



Universiteit
Leiden
The Netherlands

TripJudge: a relevance judgement test collection for TripClick health retrieval

Althammer, S.; Hofstätter, S.; Verberne, S.; Hanbury, A.

Citation

Althammer, S., Hofstätter, S., Verberne, S., & Hanbury, A. (2022). TripJudge: a relevance judgement test collection for TripClick health retrieval. *Proceedings Of The 31St Acm International Conference On Information & Knowledge Management (Cikm '22)*, 3801–3805. doi:10.1145/3511808.3557714

Version: Publisher's Version
License: [Creative Commons CC BY 4.0 license](https://creativecommons.org/licenses/by/4.0/)
Downloaded from: <https://hdl.handle.net/1887/3736325>

Note: To cite this publication please use the final published version (if applicable).



TripJudge: A Relevance Judgement Test Collection for TripClick Health Retrieval

Sophia Althammer
sophia.althammer@tuwien.ac.at
Vienna University of Technology
Vienna, Austria

Suzan Verberne
s.verberne@liacs.leidenuniv.nl
Leiden University
Leiden, Netherlands

Sebastian Hofstätter
sebastian.hofstaetter@tuwien.ac.at
Vienna University of Technology
Vienna, Austria

Allan Hanbury
hanbury@ifs.tuwien.ac.at
Vienna University of Technology
Vienna, Austria

ABSTRACT

Robust test collections are crucial for Information Retrieval research. Recently there is a growing interest in evaluating retrieval systems for domain-specific retrieval tasks, however these tasks often lack a reliable test collection with human-annotated relevance assessments following the Cranfield paradigm. In the medical domain, the TripClick collection was recently proposed, which contains click log data from the Trip search engine and includes two click-based test sets. However the clicks are biased to the retrieval model used, which remains unknown, and a previous study shows that the test sets have a low judgement coverage for the Top-10 results of lexical and neural retrieval models. In this paper we present the novel, relevance judgement test collection TripJudge for TripClick health retrieval. We collect relevance judgements in an annotation campaign and ensure the quality and reusability of TripJudge by a variety of ranking methods for pool creation, by multiple judgements per query-document pair and by an at least moderate inter-annotator agreement. We compare system evaluation with TripJudge and TripClick and find that that click and judgement-based evaluation can lead to substantially different system rankings.

CCS CONCEPTS

• Information systems → Test collections.

KEYWORDS

Test collections, health retrieval, relevance judgements

ACM Reference Format:

Sophia Althammer, Sebastian Hofstätter, Suzan Verberne, and Allan Hanbury. 2022. TripJudge: A Relevance Judgement Test Collection for TripClick Health Retrieval. In *Proceedings of the 31st ACM Int'l Conference on Information and Knowledge Management (CIKM '22)*, Oct. 17–21, 2022, Atlanta, GA, USA. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3511808.3557714>

1 INTRODUCTION

Reliable and robust evaluation of ranking systems is crucial to Information Retrieval (IR) research. Thus a great effort in IR research is put into creating reusable and robust test collections [26, 30, 33, 37] in task-specific evaluation campaigns like TREC [8, 11] or CLEF. These campaigns follow the Cranfield paradigm [9] to create relevance judgements on the pooled output of the participating systems. Recently there has been a growing interest in evaluating the retrieval performance of retrieval models for domain-specific retrieval tasks [2, 13, 14, 27, 36] including the medical domain [22, 23, 35]. Domain-specific retrieval tasks often lack a reliable test collection with human relevance judgments following the Cranfield paradigm [22, 27]. Furthermore it remains unclear how well old test collections can be used to evaluate neural retrieval models, which were not part of the pooling process [33].

