

Reliable Confidence Intervals for Information Retrieval Evaluation Using Generative A.I.

Harrie Oosterhuis*^{†‡}
Google Research
Radboud University
Amsterdam, The Netherlands
harrie.oosterhuis@ru.nl

Rolf Jagerman*[†]
Google Research
Amsterdam, The Netherlands
jagerman@google.com

Zhen Qin[†]
Google Research
New York, US
zhenqin@google.com

Xuanhui Wang[†]
Google Research
Mountain View, US
xuanhui@google.com

Michael Bendersky[†]
Google Research
Mountain View, US
bemike@google.com

Abstract

The traditional evaluation of information retrieval (IR) systems is generally very costly as it requires manual relevance annotation from human experts. Recent advancements in generative artificial intelligence –specifically large language models (LLMs)– can generate relevance annotations at an enormous scale with relatively small computational costs. Potentially, this could alleviate the costs traditionally associated with IR evaluation and make it applicable to numerous low-resource applications. However, generated relevance annotations are not immune to (systematic) errors, and as a result, directly using them for evaluation produces unreliable results.

In this work, we propose two methods based on prediction-powered inference and conformal risk control that utilize computer-generated relevance annotations to place reliable confidence intervals (CIs) around IR evaluation metrics. Our proposed methods require a small number of reliable annotations from which the methods can statistically analyze the errors in the generated annotations. Using this information, we can place CIs around evaluation metrics with strong theoretical guarantees. Unlike existing approaches, our conformal risk control method is specifically designed for ranking metrics and can vary its CIs per query and document. Our experimental results show that our CIs accurately capture both the variance and bias in evaluation based on LLM annotations, better than the typical empirical bootstrapping estimates. We hope our contributions bring reliable evaluation to the many IR applications where this was traditionally infeasible.

CCS Concepts

• **Information systems** → **Evaluation of retrieval results**; • **Computing methodologies** → *Semi-supervised learning settings*.

*Authors contributed equally to this work.

[†]Now at Google DeepMind.

[‡]Work done while Harrie Oosterhuis was working at Google 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 third-party components of this work must be honored. For all other uses, contact the owner/author(s).
KDD '24, August 25–29, 2024, Barcelona, Spain
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0490-1/24/08
<https://doi.org/10.1145/3637528.3671883>

Keywords

Information Retrieval Evaluation, Large Language Models, Confidence Intervals, Generative A.I., Conformal Prediction

ACM Reference Format:

Harrie Oosterhuis, Rolf Jagerman, Zhen Qin, Xuanhui Wang, and Michael Bendersky. 2024. Reliable Confidence Intervals for Information Retrieval Evaluation Using Generative A.I.. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '24)*, August 25–29, 2024, Barcelona, Spain. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3637528.3671883>

1 Introduction

The evaluation of information retrieval (IR) systems is an important and long-established part of the IR field [30, 71, 73]. The goal of standard IR systems is to retrieve and rank documents according to their relevance to a query and user. Accordingly, standard IR evaluation metrics (e.g., precision, recall, discounted cumulative gain (DCG), etc.) measure how relevant the top ranked items are for a set of known queries [12, 35, 36]. Accordingly, traditional evaluation requires a dataset with examples of documents, queries and annotations that indicate the relevance of documents to queries [39, 40, 57, 69]. Whilst documents and queries are often gathered by logging user interactions, relevance annotations are traditionally created through the labour of human experts, who are trained for the specific labelling task [6, 17, 30, 47]. Consequently, creating a new dataset for IR evaluation purposes is generally very costly, and as a result, no large datasets have been created for many IR settings [19, 34, 63, 74]. Thus, for these low-resource settings traditional evaluation is not available in practice.

Despite the large costs involved, there has been a continuous effort, often driven by initiatives like TREC and CLEF, to create new datasets for different IR tasks [9, 14, 31, 38, 43, 52–54, 62, 65, 67, 68, 70]. Since the foundational Cranfield collection [67], many datasets have been created for ad-hoc retrieval [31, 38, 52, 70]. However, to match the large variety of IR-related tasks, many other datasets were subsequently introduced, accordingly; For example, datasets with numerical IR features for learning-to-rank [14, 53, 54], or large collections of natural language question-answering examples such as MS MARCO [47] and BioASQ [65]. Similarly, recent years have seen the introduction of the TREC-DL [17], BEIR [62]

and Istella22 [20], among others [9, 68], specifically for the evaluation of neural IR systems. Consequently, most of the subsequent advancements in neural IR were only possible because of the availability of these datasets and the reliable benchmarking that they enable [29, 45, 46, 48]. This highlights the importance and impact of reliable evaluation on the IR field [30, 57, 69].

It is thus no surprise that potential novel data sources are received with great excitement. In the past two years, new advancements in generative artificial intelligence [18, 28, 37, 41, 51], especially the arrival of large language models (LLMs) [11, 49, 61], are speculated to bring potentially groundbreaking sources for IR evaluation [15, 25]. LLMs are trained on extremely large corpora of diverse texts for the task of generating fluent natural language [42, 60]. Importantly, LLMs are also capable at performing numerous miscellaneous tasks such as question-answering, text-summarization and text-annotation [1, 15, 27, 61, 64]. Compared to annotation by human experts, annotation via LLMs can be performed relatively cheaply and at much larger scale [44, 63]. Several existing studies have already investigated the application of LLM-generated relevance annotations to IR evaluation [15, 25, 44]. In particular, Thomas et al. [63] found that, when applied correctly to specific settings, LLMs can produce better labels than third-party assessors at a fraction of the costs. Thus, there is a clear potential for IR evaluation based on computer-generated relevance annotations.

