	Kendall's τ
Robust04	-0.0944	-0.0518	-0.1292	-0.1053
WT10g	-0.0088	0.0084	-0.0256	-0.0287
MS MARCO	0.0115	0.0307	0.0388	0.0269

and the degree of bias in document access within each collection. Further, we investigate the relationship between findability and retrievability [3]. By comparing findability scores with retrievability scores, we uncover that findability scores are independent and distinct from retrievability scores.

4.1 Datasets and Retrieval Models

We evaluate the proposed findability metric using three benchmark datasets: TREC Robust, WT10g, and MS-Marco passage. The statistics of the datasets are presented in Table 1. These datasets are commonly used in information retrieval research and provide diverse document collections for our analysis. In this study, we investigate the findability of documents using three different retrieval models, particularly BM25 [15, 17], LM-Dir [23], and DFR-PL2 [1].

4.2 Query Generation Method

As discussed in Section 3.1, the estimation of findability requires a set of relevant queries for each document in the collection. In this study, we employ a known-item query generation method, which was introduced in a previous work [2]. This method has demonstrated its effectiveness in generating query scores that are comparable to manually generated known-item queries. Following the *popular+discriminative* selection strategy method [2], we generate a set of known-item search queries with the average query length k and mixing parameter λ (as defined in [2]) set to 4 and 0 respectively.

To ensure a comprehensive assessment, the generated number of relevant queries is set to 10% of the total number of distinct terms in the document with an upper cap of 50 to maintain computational tractability.




Figure 1: Lorenz Curves for three retrieval models for each of the three collections. Note that, some of the curves are closely situated, potentially making them indistinguishable in the plots, particularly for the Robust04 and WT10g collections. To highlight the contrast between retrievability with findability, the scaled curves for retrievability are also presented in red.

4.3 Experimentation

The findability measure, defined in Equation 1, includes a parameter c that represents the maximum rank tolerance of users. While the optimal value of c can vary depending on a specific task and user preferences, we have chosen a value of $c = 100$ for this particular study³. Based on our initial study of CTR data, we found that the inverse law of convenience provides a better fit. Therefore, we have opted to utilize the inverse law form (Equation 3) for estimating the findability scores.

We report the Gini coefficient G and the mean findability $\langle f \rangle$ of the collections in Table 2. From the table, we can observe that as the Gini coefficient increases, there is a noticeable decrease in mean findability. This implies that across the three examined retrieval models, improving the overall findability of documents results in a concurrent decrease in findability bias. Consequently, retrieval models that increase the findability of documents tend to improve accessibility to the collection as a whole by enhancing the findability of the majority of documents. The inequality is graphically presented as Lorenz curve in Figure 1. In the figure 1, the curve for retrievability employing BM25 model is also presented in red to showcase the disparity with the findability bias.

Moreover, it is observed that mean findability decreases and findability bias increases with collection size. This aligns with our intuition that a larger number of documents results in heightened competition among them for higher rankings, thereby diminishing findability. Additionally, it seems that the bias imposed by the retrieval model on the collection becomes more pronounced as the collection size increase.

Among the three retrieval models, LM-Dir yields lower findability for the collection’s documents compared to the other two models. While BM25 and PL2 exhibit similar mean findability, the Gini coefficient for PL2 is reported to be lower indicating that PL2 performs better overall in terms of the findability aspect of accessibility-based model performance.

³We experimented with c from 10 to 100 varied in steps of 10; considering the space limitation, we are reporting the results for $c = 100$ in this paper.

4.4 Comparing with Retrievability

The distinction between retrievability and findability may initially appear subtle. Still, it is essential to acknowledge the substantial conceptual distinctions between them that emerge when considering user behavior and user-centric factors. To clarify the distinction, we perform a retrievability analysis using BM25 retrieval model for all three collections, utilizing a standard query set commonly used for computing retrievability scores, as mentioned in the original work [4]. Further, we compute Pearson’s and Kendall’s rank correlation coefficient. As presented in Table 3, the obtained correlations reveal an almost negligible association between the two measures. Even when utilizing the same queries for retrievability analysis as those employed for findability evaluation (known-item queries), Table 3’s third column reveals a persistent lack of relationship between findability and retrievability. These correlation results serve to establish findability as an independent measure of accessibility that was not previously encompassed by the retrievability measure.

The findability score provides a uniform interpretation and a constant range, akin to a coefficient, unlike retrievability which encounters comparability challenges across diverse studies due to variations in query sets. Moreover, findability measure is well-suited for analyzing the findability of individual documents, whereas a single retrievability score alone may not adequately represent retrievability of a document without additional context or information.