In the medical domain, a recently proposed benchmark is the publicly available TripClick collection [22], which contains large-scale click logs from Trip, an English health search engine with professional and non-professional users. It provides two test sets with labels based on the clicks of the users, either estimating relevance by the raw clicks ('Raw') or by the rate of clicks of a document over all retrieved documents for a query ('DCTR'). As the TripClick test sets are based on the clicks of the users, the test sets are biased towards the retrieval model employed by the search engine [34], which remains unknown. In a previous study [13] the test sets were shown to have a low annotation coverage of at most 41% of the Top10 results for lexical and neural retrieval models.

In this paper we address these shortcomings of click-based test collections by creating TripJudge, a relevance judgement test collection for TripClick. We collect relevance judgements by running an annotation campaign on the test set queries of TripClick. In order to increase the reusability of our test collection [6], we use three participating systems for the pool creation from Hofstätter et al. [13] employing lexical and neural retrieval models. To control the quality of the relevance assessments we monitor the annotation time per query, we employ a graded relevance scheme [1, 12] and we employ multiple relevance assessments per query-document pair (we aim for three assessments but have at least two). We reach an at least moderate inter-annotator agreement measured with Cohen's Kappa.

Furthermore we compare the click and judgement-based evaluation of various retrieval systems and investigate how the rankings of



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

CIKM '22, October 17–21, 2022, Atlanta, GA, USA
© 2022 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-9236-5/22/10.
<https://doi.org/10.1145/3511808.3557714>

the systems change. Related work about comparing clicks and judgements for evaluation come to different conclusions. While Joachims et al. [18] and Zobel [37] find reasonable agreement between the clicks and relevance judgements, Kamps et al. [19] demonstrate that system rank correlations between evaluation based on clicks or judgements is low. First we analyze the overlap of the click-based test collections with TripJudge and find that the majority of the documents that are judged as relevant are not labelled in the click-based collections and therefore are considered irrelevant during evaluation. Similar to related work [7, 29, 31], we measure system rank correlation between evaluation with different test collections with Kendalls τ correlation [20]. We find that the rankings with the evaluation of TripJudge differ from the rankings with the click-based test collections. Our contributions are the following:

- We create the relevance judgement-based test collection TripJudge for TripClick health retrieval and make it publicly available on www.github.com/sophialthammer/tripjudge;
- We ensure the quality and reusability of TripJudge by a variety of systems for pool creation, by multiple judgements per query-document pair, and by an at least moderate inter-annotator agreement in our annotation campaign;
- We compare evaluation with click-based TripClick and our judgment-based TripJudge and find that click and judgment-based evaluation can lead to different system rankings

2 METHODOLOGY

We describe how we preprocess the TripClick test queries as well as the pool creation followed by the annotation campaign.

2.1 Data and Pool Preparation

In the TripClick test sets the queries are grouped by their user interaction frequency into Head, Torso and Tail. For the annotation campaign we used the 1175 head queries, which consist of keyword queries. During the campaign we noticed duplicate queries which differ in their casing (lower/uppercasing) or queries without any text (for example "#1 or #2"). We discard these queries from our TripJudge test collection and end up with 1136 unique queries¹.

For the pool creation we use the runs from Hofstätter et al. [13]. In order to have different first stage retrieval methods we use the lexical retrieval run with BM25 [24] (run 1 in Table 2) as well as the SciBERT_{DOT} run (run 2 in Table 2) which is based on dense retrieval [3, 15]. As additional run we use the Ensemble which re-ranks BM25 Top-200 candidates using an Ensemble of BERT_{CAT} based on SciBERT, PubMedBERT-Abstract and PubMedBert-Full Text (run 7 in Table 2). We create the pool by taking the union of the query-document pairs from the Top-10 of the three runs for all test queries and keep the highest rank among the three runs. This results in a total of 29581 pairs and prioritize the annotation according to the rank of the document: all Top-n pairs have priority $10 - k(k \in [1..n])$; the higher the priority, the earlier they are selected for annotation in the annotation interface.

In order to maintain a low latency in our annotation system, we needed to truncate the documents. As the dense retrieval models rely on the text up to 512 BERT tokens, we truncate the document text to this length, which applies to 10% of the documents.