However, a fundamental issue with evaluation based on LLMs, or other generative models, is that they are bound to make errors [8, 15, 63]. Part of these errors are coincidental, since perfect relevance prediction is infeasible in practice, but other errors are systematic [25]. For instance, an LLM could systematically misestimate relevance in certain domains or on documents with particular attributes [8, 63]. In turn, these errors could affect the final evaluation metrics and result in incorrect assessments of performance. Unfortunately, generative models cannot give trustworthy insight into their own reliability [37, 51]. Thus when solely relying on LLM-based evaluation, one cannot be certain how reliable their conclusions are.

In this work, we investigate how computer-generated relevance annotations can be used for reliable evaluation, by constructing confidence intervals (CIs) around ranking metrics with them [56, 73]. Our approach requires a small number of reliable *ground truth* annotations, in order to statistically analyze the distribution of errors that exist in the generated annotations. Subsequently, we apply two state-of-the-art methodologies [3, 4] with a strong theoretical grounding to find reliable CIs. In this work, we provide two main methodological contributions:

Our first contribution is the novel application of *prediction-powered inference* (PPI) to IR evaluation [3]. PPI applies classical methods for building CIs but builds them around the error between the predicted and true values of a metric. Thereby, somewhat-reasonable predictions can lead to substantially smaller CIs than classical CI around just the metric value. The limitations of PPI is that it does not utilize the uncertainty of the generative model, and that it only provides a CI around the final metric value.

Our second contribution addresses these limitations by proposing a novel *conformal risk control* (CRC) approach [2, 4]. We introduce a novel method to place an optimistic and a pessimistic prediction around each generated relevance label, which follows

the confidence of the generative model. These predictions can be propagated to form an interval around metrics on the query or dataset level. Through CRC, our approach calibrates the intervals to guarantee that the true value lies between them with a minimum probability. In other words, our method puts lower and upper bounds around the relevance of each document that naturally translate to reliable CI on query and dataset-level metrics. Thereby, unlike PPI, our CRC approach does utilize the confidence of the generative model and can provide CI on query-level performance.

Our results on several IR benchmarks show that both our methods provide CIs around LLM-based metric predictions that accurately capture the true values, while also being significantly less wide than those of previous CI methods [16, 73]. Moreover, unlike other approaches [3], our CRC method can vary CIs per document, query or collection of queries and can thus better indicate where a generative model is more or less reliable.

To the best of our knowledge, our novel approaches are the first that leverages computer-generated relevance annotations to produce reliable CIs for IR evaluation. We hope this contribution opens up novel possibilities for reliable benchmarking of low-resource IR tasks that have been traditionally infeasible.

2 Related Work

2.1 Confidence intervals for IR evaluation

Evaluation is a well-established core part of the IR field [30, 52, 67, 71, 73]. Generally, it aims to measure how well a retrieval system can produce a list of ranked documents in response to a user query [30, 58, 69]. The most prevalent form of IR evaluation relies on datasets containing example queries, documents and human-annotated relevance labels [30, 57, 67, 69]. Accordingly, there is a long history of efforts to create such datasets in the IR community, such as TREC [17, 31, 68, 70, 71], CLEF [52], NTCIR [38] and many others [9, 14, 20, 43, 53, 54, 62, 65, 67]. Despite the enormous importance of these datasets, they are known to have limitations. For instance, expert annotators can give conflicting relevance assessments, and the actual users of an IR application can disagree with the experts as well [57]. Furthermore, the constructions of these datasets is often costly which puts constraints on their size [13, 14, 67, 70]. As a result, IR datasets can only represent a limited slice of the queries that a real IR system receives [13, 73].

Accordingly, statistical approaches to IR evaluation have been developed to deal with these limitations. For example, it has become common practice to use significance tests to ensure that observed differences in IR metrics are, with high probability, not the result of random chance [26, 59, 66]. Confidence intervals (CIs) have been used to express the uncertainty that comes from using the dataset sample of queries to estimate performance over all queries [16, 56]. Furthermore, previous work has also applied CI for relevance annotator disagreement [21, 33] and missing relevance annotations [5, 72, 75]. The statistical methods used to construct CI by previous work in IR have been based on empirical bootstrapping techniques [22, 23, 32]. To the best of our knowledge, our work is the first to consider PPI and CRC methods for IR evaluation [3, 4].

