

MacAvaney, S. and Soldaini, L. (2023) One-Shot Labeling for Automatic Relevance Estimation. In: 46th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR23), Taipei, Taiwan, 23-27 July 2023, pp. 2230-2235. ISBN 9781450394086 (doi: [10.1145/3539618.3592032](https://doi.org/10.1145/3539618.3592032))

There may be differences between this version and the published version. You are advised to consult the publisher's version if you wish to cite from it.

© The Authors 2023. This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in the Proceedings of the ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR23), Taipei, Taiwan, 23-27 July 2023, pp. 2230-2235. ISBN 9781450394086 (doi: [10.1145/3539618.3592032](https://doi.org/10.1145/3539618.3592032))

<https://eprints.gla.ac.uk/296335/>

Deposited on: 02 May 2023

One-Shot Labeling for Automatic Relevance Estimation

Sean MacAvaney*

University of Glasgow
United Kingdom

sean.macavaney@glasgow.ac.uk


Luca Soldaini*

Allen Institute for AI
United States

lucas@allenai.org

ABSTRACT

Dealing with unjudged documents (“holes”) in relevance assessments is a perennial problem when evaluating search systems with offline experiments. Holes can reduce the apparent effectiveness of retrieval systems during evaluation and introduce biases in models trained with incomplete data. In this work, we explore whether large language models can help us fill such holes to improve offline evaluations. We examine an extreme, albeit common, evaluation setting wherein only a single known relevant document per query is available for evaluation. We then explore various approaches for predicting the relevance of unjudged documents with respect to a query and the known relevant document, including nearest neighbor, supervised, and prompting techniques. We find that although the predictions of these One-Shot Labelers (1SL) frequently disagree with human assessments, the labels they produce yield a far more reliable ranking of systems than the single labels do alone. Specifically, the strongest approaches can consistently reach system ranking correlations of over 0.86 with the full rankings over a variety of measures. Meanwhile, the approach substantially increases the reliability of t -tests due to filling holes in relevance assessments, giving researchers more confidence in results they find to be significant. Alongside this work, we release an easy-to-use software package to enable the use of 1SL for evaluation of other ad-hoc collections or systems.

 github.com/seanmacavaney/autoqrels

CCS CONCEPTS

• Information systems → Evaluation of retrieval results.

KEYWORDS

few-shot learning, relevance assessments, neural networks

ACM Reference Format:

Sean MacAvaney and Luca Soldaini. 2023. One-Shot Labeling for Automatic Relevance Estimation. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '23)*, July 23–27, 2023, Taipei, Taiwan. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3539618.3592032>

* Listed alphabetically; both authors contributed equally to this research.

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 republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SIGIR '23, July 23–27, 2023, Taipei, Taiwan

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-9408-6/23/07...\$15.00

<https://doi.org/10.1145/3539618.3592032>

1 INTRODUCTION

Cranfield-style test collections [8]—consisting of a document corpus, a set of queries, and manually-assessed relevance judgements between queries and documents—are the bedrock of most information retrieval research. Since it is often prohibitively expensive to assess every document for every query (for a non-trivially sized corpus), a subset of documents are chosen for each query for assessment. This is typically accomplished through a pooling process, wherein the top results from one or more systems are selected for assessment. Decisions about pooling (e.g., how many systems, how many results per system, etc.) have a direct effect on the cost of assessment in terms of human labor. On one extreme, some collections only pool from the very top results (e.g., top 10) of a single system. This makes the cost of assessing a single query low, but results in very incomplete assessments [1] and survivorship bias [14]. At the other extreme, some collections assess the tens to hundreds of documents from dozens of systems. This comes at a much higher cost, but results in test collections that are more complete, improving test collection reusability since documents returned from new systems are more likely to already be assessed [46]. Given an identical assessment budget, arguments have been made in favor of both techniques. While reusability of deeper pools is an appealing advantage, shallower assessment means that more queries can be assessed, improving coverage of various information needs and making statistical significance easier [3]. Shallow assessments can also be more easily collected, since they do not require a diverse set of input systems, and can be inferred from query logs (e.g., [34]), making them popular in industry and in large-scale collections like MS MARCO [30]. Further, even with deep assessments, “holes” (missing assessments) still exist, especially for large corpora, where many documents can satisfy a query’s information need [45].

In this work, we explore the extent to which we can fill holes in relevance assessments using models over query and document text. To reign in the scope, we focus on a setting where only a single relevant document per query, sampled from the top results of a baseline system, is known. This setup is akin to the MS MARCO dataset, and is often used in industry. We explore several “one-shot” approaches¹ (henceforth *one-shot labelers*, or 1SL) for predicting the relevance of holes, including a k nearest neighbor search of the known relevant document, a supervised relative relevance scorer (DuoT5 [32]), and an instruction-tuned model (FlanT5 [7]). Through experiments on the TREC DL 2019, 2020, and 2021 datasets, we find that 1SL approaches are not yet able to sufficiently identify all relevant documents (best F1 score: 0.63). Therefore, 1SL is likely inadequate for estimating system recall. However, we find

¹ Zero-shot methods are simply relevance models, which could be used as ranking models directly. Since these would clearly favor systems that use them for ranking, we only consider approaches that make use of a known relevant document.

that when incorporating the estimated relevance scores into recall-agnostic C/W/L measures [27], we are able to achieve high correlation with the ranking of systems from full TREC assessments (consistently ≥ 0.86 correlation, often ≥ 0.98). Further, we find that using these approaches yields more reliable statistical tests when comparing system performance. These results suggest that automatic 1SL techniques could potentially replace expensive deep assessments when evaluating with precision-oriented measures.

2 RELATED WORK

The Cranfield evaluation paradigm [8] was introduced as a more cost-effective alternative to a user study: instead of asking a population of users to interact with IR systems, use expert assessors to judge whether documents retrieved by IR systems are relevant for a set of predefined topics, each representing an hypothetical information need. Cranfield evaluations rely on a series of assumptions [44] to justify this simplification: (i) relevance can be approximated as similarity between documents and topics, (ii) the topics are a representative approximation of user information need, and, most pertinent to our work, (iii) all relevant documents for a query are known. The latter requirement is often violated, as it is unfeasible to judge every document in a collection containing up to billions of items [44, 54]. Instead, pooling techniques are often used to create a test collection: all systems part of a shared task contribute² to a pool of documents [41]; items in the pool get judged by annotators; any document not in the pool is assumed to be irrelevant.