¹We publish the reasons for removal, the removed, and remaining queries in GitHub

# of queries	1136
# of documents in collection	2.3M
# annotated q-d pairs	12590
# of judgements	38810
avg # assessments per query	2.92
# (%), avg annotation time for Wrong	3811 (10%), 39s
# (%), avg annotation time for Topic	10901 (28%), 47s
# (%), avg annotation time for Partial	13008 (33%), 54s
# (%), avg annotation time for Perfect	11090 (29%), 44s
# annotators	135
avg # of annotations per annotator	287

Table 1: Statistics of the annotation campaign.

2.2 Annotation Campaign

We conduct the annotation campaign among 135 computer science students with a target of 300 annotations per annotator and reach an average of 287 annotations per annotator (Table 1).

The users of the Trip search engine are a mix of experts and non-experts. As the annotators are non-experts and as previous work points out the challenges with students as annotators [4, 21], we take several steps to ensure and monitor the quality of the annotations: **1)** We use a 4-graded, ordinal relevance scheme: *Wrong* (1), *Topic* (2), *Partial* (3), and *Perfect* (4), as suggested in related work [1, 12] **2)** We aim for three relevance assessments per query-document pair, on average we reach 2.92 relevance assessments per pair, where we employ majority voting or a heuristic in case of no majority. We discard the pairs with only one relevance assessment. **3)** We monitor the average annotation time per pair per relevance grade, we remove annotations with a short annotation time (below 1 second) and reach an average annotation time of 48 seconds per judgement. **4)** We collect feedback about the campaign from the students at the end of the annotation campaign.

We conducted the annotation campaign using the FiRA interface [16, 17] and ran the campaign for 7 days with a fixed deadline². To control the quality of the judgements during the campaign, we monitored the average number of judgements per 12 hours and we observe the daily average annotation time per relevance grade to detect random judgements. We reach a high average annotation time per relevance grade (Table 1).

Overall the students gave us positive feedback about the evaluation campaign: 33% rated it *Very good*, 28% *Good* and 27% *Decent*, only 12% of the students did not like it. The students could also give written feedback. The main feedback was that it was hard to distinguish between the relevance grades of *Topic* and *Partial*. This difficulty is also reflected in the average time per annotation for these relevance grades. While it took the annotators on average 39 seconds to decide on *Wrong*, the annotations for *Topic* and *Partial* took 47 and 54 seconds on average, respectively. We reached 38810 total judgements and judged at least until the Top-4 for our pool.

3 QUALITY ANALYSIS

After the annotation campaign we processed the 38810 raw relevance judgements to form the TripJudge test collection which we publish in the standard TREC format for qrels.

²We publish the annotation guidelines in our GitHub repository




Figure 1: Distribution of relevance grades for 4-grade and 2-grades, percentage of heuristic and majority voting.




Figure 2: Cohen's Kappa agreement between the annotators and the annotations aggregated with majority voting.

We removed the query-document pairs with only one judgement. We grouped the relevance grades *Wrong* and *Topic* into *Irrelevant* and *Partial* and *Perfect* into *Relevant*, in order to attain a 2-graded relevance judgement with potential higher agreement. We computed the final relevance judgement either via full agreement, if all annotators agree on the relevance grade or with majority voting, if the annotators disagree, or with the heuristic of taking the lowest relevance grade, if the annotators disagree and there is no majority for a relevance grade. We apply this heuristic with the assumption that if disagreement is high, the relevance cannot be definitely decided and therefore the document should be annotated as irrelevant. In Figure 1 we visualize the 4-grade and 2-grade distribution of the judgements and their percentage of full agreement, majority voting and the heuristic. While the full agreement is low (20% of all queries) for the 4-grade relevance judgements, the full agreement for the 2-grade relevance judgements is substantially higher with 48% of all queries. Furthermore, a high percentage of judgements is decided via majority vote. For the 4-grade relevance judgements, 22% of the queries are decided by the heuristic which indicates a high disagreement between the annotators, but the percentage of heuristic decisions for the 2-grade relevance judgements is low with 2%. This shows that the 2-grade relevance judgements are more robust and have a higher agreement between the annotators.