2.2 LLMs for relevance annotation generation

Recent advances in LLMs have demonstrated impressive capabilities on a broad range of tasks [1, 27, 61, 64]. Previous work has specifically considered using LLMs for relevance annotation in an IR context [15, 25, 44, 63]. Thomas et al. [63] propose using ground truth relevance labels from human annotators, to find a prompt that results in the most accurate LLM generated labels. They claim that this method produces relevance annotations at the same quality as third-party human assessors but at a fraction of the costs [63]. Clarke et al. [15] propose that LLM relevance-annotation should be approached as a spectrum, since the involvement of humans can be varied. For instance, one could delegate most work to an LLM but add some human verification, as a compromise of reduced costs and reliability. Faggioli et al. [25] support this approach, as they see severe risk in blindly following LLM generated relevance labels (at least for the current state-of-the-art LLMs). The danger foreseen by both is that generated labels can make systematic errors that lead to incorrect and unreliable evaluation of IR systems [8, 15, 25]. Our work addresses this problem, and is thus very related; specifically, our contribution can be seen as an approach of human verification designed to quantify uncertainty stemming from LLM usage.

3 Preliminaries

3.1 Evaluation metrics for retrieval systems

The general approach to the evaluation of a retrieval system is to consider the expected value of a ranking metric across the queries it will receive [30]. Standard ranking metrics assume that each document has certain relevance to a query [39]. For a set of labels \mathcal{R} , we use $P(R = r | d, q)$ to denote the probability that a human rater would give rating $r \in \mathcal{R}$, to the combination of document d and query q . We define *relevance* as the expected rating value over this distribution: $\mu(d | q) = \sum_{r \in \mathcal{R}} P(R = r | d, q) r$. In standard ranking settings, the goal is to place more relevant documents at higher ranks [36]. Ranking metrics capture this goal by giving a weight to each rank, which indicates how much the relevance of a document placed at that rank should contribute to the metric [35]. We will use ω to denote our weighting function which takes the rank of a document as its input. For example, *Precision@K* has the following corresponding weight function: $\omega_{\text{Prec@K}}(x) = \frac{1}{K} \mathbb{1}[x \leq K]$; and the popular DCG [35]: $\omega_{\text{DCG@K}}(x) = \frac{\mathbb{1}[x \leq K]}{\log_2(x+1)}$. Given a choice of weights and \mathcal{D}_q , the set of available documents for query q , the metric value for a single query is:

$$U(q) = \sum_{d \in \mathcal{D}_q} \omega(\text{rank}(d | q, \mathcal{D}_q)) \mu(d | q). \quad (1)$$

Let $P(q)$ denote the natural query distribution; the performance of a system in terms of the metric is:

$$U(Q) = \mathbb{E}_{q \sim Q}[U(q)] = \sum_{q \in Q} P(q | Q) U(q). \quad (2)$$

In practice, U can never be computed exactly, since $P(q)$ and $P(R = r | d, q)$ are never directly available. Thus, generally, an estimate of U is made on a large set of sampled user queries and a few relevance judgements per document-query pair [57, 73].

3.2 Problem setting

In our setting, we make the standard assumption that a large set of sampled user queries and a document collection are available. However, we do not assume that there are human relevance annotations for every document-query pair, and instead, we assume that ground truth annotations are only available for a small subset: the first n queries out of a total of N queries. Unique to our problem setting is that a generative model is available to predict relevance annotations. Furthermore, our aim is not to give a point estimate of the true performance of a system, instead our goal is to construct a reliable CI around the true value of an IR metric. Thereby, we utilize the generated relevance annotations, but still explicitly indicate the resulting uncertainty in our evaluation with CIs.

In formal terms, let $\alpha \in [0, 1]$ be a confidence parameter, we desire to find a lower bound \hat{U}_{low} and an upper bound \hat{U}_{high} s.t:

$$P(\hat{U}_{\text{low}} \leq U(Q) \leq \hat{U}_{\text{high}}) \geq 1 - \alpha. \quad (3)$$

Accordingly, α can be chosen to match the desired confidence, i.e., $\alpha = 0.05$ leads to a 95% CI. Additionally, in Section 6, we propose a CRC method that can also bound the performance per query, thereby, it can meet the following query-level CI goal:

$$P(\hat{U}_{\text{low}}(q) \leq U(q) \leq \hat{U}_{\text{high}}(q) | q \sim Q) \geq 1 - \alpha. \quad (4)$$

We assume that the available generative model predicts a distribution over possible relevance labels per query-document pair [44, 63]. Let $\hat{P}(R = r | d, q)$ indicate the predicted probability for relevance value r for the combination of document d and query q , the mean predicted relevance is then:

$$\hat{\mu}(d) = \sum_{r \in \mathcal{R}} \hat{P}(R = r | d, q) r. \quad (5)$$

Using these predicted relevances, we can construct a prediction of performance on the dataset-level from a sampled set of queries Q . This results in the following predicted metric value:

$$\hat{U}(Q) = \frac{1}{|Q|} \sum_{q \in Q} \sum_{d \in \mathcal{D}_q} \omega(\text{rank}(d | q, \mathcal{D}_q)) \hat{\mu}(d | q). \quad (6)$$

As discussed in previous work [15, 25], basing $\hat{U}(Q)$ on state-of-the-art LLMs could greatly reduce costs [63], but there are many risks involved in replacing human annotators [8]. The accuracy of $\hat{U}(Q)$ completely depends on the predictive capabilities of the generative model. Thus, without further knowledge about the reliability of the predictions, one has no indication of its trustworthiness. Our proposed methodologies use the available n ground truth query-level performances together with the many generated relevance predictions to construct reliable CIs that quantify these risks.