Metrics for Partially Assessed Collections. In response to the inevitable violation of completeness assumption in pooled evaluation, researchers have proposed metrics that take this limitation into account. Early on, Buckley and Voorhees [4] measured the effect of this violation in common metrics, such as mean Average Precision, Precision@k retrieved results, and R-Precision; further, they proposed Bpref, a metric that exhibits fewer errors in ranking systems that contributed to the evaluation pool. Follow up work proposed rank-biased precision (RBP) [28, 29] and inferred precision [36, 52]; the latter has been argued to be more effective than Bpref and RBP thanks to its superior discriminative power [37]. Fröbe et al. [13] proposed using bootstrapping to overcome holes.

Despite these efforts, the community has largely not adopted these measures;³ instead, measures that simply treat unjudged documents as irrelevant—such as nDCG, MAP, and P@k—remain in common use. Meanwhile, it has also become increasingly popular to simply report the proportion of documents in the top results that have relevance assessments (Judged@k, or conversely Hole@k) alongside these measures (e.g., [23, 43]). Recognizing the preference for traditional metrics that rely on full judgments, our 1SLs can effectively fill holes in the evaluation pool, thus mitigating violations of the completeness assumption.

Automated Assessment of Documents. Given the limitations arising from incomplete judgments, automated tools have been proposed to assist or automate the evaluation of retrieval systems. Some researchers have focused on completely automating IR evaluation [31, 40, 42, 50, 51]; however, these works have been generally

² Many strategies can be employed to sample documents from the output of a IR system in order to build a pool. ³ As a few examples, none of the TREC 2021 tracks reported measures that were robust to unjudged documents as official evaluation measures, nor did any of papers at SIGIR 2022 “Search and Ranking” session.

Table 1: Comparison of proposed gain estimators. \mathcal{X}^* indicates a system that was not directly supervised on the data, but may seen some during instruction tuning.

G	Considers			MS MARCO	
	q	d^+	$d^?$	Supervised	# Params
MaxRep-BM25	\mathcal{X}	✓	✓	\mathcal{X}	n/a
MaxRep-TCT	\mathcal{X}	✓	✓	✓	110M
DuoT5	✓	✓	✓	✓	3B
DuoPrompt	✓	✓	✓	\mathcal{X}^*	3B

criticized as evaluating retrieval systems by “popularity as opposed to performance” [2]. Further, recent studies of fully-automated relevance estimators found that careful selection of topics [15] and ensembling of estimators [35] are necessary to best employ these otherwise unstable techniques.

Rather than completely relying on models, our work instead focuses on how to best exploit a very small number of relevance judgments provided by human assessors to label an otherwise unusable collection. Previous work in this area is built upon the cluster hypothesis [19]: that is, it assumes that documents that are similar to relevant documents must also be relevant. For example, Carterette and Allan [6] used *tfidf*-weighted vector space similarity to estimate the probability of a unjudged document. Büttcher et al. [5] experimented with predicting relevance of documents using models that estimate unigram distribution over relevant and non-relevant documents. Another direction involves human-written criteria, which can be automatically assessed by QA systems [38]. Hui and Berberich [17] evaluates different strategies to measure document similarity for MaxRep, concluding that bag of words and paragraph embedding approaches to be most suitable. In this manuscript, we extend their approach by using modern contextualized embedders, and compare it with our proposed approaches.

Outside adhoc retrieval, others have proposed methods for completing partially annotated pools for temporal summarization [25], novelty-oriented retrieval [18], and question answering [47] tasks.

3 AUTOMATIC ONE-SHOT LABELING (1SL)

Consider the case where only a single relevant document d^+ per query q is known. Many common evaluation measures (e.g., those in the C/W/L framework [27]) begin by mapping a system ranking $[d_1, d_2, \dots, d_n]$ to gain values $[g_1, g_2, \dots, g_n]$. It is common practice to treat documents that are not assessed as having no gain, i.e.,

$$g_i = \begin{cases} 1 & \text{if } d_i = d^+ \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

In this work, we explore techniques for estimating the gain for an unknown document in the ranking $d^?$ (i.e., “holes”) with respect to d^+ and q using a relevance estimator function $G(q, d^+, d^?) \in [0, 1]$:⁴

$$g_i = \begin{cases} 1 & \text{if } d_i = d^+ \\ G(q, d^+, d_i) & \text{otherwise} \end{cases} \quad (2)$$

⁴ Without loss of generality, we assume the outputs are in the range of 0 to 1; in cases where they are not, the outputs could simply be normalized using various approaches.

These gains can then be used by various evaluation measures. We refer to G as a One-Shot Labeler ($\mathbb{1}SL$), given that it provides labels for a query based on a single known relevance label. We now describe several $\mathbb{1}SL$ implementations, summarized in Table 1.

MaxRep. We adapt MaxRep by Hui and Berberich [16] to a one-shot setting: given d^+ , we retrieve the k nearest neighboring documents and also treat them as relevant, with a linearly degrading gain (i.e., the i^{th} nearest neighbor receives a gain of $(k-i)/k$). Note that this approach does not use the query directly, except insofar as that d^+ is known to be relevant to q . In line with prior work [17], we explore both lexical similarity and semantic similarity. For lexical similarity, we use the top $k = 128$ BM25 results for the given document, and for semantic similarity, we use the top $k = 128$ nearest neighbors over TCT-ColBERT [20] embeddings⁵ – both of which have been recently shown as an effective signal to identify additional relevant documents when re-ranking [24].

DuoT5. The DuoT5 model [32] is a sequence-to-sequence model trained to provide a relative relevance signal between two documents, with respect to a query. It was proposed for use as a final stage re-ranker, where it can help refine the top documents in a ranked list by comparing them with one another. We recognize that it may also be suitable estimating the relevance in a one-shot setting. Specifically, by presenting the model with q and d^+ , the model estimates the probability that d^2 is more relevant to q than d^+ . Note that in contrast with when DuoT5 is used as a re-ranker, this case makes use of a *known* relevant document, and uses the information to predict the relevance of additional documents.