5 CONCLUSION AND FUTURE WORK

This paper introduces a novel metric called ‘findability’ for measuring document accessibility. Our study demonstrates that all three retrieval models exhibit comparable behavior in regard to findability. Furthermore, we compare findability with retrievability, another existing metric for document accessibility. Future research will investigate the use of the findability measure in fine-tuning IR system parameters in situations where relevance judgments are unavailable. Additionally, exploring the correlation between improved overall findability and enhanced user experience is an area of interest for further investigation.

Acknowledgement: We would like to thank the anonymous reviewer for their valuable and encouraging feedback.

REFERENCES

- [1] Gianni Amati and Cornelis Joost Van Rijsbergen. 2002. Probabilistic Models of Information Retrieval Based on Measuring the Divergence from Randomness. *ACM Trans. Inf. Syst.* 20, 4 (oct 2002), 357–389. <https://doi.org/10.1145/582415.582416>
- [2] Leif Azzopardi, Maarten de Rijke, and Krisztian Balog. 2007. Building Simulated Queries for Known-Item Topics: An Analysis Using Six European Languages. In *Proceedings of the 30th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval* (Amsterdam, The Netherlands) (*SIGIR '07*). Association for Computing Machinery, New York, NY, USA, 455–462. <https://doi.org/10.1145/1277741.1277820>
- [3] Leif Azzopardi and Vishwa Vinay. 2008. Accessibility in Information Retrieval. In *Proceedings of the IR Research, 30th European Conference on Advances in Information Retrieval* (Glasgow, UK) (*ECIR'08*). Springer-Verlag, Berlin, Heidelberg, 482–489.
- [4] Leif Azzopardi and Vishwa Vinay. 2008. Retrievability: An Evaluation Measure for Higher Order Information Access Tasks. In *Proceedings of the 17th ACM Conference on Information and Knowledge Management* (Napa Valley, California, USA) (*CIKM '08*). Association for Computing Machinery, New York, NY, USA, 561–570. <https://doi.org/10.1145/1458082.1458157>
- [5] Leif Azzopardi, Colin Wilkie, and Tony Russell-Rose. 2013. Towards Measures and Models of Findability. In *SIGIR 2013 Workshop on Modeling User Behavior for Information Retrieval Evaluation (MUBE 2013)*. Citeseer, 3.
- [6] Ed H. Chi, Peter Pirolli, and James Pitkow. 2000. The Scent of a Site: A System for Analyzing and Predicting Information Scent, Usage, and Usability of a Web Site. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (The Hague, The Netherlands) (*CHI '00*). Association for Computing Machinery, New York, NY, USA, 161–168. <https://doi.org/10.1145/332040.332423>
- [7] Anirban Dasgupta, Arpita Ghosh, Ravi Kumar, Christopher Olston, Sandeep Pandey, and Andrew Tomkins. 2007. The Discoverability of the Web. In *Proceedings of the 16th International Conference on World Wide Web* (Banff, Alberta, Canada) (*WWW '07*). Association for Computing Machinery, New York, NY, USA, 421–430. <https://doi.org/10.1145/1242572.1242630>
- [8] Xiao Fang, Paul Jen-Hwa Hu, Michael Chau, Han-Fen Hu, Zhuo Yang, and Olivia R. Liu Sheng. 2012. A Data-Driven Approach to Measure Web Site Navigability. *Journal of Management Information Systems* 29, 2 (2012), 173–212. <http://www.jstor.org/stable/41713888>
- [9] Joseph L. Gastwirth. 1972. The Estimation of the Lorenz Curve and Gini Index. *The Review of Economics and Statistics* 54, 3 (1972), 306–316. <http://www.jstor.org/stable/1937992>
- [10] Corrado Gini. 1936. On the measure of concentration with special reference to income and statistics. *Colorado College Publication, General Series* 208, 1 (1936), 73–79.
- [11] Max O Lorenz. 1905. Methods of measuring the concentration of wealth. *Publications of the American statistical association* 9, 70 (1905), 209–219. <http://www.jstor.org/stable/2276207>