4 Method 1: Prediction-Powered Inference for Information Retrieval Evaluation

Our first proposed method applies the *prediction-powered inference* (PPI) framework to IR evaluation. PPI is a very recent advancement in CI construction introduced by Angelopoulos et al. [3]. It utilizes computer-generated predictions to create smaller CI when these predictions are somewhat accurate. The core idea is to avoid estimating a variable on labelled data directly, and instead, build an estimate around the predictions which is then corrected based on

the labelled data [10]. If the predictions are found to be accurate on the labelled data, then this increases our confidence that its predictions on unlabelled data are also accurate. As predictions become available in much larger quantities, this can increase our confidence further in the overall estimate. To the best of our knowledge, we are the first to apply PPI to IR evaluation.

4.1 Classical empirical mean estimation

Before we detail our application of PPI, it is easiest to start with classical empirical estimates. As stated in Section 3.2, our aim is to place a reliable CI around the true performance $U(Q)$, and relevance annotations are available for the first n queries in Q . Therefore, we can make an empirical estimate of the mean metric performance based on these queries:

$$\hat{U}^{\text{emp}}(Q) = \frac{\sum_{i=1}^n U(q_i)}{n}, \quad \hat{\sigma}_{\text{emp}}^2 = \frac{\sum_{i=1}^n (\hat{U}^{\text{emp}}(Q) - U(q_i))^2}{n-1}, \quad (7)$$

where $\hat{\sigma}_{\text{emp}}^2$ is the estimated variance of the empirical estimate, and $U(q_i)$ the metric value for the single query q_i (Eq. 1). We note that its variance is solely reflective of the ground truth data. Obviously, this estimate does not fully utilize our problems setting, as it ignores the queries without ground truth relevance annotations and their corresponding computer-generated relevance annotations.

4.2 Prediction-powered inference

In contrast, PPI mean estimation combines ground truth and predicted values to create an estimator that has potentially much lower variance. In our setting, the PPI estimator is a combination of the estimated mean predicted query performance and the estimated mean prediction error:

$$\hat{U}^{\text{PPI}}(Q) = \underbrace{\frac{\sum_{i=1}^N \hat{U}(q_i)}{N}}_{\text{mean prediction}} + \underbrace{\frac{\sum_{i=1}^n U(q_i) - \hat{U}(q_i)}{n}}_{\text{mean prediction error}}. \quad (8)$$

In other words, PPI constructs an estimate of the query performance based on the predicted relevance annotations, and corrects it by the estimated error based on the difference between the predicted and ground truth annotations. As a result, it is unbiased:

$$\mathbb{E}_{Q \sim \mathcal{Q}}[\hat{U}^{\text{PPI}}(Q)] = \mathbb{E}_{Q \sim \mathcal{Q}}[\hat{U}^{\text{emp}}(Q)] = U(Q), \quad (9)$$

Whilst the empirical and PPI estimates have the same expected value, the key-difference is their variances. Assuming the queries are i.i.d., the estimated variance of PPI can be decomposed into a part stemming from the mean prediction and another from the prediction error:

$$\begin{aligned} \hat{\sigma}_{\text{PPI}}^2(Q) &= \hat{\sigma}_{\text{pred}}^2(Q) + \hat{\sigma}_{\text{error}}^2(Q), \\ \hat{\sigma}_{\text{pred}}^2(Q) &= \sum_{i=1}^N \frac{(\hat{U}(q_i) - \frac{1}{N} \sum_{j=1}^N \hat{U}(q_j))^2}{N-1}, \\ \hat{\sigma}_{\text{error}}^2(Q) &= \sum_{i=1}^n \frac{(U(q_i) - \hat{U}(q_i) - \frac{1}{n} \sum_{j=1}^n (U(q_j) - \hat{U}(q_j)))^2}{n-1}. \end{aligned} \quad (10)$$

This reveals how PPI can benefit from predictions and unlabelled data. We see that $\hat{\sigma}_{\text{error}}^2$ shrinks as predicted performances become

more accurate, whilst $\hat{\sigma}_{\text{pred}}^2$ shrinks as more unlabelled data becomes available (as N increases). Comparing $\hat{\sigma}_{\text{PPI}}^2$ with $\hat{\sigma}_{\text{emp}}^2$ reveals that PPI can give a lower variance estimate, but only if predictions are somewhat accurate. Conversely, when they are inaccurate the variance could actually be greater.

Finally, in order to turn the estimated mean and variance into a CI, we follow Angelopoulos et al. [3] and assume $\hat{U}^{\text{PPI}}(Q)$ follows a normal distribution. The 95% confidence interval is then:

$$\hat{U}_{\text{high/low}}(Q) = \hat{U}^{\text{PPI}}(Q) \pm 1.96 \sqrt{\frac{\hat{\sigma}_{\text{error}}^2}{n} + \frac{\hat{\sigma}_{\text{pred}}^2}{N}}. \quad (11)$$

Accordingly, one can use a different z-score than 1.96 to choose a different level of confidence. We note that this implicitly assumes the prediction error follows a symmetric distribution.