DuoPrompt. Recent advances in instruction-tuned models [39, 48, 49] have lead to the rise of models that can be prompted “in context” to complete a task. Practically, this consists of prefixing each sample one wishes to perform inference on with instructions describing the task. In our case, we prepend query and passage with a short text prompting the model to estimating relevance of passages for the given query.⁶ Akin to DuoT5, DuoPrompt prompts the model to assess whether d^2 is as relevant as d^+ for q . We teacher-force the model to generate either “yes” or “no”; the probability values obtained by computing the softmax over the logits of these two tokens is used as the relevance score.⁷ We leverage the pre-trained Flan-T5 [7], a variant of T5 [33] that has been fine-tuned on a large set of instruction datasets.⁸

4 EXPERIMENTS AND RESULTS

In this section, we establish two research questions and evaluate how one-shot modelers perform in these two settings. First, **RQ1** asks: *can $\mathbb{1}SLs$ be used to directly label unjudged documents in a shared task pool?* While appealing, we show that proposed techniques cannot reliably provide binary relevance labels. Then, **RQ2** asks (despite the limitations described above): *can $\mathbb{1}SLs$ be used*

⁵ castorini/tct_colbert-v2-hnp-msmarco-r2 ⁶ Prompt: Determine if passage B is as relevant as passage A. Passage A: <d+> Passage B: <d?> Query: <q> Is passage B as relevant as passage A? ⁷ Inspired by its success in prior work [21, 26, 53], we also experimented with adding examples of the tasks to the prompt, a practice known as in-context learning. However, we found it to be worse than not including any labeled samples. This might be due to the large variance in passages, but we leave the problem of exploring in-context learning to future work. ⁸ We note that the instruction data includes 500 samples from the MS MARCO QnA train dataset, resulting in a model that is supervised on a small amount of data from our target domain. However, the instruction prompts do not make use of the relevance labels; only a set of candidates, the queries and assessor-written answers.




Figure 1: Precision-recall curve of one-shot labelers ($\mathbb{1}SLs$) on the $qrels$ from the 2019 TREC Deep Learning Track.

to rank IR systems a shared task? Crucially, we find that $\mathbb{1}SLs$ are very accurate relative performance estimators across three adhoc passage retrieval benchmarks and three metrics. To answer both questions, we use a simulated shallow document pool, based on the task’s official BM25 baseline.⁹ We choose the first relevant document from the run, and treat it as the single d^+ for the query.¹⁰

4.1 RQ1: $\mathbb{1}SLs$ as Relevance Estimators

Setup. In order to establish whether any of the methods described in Section 3 can function as automatic binary relevance assessors, we measure their efficacy at inferring relevance labels for unknown documents. Given the simulated shallow pools, we predict the binary labels of $qrels$ from the TREC DL 2019 track [11].

Results. We report the precision-recall curves of these predictions in Figure 1. We note clear differences among the methods introduced in Section 3. First, MaxRep methods achieve notably worse F1 and Average Precision (AP) than other methods. This is likely due to MaxRep not considering query text when estimating relevance, instead only relying on intra-document similarity. On the other hand, DuoPrompt and DuoT5 are moderately effective at the tasks, achieving an F1-score of at most 0.55 and 0.59 (AP: 0.59 and 0.56) respectively. However, these results are far from meeting necessary accuracy to use DuoPrompt or DuoT5 as fully automatic labelers, as they are unable to reliably identify the long tail of relevant documents. Further, we note that their PR curves are very different, making difficult to tune a reliable operating point for the binary labels. This observation confirms that the difficulties noted by Hauff et al. [15] and Roitero et al. [35] in tuning zero-shot labelers applies to our $\mathbb{1}SLs$.

4.2 RQ2: $\mathbb{1}SLs$ for System Evaluation

Setup. In order for $\mathbb{1}SLs$ to be a viable alternative to highly skilled assessors, they should lead to similar evaluation outcomes. More specifically, given a set of systems submitting runs to a shared track, results using $\mathbb{1}SLs$ should result in a ranking of participating systems that is comparable to that of human labels and for similar outcomes from statistical significance tests.

To study if that is the case, we calculate various evaluation measures for the submissions to the TREC 2019, 2020, and 2021 Deep

⁹ bm25base_p for DL 2019, p_bm25 for DL 2020 and 2021 ¹⁰ On average, this requires 6 passages to be “examined” per query.

Table 2: $\mathbb{1}$ SLs as system evaluators. Correlations of the TREC submissions are reported in terms of Kendall’s τ , Spearman’s ρ , Rank Biased Overlap ($p = 0.9$), and the t -test ($p < 0.05$ w/ Bonferroni) false positive/negative rates of the top identified system.

Measure	Holes	msmarco-passage/trec-d1-2019					msmarco-passage/trec-d1-2020					msmarco-passage-v2/trec-d1-2021				
		τ	ρ	RBO	t -FNR	t -FPR	τ	ρ	RBO	t -FNR	t -FPR	τ	ρ	RBO	t -FNR	t -FPR
SDCG@10	Non-relevant	-0.204	-0.248	0.420	0.000	0.857	0.419	0.498	0.731	0.000	0.510	0.402	0.531	0.564	0.125	0.714
	MaxRep-BM25	0.240	0.320	0.690	0.000	0.229	0.468	0.597	0.710	0.000	0.059	0.455	0.576	0.609	0.200	0.489
	MaxRep-TCT	0.829	0.958	0.818	0.208	0.083	0.793	0.933	0.920	0.263	0.053	0.578	0.761	0.817	0.222	0.300
	DuoT5	0.889	0.972	0.812	0.000	0.417	0.837	0.944	0.939	0.000	0.895	0.859	0.963	0.880	0.000	0.571
	DuoPrompt	0.904	0.980	0.830	0.160	0.000	0.909	0.986	0.863	0.184	0.000	0.910	0.983	0.925	0.040	0.143
P@10	Non-relevant	-0.033	-0.003	0.511	0.000	0.571	0.429	0.526	0.695	0.143	0.120	0.387	0.520	0.528	0.182	0.739
	MaxRep-BM25	0.362	0.452	0.722	0.000	0.314	0.442	0.565	0.660	0.000	0.118	0.425	0.579	0.571	0.167	0.467
	MaxRep-TCT	0.870	0.971	0.874	0.208	0.083	0.792	0.928	0.769	0.214	0.067	0.579	0.767	0.781	0.154	0.290
	DuoT5	0.891	0.974	0.811	0.000	0.333	0.868	0.970	0.781	0.000	0.867	0.858	0.962	0.916	0.000	0.700
	DuoPrompt	0.891	0.981	0.872	0.000	0.250	0.907	0.986	0.817	0.143	0.000	0.868	0.972	0.903	0.000	0.200
RBP(p=0.8)	Non-relevant	-0.177	-0.220	0.445	0.000	0.857	0.437	0.510	0.840	0.000	0.353	0.387	0.516	0.535	0.167	0.706
	MaxRep-BM25	0.246	0.332	0.713	0.000	0.200	0.452	0.583	0.715	0.000	0.059	0.446	0.562	0.601	0.182	0.565
	MaxRep-TCT	0.853	0.963	0.829	0.167	0.083	0.791	0.930	0.917	0.333	0.000	0.569	0.761	0.766	0.259	0.333
	DuoT5	0.892	0.975	0.854	0.000	0.667	0.806	0.929	0.939	0.000	0.889	0.863	0.970	0.882	0.000	0.625
	DuoPrompt	0.919	0.986	0.889	0.040	0.000	0.889	0.980	0.964	0.179	0.000	0.890	0.980	0.897	0.000	0.250