- [12] Hao Ma, Raman Chandrasekar, Chris Quirk, and Abhishek Gupta. 2009. Improving Search Engines Using Human Computation Games. In *Proceedings of the 18th ACM Conference on Information and Knowledge Management* (Hong Kong, China) (*CIKM '09*). Association for Computing Machinery, New York, NY, USA, 275–284. <https://doi.org/10.1145/1645953.1645990>
- [13] Peter Morville. 2005. *Ambient Findability: What We Find Changes Who We Become*. O'Reilly Media, Inc.
- [14] Shashank Pandit and Christopher Olston. 2007. Navigationaided Retrieval. In *Proceedings of the 16th International Conference on World Wide Web* (Banff, Alberta, Canada) (*WWW '07*). Association for Computing Machinery, New York, NY, USA, 391–400. <https://doi.org/10.1145/1242572.1242626>
- [15] Stephen E. Robertson, Steve Walker, Steve Jones, Micheline M. Hancock-Beaulieu, Michael Gatford, and Michel M. H. Beaulieu. 1994. Okapi at TREC-3. In *Proceedings of the Third Text REtrieval Conference (TREC-3)*.
- [16] Thaer Samar, Myriam C. Traub, Jacco van Ossensbruggen, Lynda Hardman, and Arjen P. de Vries. 2018. Quantifying retrieval bias in Web archive search. *Int. J. Digit. Libr.* 19, 1 (2018), 57–75. <https://doi.org/10.1007/s00799-017-0215-9>
- [17] K. Sparck Jones, S. Walker, and S.E. Robertson. 2000. A probabilistic model of information retrieval: development and comparative experiments: Part 1. *Information Processing & Management* 36, 6 (2000), 779–808. [https://doi.org/10.1016/S0306-4573\(00\)00015-7](https://doi.org/10.1016/S0306-4573(00)00015-7)
- [18] Trystan Upstill, Nick Craswell, and David Hawking. 2002. Buying bestsellers online: A case study in Search & Searchability. In *ADCS 2002, Proceedings of the Seventh Australasian Document Computing Symposium, Sydney, Australia, December 16, 2002*. <http://www.cie.ict.csiro.au/adcs2002/papers/upstill-craswell-hawking.pdf>
- [19] Ryen W. White and Resa A. Roth. 2009. *Exploratory Search: Beyond the Query-Response Paradigm*. Morgan & Claypool Publishers. <https://doi.org/10.2200/S00174ED1V01Y200901ICR003>
- [20] Colin Wilkie and Leif Azzopardi. 2014. Best and Fairest: An Empirical Analysis of Retrieval System Bias. In *Advances in Information Retrieval*. Springer International Publishing, Cham, 13–25. https://doi.org/10.1007/978-3-319-06028-6_2
- [21] Colin Wilkie and Leif Azzopardi. 2015. 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 (Lecture Notes in Computer Science, Vol. 9022)*. 209–214. https://doi.org/10.1007/978-3-319-16354-3_22
- [22] Mark D. Wilkinson, Michel Dumontier, IJsbrand Jan Aalbersberg, Gabrielle Appleton, Myles Axton, Arie Baak, Niklas Blomberg, Jan-Willem Boiten, Luiz Bonino da Silva Santos, Philip E. Bourne, Jildau Bouwman, Anthony J. Brookes, Tim Clark, Mercè Crosas, Ingrid Dillo, Olivier Dumon, Scott Edmunds, Chris T. Evelo, Richard Finkers, Alejandra Gonzalez-Beltran, Alasdair J.G. Gray, Paul Groth, Carole Goble, Jeffrey S. Grethe, Jaap Heringa, Peter A.C. 't Hoen, Rob Hooft, Tobias Kuhn, Ruben Kok, Joost Kok, Scott J. Lusher, Maryann E. Martone, Albert Mons, Abel L. Packer, Bengt Persson, Philippe Rocca-Serra, Marco Roos, Rene van Schaik, Susanna-Assunta Sansone, Erik Schultes, Thierry Sengstag, Ted Slater, George Strawn, Morris A. Swertz, Mark Thompson, Johan van der Lei, Erik van Mulligen, Jan Velterop, Andra Waagmeester, Peter Wittenburg, Katherine Wolstencroft, Jun Zhao, and Barend Mons. 2016. The FAIR Guiding Principles for scientific data management and stewardship. *Scientific Data* 3, 1 (March 2016). <https://doi.org/10.1038/sdata.2016.18>
- [23] ChengXiang Zhai and John D. Lafferty. 2001. A Study of Smoothing Methods for Language Models Applied to Ad Hoc Information Retrieval. In *SIGIR 2001: Proceedings of the 24th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, September 9-13, 2001, New Orleans, Louisiana, USA*. ACM, 334–342. <https://doi.org/10.1145/383952.384019>
- [24] Yuming Zhou, Hareton Leung, and P. Winoto. 2007. MNav: A Markov Model-Based Web Site Navigability Measure. *IEEE Transactions on Software Engineering* 33, 12 (Dec. 2007), 869–890. <https://doi.org/10.1109/tse.2007.70743>