Figure 3: Relevance judgements from TripJudge for the Top-4 of the pool vs TripClick click-based labels from the DCTR or Raw test collection. Green bars denote agreement between the relevance judgement of TripJudge and the click label from TripClick, red bars denote disagreement.

We also study the inter-annotator agreement between the annotators for all relevance judgements. We measure the inter-annotator agreement with Cohen's Kappa κ [10], which is a standard metric to compare multiple sets of judgements and to measure the subjectivity in the assessments. As the 4-grade annotations are ordinal, we use a linear weighted Kappa for them.

In Figure 2 the 2-grade and 4-grade agreement with Cohen's Kappa is visualized with an average κ of 0.63 for the 2-grade and an average weighted κ of 0.60 for the 4-grade relevance judgements, which indicates moderate agreement for the 4-grade and substantial agreement for the 2-grade. For both the 4-grade and 2-grade agreement we reach an at least moderate agreement with 50% of the kappa values between 0.50 and 0.70. Furthermore, a certain level of disagreement in the relevance judgement is expected due to the subjectivity of the individual annotators [5, 28] and our agreement levels align with previous work [1], which also employs non-expert annotators for judgements of TREC collections.

4 TRIPJUDGE VS TRIPCLICK

We compare the relevance judgements of TripJudge with the click-based labels of TripClick and investigate the system ranking difference of the two test collections. Due to the higher inter-annotator agreement we consider the 2-graded judgements of TripJudge.

4.1 Coverage and Intersection

We analyze the coverage and intersection of the annotated query-document pairs between TripJudge and TripClick DCTR and Raw test collection. Figure 3 visualizes the percentage of relevant and irrelevant relevance judgements from TripJudge. The different patterns of the bar visualize the label from the TripClick DCTR or Raw test collection. The labels 1/2/3 from TripClick refer to relevant documents, Label 0 denotes irrelevant documents and unlabelled

Model	TripJudge						TripClick (DCTR)					TripClick (Raw)				
	J@5	J@10	n@5-j	n@5	n@10	R@100	J@5	J@10	n@5	n@10	R@100	J@5	J@10	n@5	n@10	R@100
First stage retrieval																
1 BM25	78%	47%	.761	.694	.570	.771	33%	31%	.122	.140	.499	30%	27%	.199	.198	.464
2 SciBERT _{DOT}	87%	50%	.602	.540	.456	.636	48%	41%	.232	.243	.562	44%	38%	.362	.328	.496
3 PubMedBERT _{DOT}	49%	36%	.652	.377	.356	.649	45%	40%	.223	.235	.582	42%	37%	.345	.318	.518
Re-Ranking (BM25 Top-200)																
4 ColBERT (SciBERT)	64%	44%	.758	.538	.501	.790	52%	47%	.254	.270	.589	49%	43%	.395	.367	.547
5 ColBERT (PubMedBERT)	63%	44%	.758	.527	.493	.777	55%	49%	.261	.278	.595	52%	45%	.412	.382	.551
6 BERT _{CAT}	64%	45%	.757	.540	.506	.818	56%	50%	.271	.287	.594	53%	46%	.421	.389	.552
7 Ensemble	88%	51%	.756	.698	.592	.814	58%	52%	.285	.303	.600	55%	48%	.443	.409	.556

Table 2: Effectiveness results and judgement coverage for judgement-based TripJudge and click-based TripClick DCTR/Raw test collection. J@m denotes the judgement coverage at rank m, n@m denotes the nDCG at cutoff m, -j denotes the j-option in trec_eval when only evaluating on the judged query-document pairs. Top-10 of run 1,2,7 create the pool for TripJudge.