Learning tracks [9–11] (37, 58, and 58 runs, respectively) using our simulated shallow pools. We treat holes as either non-relevant (0-gain, the baseline), or by replacing the relevance gains using $\mathbb{1}$ SL scores. Since we found that $\mathbb{1}$ SLs are not yet capable of providing a reliable measure of recall, we instead focus on three recall-agnostic evaluation measures from the C/W/L framework: SDCG@10, P@10, and RBP(p=0.8) using *ir*-measures [22].¹¹ We use variants of P@10 (*weighted-precision*) and RBP that make use of partial gains [27].

Results. We report our main findings in Table 2. Here the first row in each section corresponds to estimating system ranks using only the one labeled document from BM25; subsequent rows show how $\mathbb{1}$ SLs perform. Overall, we observe that DuoPrompt usually leads to more accurate estimation across all years and metrics: it achieves a correlation of 0.87 to 0.92, and a rank correlation of 0.97 to 0.98. More importantly, t -tests conducted using the method are far more reliable than other methods; it rarely gives false positive results (as compared to the full *qrels*), and infrequently yields false negatives. DuoT5 also often achieves high correlations with the full *qrels*, but results in more false positive t -tests. On the other hand, MaxRep [16] does not offer reliable estimation of labels, leading to inferior correlations when compared to DuoT5 and DuoPrompt. MaxRep-BM25 is particularly unreliable, with all its correlation < 0.5 ; MaxRep-TCT, while more accurate than MaxRep-TCT, still suffers from higher false positive rates in most cases than DuoPrompt. Finally, while unsurprising, we note that, ordering of systems induced by using only one positive labels (“not relevant” in Table 2) is a very inaccurate approach, leading to very low or even negative correlations and a high proportion of false positive tests. Clearly, in this case, the single system that contributed to the *qrels* is unfairly favored above the other systems, leading to artificially higher evaluation measures than it would under deeper pools.

5 CONCLUSIONS AND RECOMMENDATIONS

In this work, we presented a new approach to replace deep manual relevance assessments with prediction from automatic one-shot

¹¹ nDCG@10 and P@10 are official TREC DL task metrics. SDCG@10 is similar to nDCG@10, but assumes there are 10 fully-relevant documents available (and is therefore unaffected by the introduction of new relevant documents to the *qrels*).

labelers ($\mathbb{1}$ SLs). These models rely on one relevant document per topic to evaluate a pool of documents retrieved by IR systems. This approach can reliably predict ordering of systems across multiple shared tasks and metrics. One obvious question is: **Can shared tasks organizers largely do away with human labelers?**¹²

First, we emphatically stress that whether to completely replace human assessors **is not** the right question to ask; rather, accurate automatic labelers allow re-distributing existing annotation budgets in ways that unlock new evaluation practices. First, our approach requires at least one high-quality relevant document. More importantly, $\mathbb{1}$ SLs can be used in conjunction with human annotators to obtain a pool that is both wide and deep: human annotators can be tasked to label a few documents for each topic in a large set of queries, with automatic labelers inferring judgements on the remainder of the pool. This would lead easier measurement of significant improvements [3], and the ability to study more complex information needs. Moreover, automatic labelers could increase value of existing test collections: an estimation of the relevance of any previously unjudged document retrieved by new systems could be used instead of treating them as irrelevant.

We note that the scope of this study has limitations that would need to be explored before this approach be used in practice. We only explore passage retrieval; providing relative relevance estimates over entire documents is likely a far more challenging task, especially with context window limitations of neural language models. Further, we only explore situations where a single relevant document is known; rather, we often know a set of relevant documents (of potentially differing relevance grades). Techniques for using multiple labels (e.g., aggregating scores from multiple relevant documents) need to be explored. We also dodged the issue of measures that rely on recall by using the utility-based C/W/L measures. Although such measures are practical in our setting, measures that depend on system recall (such as nDCG, MAP, etc.) are commonly used in evaluations. Finally, more work is needed to explore potential biases these models introduce, especially when the systems share a base language model.

¹² We refer the readers to Faggioli et al. [12], a contemporaneous article that explores this question in greater detail. Here, we answer this question with respect to our study.