This concludes our description of our PPI method. Its biggest advantage is its simplicity and straightforward application, making it attractive for practical usage. A limitation is that it only gives a CI of the overall performance (dataset-level). Therefore, PPI cannot be used to place CI around individual query performances, and similarly, it cannot vary its confidence for different queries.

5 Background: Conformal Prediction and Conformal Risk Control

This section provides the necessary background on conformal prediction and conformal risk control (CRC) [2, 4, 7, 50], before Section 6 introduces our CRC approach for IR evaluation.

5.1 Conformal prediction

Conformal prediction provides a unique approach to uncertainty quantification in predictions [7, 50]. The key characteristic of conformal prediction is that its predictions are not individual labels but sets of labels. For instance, the most basic version of this approach constructs a prediction set C by including all labels that have a predicted probability above a threshold $\lambda \in [0, 1]$ [2]. Let \hat{P} indicate a predicted probability, X contextual features and Y a corresponding label, this basic prediction set is then:

$$C_{\text{basic}}(X | \lambda) = \{y : \hat{P}(Y = y | X) > \lambda\}. \quad (12)$$

Given a set of i.i.d. calibration data, conformal prediction can set λ so that C_{basic} contains the true label with high probability:

$$P(Y \in C_{\text{basic}}(X | \lambda)) > 1 - \alpha. \quad (13)$$

Thereby, C_{basic} can capture the uncertainty in the prediction of X , with strong theoretical guarantees, when applied to the same distribution from which the calibration data was sampled [2].

5.2 Conformal risk control

For purposes other than label prediction, there is a more general version of this approach: *conformal risk control* (CRC) [4]. Let $C(X | \lambda)$ be an arbitrary function that constructs sets that increase with λ , \mathcal{L} a bounded loss function that shrinks as C grows, and in this context $\alpha \in \mathbb{R}$, CRC aims to guarantee the expected loss is bounded:

$$\mathbb{E}_{(x,y) \sim \mathcal{P}(X,Y)}[\mathcal{L}(C(X = x | \lambda), Y = y)] < 1 - \alpha. \quad (14)$$

We can see that this is a generalized version of conformal prediction, since it is equivalent to Eq. 13 when: $\mathcal{L}(C, Y) = \mathbb{1}[Y \notin C]$ [2].




Figure 1: Three different predicted relevance distributions (left) and their corresponding $\hat{\mu}_{CRC}(d, \lambda)$ curves (right).

CRC can guarantee Eq. 14, by finding a value of λ based on a set of n i.i.d. calibration data-points such that:

$$\frac{1}{n} \sum_{i=1}^n \mathcal{L}(C(X_i | \lambda), Y_i) < \alpha - \frac{B - \alpha}{n}, \quad (15)$$

where B is the maximum possible value of \mathcal{L} . Under the assumption that the calibration data was sampled from the same distribution $(P(X, Y))$, CRC is proven to provide the bound guarantee stated in Eq. 14 [2, 4]. We note that it is possible that no λ value exists that can satisfy Eq. 15 because the number of data-points n is too small. In this case, the method explicitly fails to provide a CI, thereby, CRC indicates when it is unable to guarantee reliable CIs. The generality and flexibility of the CRC framework enables us to build our own CI method for IR evaluation on top of it.

6 Method 2: Conformal Risk Control for Information Retrieval Evaluation

Our second proposed method uses *conformal risk control* (CRC) for CI construction [4]. In contrast with PPI, it can provide both CI around mean performance and per query performance. It also relies on different assumptions than PPI and empirical bootstrapping.

Our description of the method is divided into three parts: firstly, we introduce our C function, secondly, we describe how calibration data is gathered, and thirdly, we propose our alternative dual-calibration approach specific for CIs.

6.1 Optimistic and pessimistic estimation

For our purposes, C will construct CIs for the relevance of each individual document, that are then translated into CIs on query and dataset-level performance. Thus, our CRC method treats each CI as a set that includes all values between its minimum and maximum. Accordingly, we must predict the boundaries of CIs on a document-level, therefore, we propose two functions: $\hat{\mu}_{high}$ and $\hat{\mu}_{low}$, that provide more optimistic and pessimistic predictions than $\hat{\mu}$, respectively. We wish the optimism/pessimism to follow the confidence of the generative model, thus, we take the predicted distribution \hat{P} and remove λ probability from the top or bottom labels:¹

$$\begin{aligned} \hat{Q}_{high}(R=r|d, \lambda) &= \hat{P}(R=r|d) - \max\left(0, \lambda - \sum_{r' \in \mathcal{R}: r' < r} \hat{P}(R=r'|d)\right), \\ \hat{Q}_{low}(R=r|d, \lambda) &= \hat{P}(R=r|d) - \max\left(0, \lambda - \sum_{r' \in \mathcal{R}: r' > r} \hat{P}(R=r'|d)\right). \end{aligned} \quad (16)$$

We note that when λ is greater than the predicted probability for the lowest/highest label, the remainder is subtracted from the next lowest/highest label, and so forth. The results are normalized to

¹For brevity, we omit q from our notation: $\hat{P}(R=r|d, q) = \hat{P}(R=r|d)$.