Measure	TripJudge-DCTR	TripJudge-Raw	DCTR-Raw
nDCG@5	0.333	0.333	1.000
nDCG@10	0.428	0.428	1.000
MRR@10	0.238	0.238	1.000
Recall@100	0.238	0.333	0.714

Table 3: Kendall tau correlation between system rankings of TripJudge and TripClick DCTR/Raw for four metrics.

documents are considered as irrelevant during evaluation. The green bars denote agreement between the annotation from TripJudge and TripClick, the red bars denote disagreement. It is striking that all of the relevant documents from the Top-4 of TripJudge are unlabelled in DCTR and therefore considered as irrelevant. Furthermore there is high disagreement between the judgements of TripJudge and click-based labels of TripClick DCTR and Raw.

4.2 System evaluation

We compare the system rankings and the coverage of the relevance judgements for the runs for TripJudge and TripClick DCTR and Raw. For TripJudge and TripClick unlabelled documents are considered as irrelevant. In Table 2 the effectiveness metrics as well as the judgements coverage measured as J at rank n is displayed for various lexical and neural retrieval systems from Hofstätter et al. [13]. For TripJudge we see that the coverage measure with $J@5$ for the runs in the pool (run 1,2,7) is high (around 80%) compared to the coverage of the runs which did not participate in the pooling. However the coverage of TripJudge is higher than the coverage of TripClick DCTR and Raw for the respective runs. The Ensemble of BERT_{CAT} based on SciBERT, PubMedBERT-Abstract and PubMedBert-Full Text (run 7) consistently reaches the highest retrieval performance for the judgement and click-based test collections. Interestingly the dense retrieval model SciBERT_{DOT} (run 2) underperforms BM25 (run 1) when evaluated with TripJudge, although showing substantially higher retrieval effectiveness for the click-based test collections. This suggests that the higher retrieval effectiveness of run 2 compared to run 1 for the click-based collections is due to the higher coverage of annotations.

Furthermore, we compare the difference in system rankings between two test collection with Kendall τ [20], which is a common measure to compare the correlation between two system rankings

[25, 29]. In Table 3 are the Kendall tau correlations between two rankings of two test collections regarding four metrics. Test collections with $\tau > 0.9$ are considered equivalent [29].

For the comparison of TripJudge with the click-based test collections, we see a low correlation of the system rankings for all 4 evaluation metrics. For the click-based test collections we see that they are equivalent for most of the metrics. We conclude that the system rankings differ drastically between TripJudge and TripClick and that TripJudge offers a valuable and reusable relevance judgement set beside the TripClick test sets for evaluating retrieval systems.

5 CONCLUSION AND FUTURE WORK

We present the TripJudge test collection with 38810 relevance judgements for TripClick health retrieval. We describe the annotation campaign for creating the relevance judgements. For increased reusability we used lexical and neural retrieval systems for pool creation. We reach an at least moderate inter-annotator agreement. When comparing the relevance judgements of TripJudge with the click-based annotations from the TripClick test collections, we find that a majority of judged relevant documents were unlabelled in TripClick, and there is a high disagreement between the relevance judgements and the click-based annotations. We re-evaluate lexical and neural models and find a higher judgement coverage for the retrieval runs for TripJudge than for the TripClick test collections. The system rankings substantially differ between the evaluation with the relevance-based and click-based collections.

One limitation of TripJudge is the depth of the relevance judgements as usually a half [11] or a third of the annotated query-document pairs should be relevant [30, 32] while we have 65% relevant. Therefore we see possible future work in annotating to a higher depth. Nevertheless we view TripJudge as a valuable resource for evaluation of the domain specific task of health retrieval. In conclusion, we argue that there must be more effort put into creating relevance judgements based test collections for domain specific retrieval tasks, in order to evaluate different systems in a robust and conclusive way.

ACKNOWLEDGEMENTS

This work was supported by EU Horizon 2020 ITN/ETN on Domain Specific Systems for Information Extraction and Retrieval (ID:860721).