REFERENCES

- [1] Negar Arabzadeh, Alexandra Vtyurina, Xinyi Yan, and Charles L. A. Clarke. 2022. Shallow pooling for sparse labels. *Inf. Retr. J.* 25, 4 (2022), 365–385. <https://doi.org/10.1007/s10791-022-09411-0>
- [2] Javed A. Aslam and Robert Savell. 2003. On the effectiveness of evaluating retrieval systems in the absence of relevance judgments. In *SIGIR 2003: Proceedings of the 26th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, July 28 - August 1, 2003, Toronto, Canada*, Charles L. A. Clarke, Gordon V. Cormack, Jamie Callan, David Hawking, and Alan F. Smeaton (Eds.). ACM, 361–362. <https://doi.org/10.1145/860435.860501>
- [3] Leonid Boytsov, Anna Belova, and Peter Westfall. 2013. Deciding on an adjustment for multiplicity in IR experiments. In *The 36th International ACM SIGIR conference on research and development in Information Retrieval, SIGIR '13, Dublin, Ireland - July 28 - August 01, 2013*, Gareth J. F. Jones, Paraic Sheridan, Diane Kelly, Maarten de Rijke, and Tetsuya Sakai (Eds.). ACM, 403–412. <https://doi.org/10.1145/2484028.2484034>
- [4] Chris Buckley and Ellen M. Voorhees. 2004. Retrieval evaluation with incomplete information. In *SIGIR 2004: Proceedings of the 27th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, Sheffield, UK, July 25–29, 2004*, Mark Sanderson, Kalervo Järvelin, James Allan, and Peter Bruza (Eds.). ACM, 25–32. <https://doi.org/10.1145/1008992.1009000>
- [5] Stefan Büttcher, Charles L. A. Clarke, Peter C. K. Yeung, and Ian Soboroff. 2007. Reliable information retrieval evaluation with incomplete and biased judgements. In *SIGIR 2007: Proceedings of the 30th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, Amsterdam, The Netherlands, July 23–27, 2007*, Wessel Kraaij, Arjen P. de Vries, Charles L. A. Clarke, Norbert Fuhr, and Noriko Kando (Eds.). ACM, 63–70. <https://doi.org/10.1145/1277741.1277755>
- [6] Ben Carterette and James Allan. 2007. Semiautomatic evaluation of retrieval systems using document similarities. In *Proceedings of the Sixteenth ACM Conference on Information and Knowledge Management, CIKM 2007, Lisbon, Portugal, November 6–10, 2007*, Mário J. Silva, Alberto H. F. Laender, Ricardo A. Baeza-Yates, Deborah L. McGuinness, Bjørn Olstad, Øystein Haug Olsen, and André O. Falcao (Eds.). ACM, 873–876. <https://doi.org/10.1145/1321440.1321564>
- [7] Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Eric Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, Albert Webson, Shixiang Shane Gu, Zhuyun Dai, Mirac Suzgun, Xinyun Chen, Aakanksha Chowdhery, Sharan Narang, Gaurav Mishra, Adams Yu, Vincent Y. Zhao, Yanping Huang, Andrew M. Dai, Hongkun Yu, Slav Petrov, Ed H. Chi, Jeff Dean, Jacob Devlin, Adam Roberts, Denny Zhou, Quoc V. Le, and Jason Wei. 2022. Scaling Instruction-Finetuned Language Models. *CoRR* abs/2210.11416 (2022). <https://doi.org/10.48550/arXiv.2210.11416> arXiv:2210.11416
- [8] Cyril W. Cleverdon. 1967. The Cranfield tests on index language devices. *Aslib Proceedings* 19, 6 (1967), 173–194. <https://doi.org/10.1108/eb050097>
- [9] Nick Craswell, Bhaskar Mitra, Emine Yilmaz, and Daniel Campos. 2020. Overview of the TREC 2020 Deep Learning Track. In *Proceedings of the Twenty-Ninth Text REtrieval Conference, TREC 2020, Virtual Event [Gaithersburg, Maryland, USA], November 16–20, 2020 (NIST Special Publication, Vol. 1266)*, Ellen M. Voorhees and Angela Ellis (Eds.). National Institute of Standards and Technology (NIST). <https://trec.nist.gov/pubs/trec29/papers/OVERVIEW.DL.pdf>
- [10] Nick Craswell, Bhaskar Mitra, Emine Yilmaz, Daniel Campos, and Jimmy Lin. 2021. Overview of the TREC 2021 Deep Learning Track. In *Proceedings of the Thirtieth Text REtrieval Conference, TREC 2021, Virtual Event [Gaithersburg, Maryland, USA], November, 2021 (NIST Special Publication)*, Ian Soboroff and Angela Ellis (Eds.). National Institute of Standards and Technology (NIST). <https://trec.nist.gov/pubs/trec30/papers/Overview-DL.pdf>
- [11] Nick Craswell, Bhaskar Mitra, Emine Yilmaz, Daniel Campos, and Ellen M. Voorhees. 2019. Overview of the TREC 2019 Deep Learning Track. In *Proceedings of the Twenty-Eighth Text REtrieval Conference, TREC 2019, Gaithersburg, Maryland, USA, November 13–15, 2019 (NIST Special Publication, Vol. 1250)*, Ellen M. Voorhees and Angela Ellis (Eds.). National Institute of Standards and Technology (NIST). <https://trec.nist.gov/pubs/trec28/papers/OVERVIEW.DL.pdf>
- [12] Guglielmo Faggioli, Laura Dietz, Charles L. A. Clarke, Gianluca Demartini, Matthias Hagen, Claudia Hauff, Noriko Kando, Evangelos Kanoulas, Martin Potthast, Benno Stein, and Henning Wachsmuth. 2023. Perspectives on Large Language Models for Relevance Judgment. *CoRR* abs/2304.09161 (2023). <https://doi.org/10.48550/arXiv.2304.09161> arXiv:2304.09161
- [13] Maik Fröbe, Lukas Gienapp, Martin Potthast, and Matthias Hagen. 2023. Bootstrapped nDCG Estimation in the Presence of Unjudged Documents. In *Advances in Information Retrieval - 45th European Conference on Information Retrieval, ECIR 2023, Dublin, Ireland, April 2–6, 2023, Proceedings, Part I (Lecture Notes in Computer Science, Vol. 13980)*, Jaap Kamps, Lorraine Goeriot, Fabio Crestani, Maria Maistro, Hideo Joho, Brian Davis, Cathal Gurrin, Udo Kruschwitz, and Annalina Caputo (Eds.). Springer, 313–329. https://doi.org/10.1007/978-3-031-28244-7_20
- [14] Prashansa Gupta and Sean MacAvaney. 2022. On Survivorship Bias in MS MARCO. In *SIGIR '22: The 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, Madrid, Spain, July 11 - 15, 2022*, Enrique Amigó, Pablo Castells, Julio Gonzalo, Ben Carterette, J. Shane Culpepper, and Gabriella Kazai (Eds.). ACM, 2214–2219. <https://doi.org/10.1145/3477495.3531832>