produce the valid probability distributions; \hat{P}_{high} and \hat{P}_{low} :

$$\hat{P}_{high/low}(R=r|d, \lambda) = \frac{\hat{Q}_{high/low}(R=r|d, \lambda)}{\sum_{r' \in \mathcal{R}} \hat{Q}_{high/low}(R=r'|d, \lambda)}. \quad (17)$$

Due to possible bias in the predicted relevance annotations, e.g., all predictions could be severe over or underestimates, we want to enable both boundaries of CIs to be optimistic or pessimistic. For elegance, we let $\lambda \in (-1, 1)$ and our perturbed distribution is either optimistic or pessimistic based on the sign of λ :

$$\hat{P}_{CRC}(R=r|d, \lambda) = \begin{cases} \hat{P}_{high}(R=r|d, \lambda) & \text{if } \lambda \geq 0, \\ \hat{P}_{low}(R=r|d, -\lambda) & \text{otherwise.} \end{cases} \quad (18)$$

The final optimistic or pessimistic estimates are the expected values over these perturbed distributions:

$$\hat{\mu}_{CRC}(d, \lambda) = \sum_{r \in \mathcal{R}} \hat{P}_{CRC}(R=r|d, \lambda) r. \quad (19)$$

Figure 1 visualizes how $\hat{\mu}_{CRC}$ varies over different λ values for three different predicted relevance distributions. We see that low predicted probabilities for the largest labels mean that λ has to be greater for $\hat{\mu}_{CRC}$ to reach high values, and vice versa, λ has to be lower for low probabilities for the lowest label values to reach low values. In other words, it takes more extreme λ values for $\hat{\mu}_{CRC}$ to be heavily optimistic when the generative model is very confidently pessimistic, and vice versa.

The document-level $\hat{\mu}_{CRC}$ are translated to performance estimates following Eq. 1 & 6 but with $\hat{\mu}(d)$ replaced by $\hat{\mu}_{CRC}(d, \lambda)$. Finally, to construct CIs, we use two parameters: $\lambda_{high} \in (-1, 1)$ and $\lambda_{low} \in (-1, 1)$, s.t. $\lambda_{low} < \lambda_{high}$, to obtain $\hat{U}_{CRC}(Q, \lambda_{low})$ and $\hat{U}_{CRC}(Q, \lambda_{high})$. The predicted CI is the range between the perturbed estimates:

$$C(Q, \lambda_{high}, \lambda_{low}) = [\hat{U}_{CRC}(Q, \lambda_{low}), \hat{U}_{CRC}(Q, \lambda_{high})]. \quad (20)$$

Our proposed C function has several significant properties that enable it to function well as CI: When $\lambda_{low} = \lambda_{high} = 0$, it only contains the predicted $\hat{U}(Q)$ value, since: $\hat{U}_{CRC}(Q, 0) = \hat{U}(Q)$. As the λ approach one and minus one, the perturbed estimates become the minimum and maximal possible metric values:

$$\lim_{\lambda_{high} \rightarrow 1, \lambda_{low} \rightarrow -1} C(Q, \lambda_{high}, \lambda_{low}) = [\max U(\cdot), \min U(\cdot)]. \quad (21)$$

Consequently, there always exists values for λ_{high} and λ_{low} to bound the true performance $U(Q)$, since it must lie between the minimal and maximal possible metric values:

$$\exists \lambda_{high} \in (-1, 1), \lambda_{low} \in (-1, \lambda_{high}]; \quad U(Q) \in C(Q, \lambda_{high}, \lambda_{low}).$$

To summarize, we have proposed a novel $C(Q, \lambda)$ function that creates a CI based on the relevance annotations of a generative model. It follows the confidence of the underlying generative model

by perturbing the predicted relevance distributions in an optimistic or pessimistic manner. The remainder of this section explains how we determine the values of λ_{high} and λ_{low} such that a reliable CI is found that captures the true metric value with high confidence.

6.2 Data sampling and bootstrapping

In order to perform CRC calibration, a set of ground truth examples is required to serve as calibration data. In our setting, we aim to estimate the mean over the true query-distribution Q based on the sampled set of queries Q . Accordingly, a set of examples of mean estimates based on sampled set from Q is required; we create M examples by sampling from the n queries in Q with ground truth annotations: $\bar{Q}_i \subset \{q_1, q_2, \dots, q_n\}$. The collection of these M sets should mimic the distribution of Q : $\bar{Q} = \{\bar{Q}_1, \bar{Q}_2, \dots, \bar{Q}_M\}$.

There are many options to construct \bar{Q} , for instance, one could sample queries with or without replacement, the size of the sampled sets could be varied, etc. Moreover, if one wants to create CIs around the performance of each query, they can choose the sets to contain a single query: $\bar{Q}_i = \{q_i\}$. Another option is to sample queries and subsets of the document to be ranked, to artificially increase the variety in candidate documents available per query. Choices that increase the number of examples M have the potential to decrease CI width. However, if the resulting \bar{Q} is no longer representative of the true distribution Q , the reliability of the CIs will decrease.