REFERENCES

- [1] Omar Alonso and Stefano Mizzaro. 2012. Using crowdsourcing for TREC relevance assessment. *Information Processing & Management* 48 (11 2012), 1053–1066. <https://doi.org/10.1016/j.ipm.2012.01.004>
- [2] Sophia Althammer, Sebastian Hofstätter, and Allan Hanbury. 2021. Cross-domain Retrieval in the Legal and Patent Domains: a Reproducibility Study. In *Advances in Information Retrieval, 43rd European Conference on IR Research, ECIR 2021*.
- [3] Sophia Althammer, Sebastian Hofstätter, Mete Sertkan, Suzan Verberne, and Allan Hanbury. 2022. Paragraph Aggregation Retrieval Model (PARM) for Dense Document-to-Document Retrieval. In *Advances in Information Retrieval, 44rd European Conference on IR Research, ECIR 2022*.
- [4] Peter Bailey, Nick Craswell, Ian Soboroff, Paul Thomas, Arjen P. de Vries, and Emine Yilmaz. 2008. Relevance Assessment: Are Judges Exchangeable and Does It Matter. In *Proceedings of the 31st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval* (Singapore, Singapore) (SIGIR '08). Association for Computing Machinery, New York, NY, USA, 667–674. <https://doi.org/10.1145/1390334.1390447>
- [5] Pia Borlund. 2003. The Concept of Relevance in IR. *J. Am. Soc. Inf. Sci. Technol.* 54, 10 (aug 2003), 913–925. <https://doi.org/10.1002/asi.10286>
- [6] Chris Buckley, Darrin Dimmick, Ian Soboroff, and Ellen Voorhees. 2007. Bias and the Limits of Pooling for Large Collections. *Inf. Retr.* 10, 6 (dec 2007), 491–508. <https://doi.org/10.1007/s10791-007-9032-x>
- [7] Chris Buckley and Ellen M. Voorhees. 2004. Retrieval Evaluation with Incomplete Information. In *Proceedings of the 27th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval* (Sheffield, United Kingdom) (SIGIR '04). Association for Computing Machinery, New York, NY, USA, 25–32. <https://doi.org/10.1145/1008992.1009000>
- [8] Charles Clarke, Nick Craswell, and Ellen M. Voorhees. 2012. Overview of the TREC 2012 Web Track. In *Text Retrieval Conference (TREC)*. <http://trec.nist.gov/pubs/trec21/papers/WEB12.overview.pdf>
- [9] Cyril W. Cleverdon. 1991. The Significance of the Cranfield Tests on Index Languages. In *Proceedings of the 14th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval* (Chicago, Illinois, USA) (SIGIR '91). Association for Computing Machinery, New York, NY, USA, 3–12. <https://doi.org/10.1145/122860.122861>
- [10] J. Cohen. 1960. A Coefficient of Agreement for Nominal Scales. *Educational and Psychological Measurement* 20, 1 (1960), 37.
- [11] Nick Craswell, Bhaskar Mitra, Emine Yilmaz, Daniel Campos, Ellen M. Voorhees, and Ian Soboroff. 2021. *TREC Deep Learning Track: Reusable Test Collections in the Large Data Regime*. Association for Computing Machinery, New York, NY, USA, 2369–2375. <https://doi.org/10.1145/3404835.3463249>
- [12] Catherine Grady and Matthew Lease. 2010. Crowdsourcing Document Relevance Assessment with Mechanical Turk. In *Proceedings of the NAACL HLT 2010 Workshop on Creating Speech and Language Data with Amazon's Mechanical Turk*. Association for Computational Linguistics, Los Angeles, 172–179. <https://aclanthology.org/W10-0727>
- [13] Sebastian Hofstätter, Sophia Althammer, Mete Sertkan, and Allan Hanbury. 2022. Establishing Strong Baselines for TripClick Health Retrieval. In *Advances in Information Retrieval, 44th European Conference on IR Research, ECIR 2022*.