- [15] Claudia Hauff, Djoerd Hiemstra, Leif Azzopardi, and Franciska de Jong. 2010. A Case for Automatic System Evaluation. In *Advances in Information Retrieval, 32nd European Conference on IR Research, ECIR 2010, Milton Keynes, UK, March 28–31, 2010. Proceedings (Lecture Notes in Computer Science, Vol. 5993)*, Cathal Gurrin, Yulan He, Gabriella Kazai, Udo Kruschwitz, Suzanne Little, Thomas Roelleke, Stefan M. Rieger, and Keith van Rijsbergen (Eds.). Springer, 153–165. https://doi.org/10.1007/978-3-642-12275-0_16
- [16] Kai Hui and Klaus Berberich. 2015. Selective Labeling and Incomplete Label Mitigation for Low-Cost Evaluation. In *String Processing and Information Retrieval - 22nd International Symposium, SPIRE 2015, London, UK, September 1–4, 2015, Proceedings (Lecture Notes in Computer Science, Vol. 9309)*, Costas S. Iliopoulos, Simon J. Puglisi, and Emine Yilmaz (Eds.). Springer, 137–148. https://doi.org/10.1007/978-3-319-23826-5_14
- [17] Kai Hui and Klaus Berberich. 2016. Cluster Hypothesis in Low-Cost IR Evaluation with Different Document Representations. In *Proceedings of the 25th International Conference on World Wide Web, WWW 2016, Montreal, Canada, April 11–15, 2016, Companion Volume*, Jacqueline Bourdeau, Jim Hendler, Roger Nkambou, Ian Horrocks, and Ben Y. Zhao (Eds.). ACM, 47–48. <https://doi.org/10.1145/2872518.2889370>
- [18] Kai Hui, Klaus Berberich, and Ida Mele. 2017. Dealing with Incomplete Judgments in Cascade Measures. In *Proceedings of the ACM SIGIR International Conference on Theory of Information Retrieval, ICTIR 2017, Amsterdam, The Netherlands, October 1–4, 2017*, Jaap Kamps, Evangelos Kanoulas, Maarten de Rijke, Hui Fang, and Emine Yilmaz (Eds.). ACM, 83–90. <https://doi.org/10.1145/3121050.3121064>
- [19] N. Jardine and Cornelis Joost van Rijsbergen. 1971. The use of hierarchic clustering in information retrieval. *Inf. Storage Retr.* 7, 5 (1971), 217–240. [https://doi.org/10.1016/0020-0271\(71\)90051-9](https://doi.org/10.1016/0020-0271(71)90051-9)
- [20] Sheng-Chieh Lin, Jheng-Hong Yang, and Jimmy Lin. 2021. In-Batch Negatives for Knowledge Distillation with Tightly-Coupled Teachers for Dense Retrieval. In *Proceedings of the 6th Workshop on Representation Learning for NLP, Repl4NLP@ACL-JCNLP 2021, Online, August 6, 2021*, Anna Rogers, Iacer Calixto, Ivan Vulic, Naomi Saphra, Nora Kassner, Oana-Maria Camburu, Trapit Bansal, and Vered Shwartz (Eds.). Association for Computational Linguistics, 163–173. <https://doi.org/10.18653/v1/2021.repl4nlp-1.17>
- [21] Jiachang Liu, Dinghan Shen, Yizhe Zhang, Bill Dolan, Lawrence Carin, and Weizhu Chen. 2022. What Makes Good In-Context Examples for GPT-3?. In *Proceedings of Deep Learning Inside Out: The 3rd Workshop on Knowledge Extraction and Integration for Deep Learning Architectures, DeelIO@ACL 2022, Dublin, Ireland and Online, May 27, 2022*, Eneko Agirre, Marianna Apidianaki, and Ivan Vulic (Eds.). Association for Computational Linguistics, 100–114. <https://doi.org/10.18653/v1/2022.deelio-1.10>
- [22] Sean MacAvaney, Craig Macdonald, and Iadh Ounis. 2022. Streamlining Evaluation with ir-measures. In *Advances in Information Retrieval - 44th European Conference on IR Research, ECIR 2022, Stavanger, Norway, April 10–14, 2022, Proceedings, Part II (Lecture Notes in Computer Science, Vol. 13186)*, Matthias Hagen, Suzan Verberne, Craig Macdonald, Christin Seifert, Krisztian Balog, Kjetil Nørvåg, and Vinay Setty (Eds.). Springer, 305–310. https://doi.org/10.1007/978-3-030-99739-7_38
- [23] Sean MacAvaney, Luca Soldaini, and Nazli Goharian. 2020. Teaching a New Dog Old Tricks: Resurrecting Multilingual Retrieval Using Zero-Shot Learning. In *Advances in Information Retrieval - 42nd European Conference on IR Research, ECIR 2020, Lisbon, Portugal, April 14–17, 2020, Proceedings, Part II (Lecture Notes in Computer Science, Vol. 12036)*, Joemon M. Jose, Emine Yilmaz, João Magalhães, Pablo Castells, Nicola Ferro, Mário J. Silva, and Flávio Martins (Eds.). Springer, 246–254. https://doi.org/10.1007/978-3-030-45442-5_31
- [24] Sean MacAvaney, Nicola Tonello, and Craig Macdonald. 2022. Adaptive Re-Ranking with a Corpus Graph. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management, Atlanta, GA, USA, October 17–21, 2022*, Mohammad Al Hasan and Li Xiong (Eds.). ACM, 1491–1500. <https://doi.org/10.1145/3511808.3557231>
- [25] Richard McCreadie, Craig Macdonald, and Iadh Ounis. 2018. Automatic Ground Truth Expansion for Timeline Evaluation. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval, SIGIR 2018, Ann Arbor, MI, USA, July 08–12, 2018*, Kevyn Collins-Thompson, Qiaozhu Mei, Brian D. Davison, Yiqun Liu, and Emine Yilmaz (Eds.). ACM, 685–694. <https://doi.org/10.1145/3209978.3210034>
- [26] Sewon Min, Xinxi Lyu, Ari Holtzman, Mikel Artetxe, Mike Lewis, Hamaneh Hajishirzi, and Luke Zettlemoyer. 2022. Rethinking the Role of Demonstrations: What Makes In-Context Learning Work?. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7–11, 2022*, Yoav Goldberg, Zornitsa Kozareva, and Yue Zhang (Eds.). Association for Computational Linguistics, 11048–11064. <https://aclanthology.org/2022.emnlp-main.759>
- [27] Alistair Moffat, Peter Bailey, Falk Scholer, and Paul Thomas. 2017. Incorporating User Expectations and Behavior into the Measurement of Search Effectiveness. *ACM Trans. Inf. Syst.* 35, 3 (2017), 24:1–24:38. <https://doi.org/10.1145/3052768>