6.3 Dual-calibration for confidence intervals

With our definition of $C(Q, \lambda_{\text{high}}, \lambda_{\text{low}})$ and the calibration data \bar{Q} , all that remains is to calibrate λ_{high} and λ_{low} . However, standard CRC is designed for the calibration of a single parameter. Luckily, for the purpose of construction a CI, we can apply CRC calibration sequentially. Because for any $\hat{U}_{\text{low}} < \hat{U}_{\text{high}}$, the following holds:

$$\begin{aligned} \left(P(U \leq \hat{U}_{\text{low}}) \leq \frac{\alpha}{2} \wedge P(U \geq \hat{U}_{\text{high}}) \leq \frac{\alpha}{2} \right) \\ \longrightarrow P(\hat{U}_{\text{low}} \leq U \leq \hat{U}_{\text{high}}) \leq 1 - \alpha. \end{aligned} \quad (22)$$

Therefore, we can first calibrate one of the bounds with CRC, and the other afterwards. Accordingly, we propose two loss functions:

$$\begin{aligned} \mathcal{L}_{\text{high}}(C_{\text{CRC}}(Q, \lambda_{\text{high}}), U(Q)) &= \mathbb{1}[\hat{U}_{\text{CRC}}(Q, \lambda_{\text{high}}) < U(Q)], \\ \mathcal{L}_{\text{low}}(C_{\text{CRC}}(Q, \lambda_{\text{low}}), U(Q)) &= \mathbb{1}[\hat{U}_{\text{CRC}}(Q, \lambda_{\text{low}}) > U(Q)]. \end{aligned} \quad (23)$$

Through applying two binary search procedures, we find the values for $\lambda_{\text{high}} \in (-1, 1)$ and $\lambda_{\text{low}} \in (-1, 1)$ such that $\lambda_{\text{low}} < \lambda_{\text{high}}$ and:

$$\frac{1}{M} \sum_{i=1}^M \mathcal{L}_{\text{high/low}}(C_{\text{CRC}}(\bar{Q}_i, \lambda_{\text{high/low}}), U(\bar{Q}_i)) < \frac{1}{2} \left(\alpha - \frac{1-\alpha}{M} \right). \quad (24)$$

Consequently, according to Eq. 22, it must be the case that the CRC requirement for the complete interval holds:

$$\frac{1}{M} \sum_{i=1}^M \mathbb{1}[U(\bar{Q}_i) \in C_{\text{CRC}}(\bar{Q}_i, \lambda_{\text{high}}, \lambda_{\text{low}})] < \alpha - \frac{1-\alpha}{M}. \quad (25)$$

Therefore, the resulting CI has the desired reliability, when applied to the distribution underlying \bar{Q} :

$$P(\hat{U}_{\text{CRC}}(\bar{Q}, \lambda_{\text{low}}) \leq U(\bar{Q}) \leq \hat{U}_{\text{CRC}}(\bar{Q}, \lambda_{\text{high}}) \mid \bar{Q} \sim \bar{Q}) > 1 - \alpha. \quad (26)$$

Table 1: DCG@10 performance of different rankers as measured by human-annotated labels and LLM-generated labels. Each approach ranks the top-100 results retrieved by BM25.

	TREC-DL		Robust04	
	Human	LLM	Human	LLM
Random	3.16	6.86	0.99	2.96
BM25	8.25	12.93	2.71	4.32
LLM	12.81	23.73	3.23	7.17
Perfect	19.00	17.44	5.70	4.65

Accordingly, we assume that this (Eq. 26) implies the following:

$$P(\hat{U}_{\text{CRC}}(Q, \lambda_{\text{low}}) \leq U(Q) \leq \hat{U}_{\text{CRC}}(Q, \lambda_{\text{high}}) \mid Q \sim Q) > 1 - \alpha. \quad (27)$$

This is a very standard assumption made in CI literature, and at the core of many bootstrapping methods [22, 24]. If \bar{Q} is created by standard sampling from Q , then this is a relatively safe assumption.

6.4 Overview

Finally, we give an overview of the different components in our CRC approach: Our CI are created with the $C_{\text{CRC}}(Q, \lambda_{\text{high}}, \lambda_{\text{low}})$ function (Eq. 20), where Q are all available queries (no ground truth annotations required). We note that when the set Q contains a single query, it produces a CI for query-level performance.

The resulting CI are only reliable if λ_{high} and λ_{low} are properly calibrated. We do so by first sampling a collection of query-sets \bar{Q} (Section 6.2) and calibrating each parameter independently (Eq. 24). Due to the nature of CI (Eq. 22), this guarantees the CRC requirement is met (Eq. 25), and assuming \bar{Q} is representative of Q , this guarantees that our CI are reliable with a given probability (Eq. 27).

7 Experimental Setup

Our experiments compare the confidence intervals produced by PPI, CRC and classical empirical bootstrapping on benchmark IR datasets, by answering the following research questions:²

- RQ1:** How many human-annotated labels are required to produce informative confidence intervals?
- RQ2:** How resilient are the confidence intervals to systematic mistakes made by LLM labelers?
- RQ3:** What benefits could PPI and CRC get from potential improvements in the accuracy of label generation?
- RQ4:** Can CRC capture differences in uncertainty per query?