- [14] Sebastian Hofstätter, Omar Khattab, Sophia Althammer, Mete Sertkan, and Allan Hanbury. 2022. Introducing Neural Bag of Whole-Words with COLBERT: Contextualized Late Interactions using Enhanced Reduction. ArXiv. <https://doi.org/10.48550/ARXIV.2203.13088>
- [15] Sebastian Hofstätter, Sheng-Chieh Lin, Jheng-Hong Yang, Jimmy Lin, and Allan Hanbury. 2021. Efficiently Teaching an Effective Dense Retriever with Balanced Topic Aware Sampling. In *Proc. of SIGIR*.
- [16] Sebastian Hofstätter, Markus Zlabinger, and Allan Hanbury. 2020. Neural-IR-Explorer: A Content-Focused Tool to Explore Neural Re-Ranking Results. In *Advances in Information Retrieval: 42nd European Conference on IR Research, ECIR 2020, Lisbon, Portugal, April 14–17, 2020, Proceedings, Part II* (Lisbon, Portugal). Springer-Verlag, Berlin, Heidelberg, 459–464. https://doi.org/10.1007/978-3-030-45442-5_58
- [17] Sebastian Hofstätter, Markus Zlabinger, Mete Sertkan, Michael Schröder, and Allan Hanbury. 2020. Fine-Grained Relevance Annotations for Multi-Task Document Ranking and Question Answering. In *Proc. of CIKM*.
- [18] Thorsten Joachims, Laura Granka, Bing Pan, Helene Hembrooke, and Geri Gay. 2005. Accurately Interpreting Clickthrough Data as Implicit Feedback. In *Proceedings of the 28th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval* (Salvador, Brazil) (SIGIR '05). Association for Computing Machinery, New York, NY, USA, 154–161. <https://doi.org/10.1145/1076034.1076063>
- [19] Jaap Kamps, Marijn Koolen, and Andrew Trotman. 2009. Comparative analysis of clicks and judgments for IR evaluation. *Journal of Visual Communication and Image Representation - JVCIR* (01 2009), 80–87. <https://doi.org/10.1145/1507509.1507522>
- [20] M. G. Kendall. 1938. A New Measure of Rank Correlation. *Biometrika* 30, 1/2 (1938), 81–93. <http://www.jstor.org/stable/2332226>
- [21] Joao Palotti, Guido Zuccon, Johannes Bernhardt, Allan Hanbury, and Lorraine Goeuriot. 2016. Assessors Agreement: A Case Study Across Assessor Type, Payment Levels, Query Variations and Relevance Dimensions. In *Experimental IR Meets Multilinguality, Multimodality, and Interaction*, Norbert Fuhr, Paulo Quaresma, Teresa Gonçalves, Birger Larsen, Krisztian Balog, Craig Macdonald, Linda Cappellato, and Nicola Ferro (Eds.). Springer International Publishing, Cham, 40–53.
- [22] Navid Rekabsaz, Oleg Lesota, Markus Schedl, Jon Brassey, and Carsten Eickhoff. 2021. *TripClick: The Log Files of a Large Health Web Search Engine*. Association for Computing Machinery, New York, NY, USA, 2507–2513. <https://doi.org/10.1145/3404835.3463242>
- [23] Kirk Roberts, Dina Demner-Fushman, Ellen Voorhees, William Hersh, Steven Bedrick, Alexander Lazar, and Shubham Pant. 2017. Overview of the TREC 2017 Precision Medicine Track. *The 26th text REtrieval conference: TREC. Text REtrieval Conference 26* (11 2017).
- [24] Stephen Robertson and Hugo Zaragoza. 2009. The Probabilistic Relevance Framework: BM25 and Beyond. *Found. Trends Inf. Retr.* 3, 4 (April 2009), 333–389. <https://doi.org/10.1561/15000000019>
- [25] Mark Sanderson and Ian Soboroff. 2007. Problems with Kendall's Tau. 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, 839–840. <https://doi.org/10.1145/1277741.1277935>