- [28] Alistair Moffat, William Webber, and Justin Zobel. 2007. Strategic system comparisons via targeted relevance judgments. In *SIGIR 2007: Proceedings of the 30th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, Amsterdam, The Netherlands, July 23-27, 2007*, Wessel Kraaij, Arjen P. de Vries, Charles L. A. Clarke, Norbert Fuhr, and Noriko Kando (Eds.). ACM, 375–382. <https://doi.org/10.1145/1277741.1277806>
- [29] Alistair Moffat and Justin Zobel. 2008. Rank-biased precision for measurement of retrieval effectiveness. *ACM Trans. Inf. Syst.* 27, 1 (2008), 2:1–2:27. <https://doi.org/10.1145/1416950.1416952>
- [30] Tri Nguyen, Mir Rosenberg, Xia Song, Jianfeng Gao, Saurabh Tiwary, Rangan Majumder, and Li Deng. 2016. MS MARCO: A Human Generated MACHine Reading Comprehension Dataset. In *Proceedings of the Workshop on Cognitive Computation: Integrating neural and symbolic approaches 2016 co-located with the 30th Annual Conference on Neural Information Processing Systems (NIPS 2016), Barcelona, Spain, December 9, 2016 (CEUR Workshop Proceedings, Vol. 1773)*, Tarek Richard Besold, Antoine Bordes, Artur S. d'Ávila Garcez, and Greg Wryns (Eds.). CEUR-WS.org. https://ceur-ws.org/Vol-1773/CoCoNIPS_2016_paper9.pdf
- [31] Rabia Nuray and Fazli Can. 2006. Automatic ranking of information retrieval systems using data fusion. *Inf. Process. Manag.* 42, 3 (2006), 595–614. <https://doi.org/10.1016/j.ipm.2005.03.023>
- [32] Ronak Pradeep, Rodrigo Frassetto Nogueira, and Jimmy Lin. 2021. The Expando-Mono-Duo Design Pattern for Text Ranking with Pretrained Sequence-to-Sequence Models. CoRR abs/2101.05667 (2021). arXiv:2101.05667 <https://arxiv.org/abs/2101.05667>
- [33] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer. *J. Mach. Learn. Res.* 21 (2020), 140:1–140:67. <http://jmlr.org/papers/v21/20-074.html>
- [34] Navid Rekasaz, Oleg Lesota, Markus Schedl, Jon Brasseur, and Carsten Eickhoff. 2021. TripClick: The Log Files of a Large Health Web Search Engine. In *SIGIR '21: The 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, Virtual Event, Canada, July 11-15, 2021*, Fernando Diaz, Chirag Shah, Torsten Suel, Pablo Castells, Rosie Jones, and Tetsuya Sakai (Eds.). ACM, 2507–2513. <https://doi.org/10.1145/3404835.3463242>
- [35] Kevin Roitero, Andrea Brunello, Giuseppe Serra, and Stefano Mizzaro. 2020. Effectiveness evaluation without human relevance judgments: A systematic analysis of existing methods and of their combinations. *Inf. Process. Manag.* 57, 2 (2020), 102149. <https://doi.org/10.1016/j.ipm.2019.102149>
- [36] Tetsuya Sakai. 2007. Alternatives to Bpref. In *SIGIR 2007: Proceedings of the 30th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, Amsterdam, The Netherlands, July 23-27, 2007*, Wessel Kraaij, Arjen P. de Vries, Charles L. A. Clarke, Norbert Fuhr, and Noriko Kando (Eds.). ACM, 71–78. <https://doi.org/10.1145/1277741.1277756>
- [37] Tetsuya Sakai and Noriko Kando. 2008. On information retrieval metrics designed for evaluation with incomplete relevance assessments. *Inf. Retr.* 11, 5 (2008), 447–470. <https://doi.org/10.1007/s10791-008-9059-7>
- [38] David P. Sander and Laura Dietz. 2021. EXAM: How to Evaluate Retrieve-and-Generate Systems for Users Who Do Not (Yet) Know What They Want. In *Proceedings of the Second International Conference on Design of Experimental Search & Information Retrieval Systems, Padova, Italy, September 15-18, 2021 (CEUR Workshop Proceedings, Vol. 2950)*, Omar Alonso, Stefano Marchesin, Marc Najork, and Gianmaria Silvello (Eds.). CEUR-WS.org, 136–146. <https://ceur-ws.org/Vol-2950/paper-16.pdf>
- [39] Victor Sanh, Albert Webson, Colin Raffel, Stephen H. Bach, Lintang Sutawika, Zaid Alyafeai, Antoine Chaffin, Arnaud Stiegler, Arun Raja, Manan Dey, M Saiful Bari, Canwen Xu, Urmish Thakker, Shanya Sharma Sharma, Eliza Szczechla, Taewoon Kim, Gunjan Chhablani, Nihal V. Nayak, Debajyoti Datta, Jonathan Chang, Mike Tian-Jian Jiang, Han Wang, Matteo Manica, Sheng Shen, Zheng Xin Yong, Harshit Pandey, Rachel Bawden, Thomas Wang, Trishala Neeraj, Jos Rozen, Abheesh Sharma, Andrea Santilli, Thibault Févry, Jason Alan Fries, Ryan Teehan, Teven Le Scao, Stella Biderman, Leo Gao, Thomas Wolf, and Alexander M. Rush. 2022. Multitask Prompted Training Enables Zero-Shot Task Generalization. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net. <https://openreview.net/forum?id=9Vrb9D0W14>
- [40] Ian Soboroff, Charles K. Nicholas, and Patrick Cahan. 2001. Ranking Retrieval Systems without Relevance Judgments. 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*, W. Bruce Croft, David J. Harper, Donald H. Kraft, and Justin Zobel (Eds.). ACM, 66–73. <https://doi.org/10.1145/383952.383961>
- [41] Karen Sparck-Jones and C. J. Van Rijsbergen. 1975. Report on the need for and provision of an 'ideal' information retrieval test collection. *Computer Laboratory, University of Cambridge, British Library Research and Development Report No. 5266* (1975). https://sigir.org/files/museum/pub-14/pub_14.pdf
- [42] Anselm Spoerri. 2007. Using the structure of overlap between search results to rank retrieval systems without relevance judgments. *Inf. Process. Manag.* 43, 4 (2007), 1059–1070. <https://doi.org/10.1016/j.ipm.2006.09.009>
- [43] Nandan Thakur, Nils Reimers, Andreas Rücklé, Abhishek Srivastava, and Iryna Gurevych. 2021. BEIR: A Heterogenous Benchmark for Zero-shot Evaluation of Information Retrieval Models. CoRR abs/2104.08663 (2021). arXiv:2104.08663 <https://arxiv.org/abs/2104.08663>
- [44] Ellen M. Voorhees. 2001. The Philosophy of Information Retrieval Evaluation. In *Evaluation of Cross-Language Information Retrieval Systems, Second Workshop of the Cross-Language Evaluation Forum, CLEF 2001, Darmstadt, Germany, September 3-4, 2001, Revised Papers (Lecture Notes in Computer Science, Vol. 2406)*, Carol Peters, Martin Braschler, Julio Gonzalo, and Michael Kluck (Eds.). Springer, 355–370. https://doi.org/10.1007/3-540-45691-0_34
- [45] Ellen M. Voorhees, Nick Craswell, and Jimmy Lin. 2022. Too Many Relevantants: Whither Cranfield Test Collections?. In *SIGIR '22: The 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, Madrid, Spain, July 11 - 15, 2022*, Enrique Amigó, Pablo Castells, Julio Gonzalo, Ben Carterette, J. Shane Culpepper, and Gabriella Kazai (Eds.). ACM, 2970–2980. <https://doi.org/10.1145/3477495.3531728>
- [46] 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
- [47] Thuy Vu and Alessandro Moschitti. 2021. AVA: an Automatic eVALuation Approach for Question Answering Systems. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2021, Online, June 6-11, 2021*, Kristina Toutanova, Anna Rumshisky, Luke Zettlemoyer, Dilek Hakkani-Tür, Iz Beltagy, Steven Bethard, Ryan Cotterell, Tanmoy Chakraborty, and Yichao Zhou (Eds.). Association for Computational Linguistics, 5223–5233. <https://doi.org/10.18653/v1/2021.naacl-main.412>
- [48] Yizhong Wang, Swaroop Mishra, Pegah Alipoormolabashi, Yeganeh Kordi, Amirreza Mirzaei, Atharva Naik, Arjun Ashok, Arut Selvan Dhanasekaran, Anjana Arunkumar, David Stap, Eshaan Pathak, Giannis Karamanolakis, Haizhi Gary Lai, Ishan Purohit, Ishani Mondal, Jacob Anderson, Kirby Kuznia, Krima Doshi, Kuntal Kumar Pal, Maitreya Patel, Mehrad Moradshahi, Mihir Parmar, Mirali Purohit, Neeraj Varshney, Phani Rohitha Kaza, Pulkit Verma, Ravsehaj Singh Puri, Rushang Karia, Savan Doshi, Shailaja Keyur Sampat, Siddhartha Mishra, Sujan Reddy A, Sumanta Patro, Tanay Dixit, and Xudong Shen. 2022. Super-NaturalInstructions: Generalization via Declarative Instructions on 1600+ NLP Tasks. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7-11, 2022*, Yoav Goldberg, Zornitsa Kozareva, and Yue Zhang (Eds.). Association for Computational Linguistics, 5085–5109. <https://aclanthology.org/2022.emnlp-main.340>
- [49] Jason Wei, Maarten Bosma, Vincent Y. Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M. Dai, and Quoc V. Le. 2022. Finetuned Language Models are Zero-Shot Learners. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net. <https://openreview.net/forum?id=gEzrGCozdQR>
- [50] Shengli Wu and Fabio Crestani. 2002. Data fusion with estimated weights. In *Proceedings of the 2002 ACM CIKM International Conference on Information and Knowledge Management, McLean, VA, USA, November 4-9, 2002*. ACM, 648–651. <https://doi.org/10.1145/584792.584908>
- [51] Shengli Wu and Fabio Crestani. 2003. Methods for Ranking Information Retrieval Systems Without Relevance Judgments. In *Proceedings of the 2003 ACM Symposium on Applied Computing (SAC), March 9-12, 2003, Melbourne, FL, USA*, Gary B. Lamont, Hisham Haddad, George A. Papadopoulos, and Brajendra Panda (Eds.). ACM, 811–816. <https://doi.org/10.1145/952532.952693>
- [52] Emine Yilmaz and Javed A. Aslam. 2006. Estimating average precision with incomplete and imperfect judgments. In *Proceedings of the 2006 ACM CIKM International Conference on Information and Knowledge Management, Arlington, Virginia, USA, November 6-11, 2006*, Philip S. Yu, Vassilis J. Tsotras, Edward A. Fox, and Bing Liu (Eds.). ACM, 102–111. <https://doi.org/10.1145/1183614.1183633>
- [53] Zihao Zhao, Eric Wallace, Shi Feng, Dan Klein, and Sameer Singh. 2021. Calibrate Before Use: Improving Few-shot Performance of Language Models. In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event (Proceedings of Machine Learning Research, Vol. 139)*, Marina Meila and Tong Zhang (Eds.). PMLR, 12697–12706. <http://proceedings.mlr.press/v139/zhao21c.html>
- [54] Justin Zobel. 1998. How Reliable Are the Results of Large-Scale Information Retrieval Experiments?. In *SIGIR '98: Proceedings of the 21st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, August 24-28 1998, Melbourne, Australia*, W. Bruce Croft, Alistair Moffat, C. J. van Rijsbergen, Ross Wilkinson, and Justin Zobel (Eds.). ACM, 307–314. <https://doi.org/10.1145/290941.291014>