- [26] Ian Soboroff. 2017. Building Test Collections: An Interactive Guide for Students and Others Without Their Own Evaluation Conference Series. In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval* (Shinjuku, Tokyo, Japan) (SIGIR '17). Association for Computing Machinery, New York, NY, USA, 1407–1410. <https://doi.org/10.1145/3077136.3082064>
- [27] Nandan Thakur, Nils Reimers, Andreas Rücklé, Abhishek Srivastava, and Iryna Gurevych. 2021. BEIR: A Heterogeneous Benchmark for Zero-shot Evaluation of Information Retrieval Models. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*. <https://openreview.net/forum?id=wCu6T5xFje>
- [28] Suzan Verberne, Antal van den Bosch, Sander Wubben, and Emiel Kraemer. 2017. Automatic summarization of domain-specific forum threads: collecting reference data. In *Proceedings of the 2017 Conference on Conference Human Information Interaction and Retrieval*, 253–256.
- [29] Ellen M. Voorhees. 1998. Variations in Relevance Judgments and the Measurement of Retrieval Effectiveness. In *Proceedings of the 21st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval* (Melbourne, Australia) (SIGIR '98). Association for Computing Machinery, New York, NY, USA, 315–323. <https://doi.org/10.1145/290941.291017>
- [30] Ellen M. Voorhees. 2018. On Building Fair and Reusable Test Collections Using Bandit Techniques. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management* (Torino, Italy) (CIKM '18). Association for Computing Machinery, New York, NY, USA, 407–416. <https://doi.org/10.1145/3269206.3271766>
- [31] Ellen M. Voorhees and Chris Buckley. 2002. The Effect of Topic Set Size on Retrieval Experiment Error. In *Proceedings of the 25th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval* (Tampere, Finland) (SIGIR '02). Association for Computing Machinery, New York, NY, USA, 316–323. <https://doi.org/10.1145/564376.564432>
- [32] Ellen M. Voorhees and Kirk Roberts. 2021. *On the Quality of the TREC-COVID IR Test Collections*. Association for Computing Machinery, New York, NY, USA, 2422–2428. <https://doi.org/10.1145/3404835.3463244>
- [33] Ellen M. Voorhees, Ian Soboroff, and Jimmy Lin. 2022. Can Old TREC Collections Reliably Evaluate Modern Neural Retrieval Models? *CoRR* abs/2201.11086 (2022). arXiv:2201.11086 <https://arxiv.org/abs/2201.11086>
- [34] Ryen White. 2013. Beliefs and Biases in Web Search. In *Proceedings of the 36th International ACM SIGIR Conference on Research and Development in Information Retrieval* (Dublin, Ireland) (SIGIR '13). Association for Computing Machinery, New York, NY, USA, 3–12. <https://doi.org/10.1145/2484028.2484053>
- [35] Chenyan Xiong, Zhenghao Liu, Si Sun, Zhuyun Dai, Kaitao Zhang, Shi Yu, Zhiyuan Liu, Hoifung Poon, Jianfeng Gao, and Paul Bennett. 2020. CMT in TREC-COVID Round 2: Mitigating the Generalization Gaps from Web to Special Domain Search. arXiv:2011.01580 [cs.IR]
- [36] Jingtao Zhan, Xiaohui Xie, Jiabin Mao, Yiqun Liu, Min Zhang, and Shaoping Ma. 2022. Evaluating Extrapolation Performance of Dense Retrieval. (2022). <https://doi.org/10.48550/ARXIV.2204.11447>
- [37] Justin Zobel. 1998. How Reliable Are the Results of Large-Scale Information Retrieval Experiments?. In *Proceedings of the 21st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval* (Melbourne, Australia) (SIGIR '98). Association for Computing Machinery, New York, NY, USA, 307–314. <https://doi.org/10.1145/290941.291014>