LLM-generated relevance labels. For each query-document pair, a prompt is constructed that asks the LLM to assess the relevance according to the relevance scales of the dataset, in our case: 0–2 (Robust04) and 0–3 (TREC-DL). The LLM is provided with clear definitions of the different relevance labels, similar to [63]. Specifically, instructions that give definitions for relevance labels in each prompt. We chose prompts that mimic the instructions for human annotators as closely as possible, hereby, we hope to

²Our experimental implementation and our dataset of generated LLM labels are available at: https://github.com/google-research/google-research/tree/master/high_confidence_ir_eval_using_genai

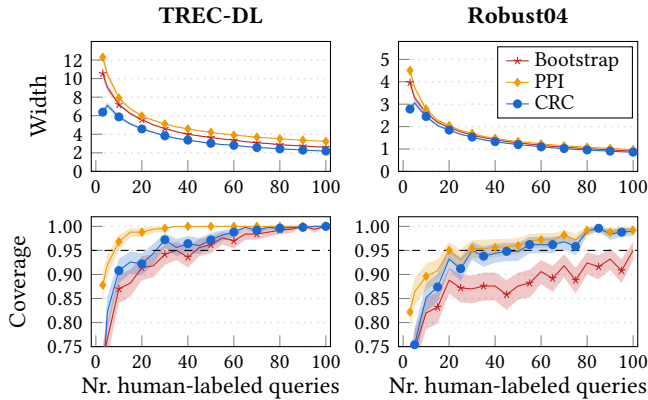


Figure 2: Width (top) and coverage (bottom) of the confidence intervals produced by the methods. The dashed line in the bottom plots is the 95% coverage target. Shaded areas indicate 95% prediction intervals over 500 independent runs.

precisely simulate the manual labeling process for each dataset. The exact prompts are provided in Appendix A.

To obtain relevance labels, we run the LLMs in ‘scoring mode’ [76].

That is, for each relevance label $r \in \mathcal{R}$, we compute the log-probability of the LLM outputting the relevance rating r . The log-probabilities are then normalized via a softmax function so that we obtain a probability distribution that represents the LLM’s confidence in assigning each relevance label to the query-document pair.

As our LLM model, we choose to use Flan-UL2 [42, 60], because it is open source and has demonstrated strong performance on ranking tasks [55]. It is worth noting that larger, more powerful, LLMs exist [61], and that we do not utilize any prompt-engineering [63]. These choices were made because the goal of our experiments is not to find the best LLM-generated labels, but to confirm whether the confidence intervals proposed by our methods accurately capture the uncertainty in LLM-generated relevance labels. Since advancements in LLM techniques result in rapid changes in the state-of-the-art, we choose to focus on the established human annotator setting instead [17, 71].

Datasets. Our evaluation is based on two established benchmark datasets: TREC-DL [17] and TREC-Robust04 [71]. Both datasets are comprised of documents and queries together with human-annotated relevance judgments. For each dataset, we perform a random 50:50 split to obtain a validation and test set where the validation set is used for calibration of the methods. (A training set is not required in our setting.) To avoid distribution shifts, for TREC-DL, we create a stratified sample over four years (2019 - 2022) that ensures each year is equally represented in each split. As the ranker to evaluate, we choose BM25, as the metric we choose DCG@10 [35]. In other words, our methods will construct CIs around the DCG@10 of BM25 on both datasets. Table 1 displays the ranking performance of BM25 and the LLM-generated labels. To match the gain function of DCG all labels were transformed accordingly: $r' = 2^r - 1$, for all performance estimations.

Methods in comparison. The methods included in our comparison are: (i) empirical bootstrapping [22], (ii) *prediction-powered inference* (PPI) (Section 4), and (iii) *conformal risk control* (CRC) (Section 6). The empirical bootstrap approach acts as a baseline




Figure 3: Width of the confidence intervals for increasing levels of LLM bias (β , top-row) and oracle-enhanced LLM accuracy ($\tau \rightarrow 1$, bottom row) with $n = 112$ on TREC-DL and $n = 125$ on Robust04. Shaded areas indicate 95% prediction intervals over 500 independent runs. Coverage plots are omitted since all methods maintain >95% coverage.

that only considers the available human-labeled data, this is a standard approach in previous IR literature [5, 56, 72, 73, 75]. All our empirical bootstrap CI are based on 10,000 bootstrap samples. PPI is computed by applying Eq. 11 to both the validation set (the first n queries) and the test set (the remaining $N - n$ queries), it utilizes both human and LLM-generated labels. Finally, our CRC approach also utilizes both, we use the validation set to calibrate the λ parameters and then compute the CI using only the LLM-generated labels on the test-set. For calibration, CRC is provided $M = 10,000$ batches each consisting of n queries that were sampled with replacement from the validation set (see Section 6.2). We note that the batch size depends on the number of available queries with human annotations, which is varied in our experiments. For the CIs to be evaluated, the CI is applied to the entire test-set to obtain a dataset-level CI, i.e., we compute $C(Q_{\text{test}}, \lambda_{\text{high}}, \lambda_{\text{low}})$ (Eq. 20). Some of our experiments consider CRC CIs around query-level performance, in these cases, λ is not calibrated on bootstrapped batches but on n batches that each contain a single query.

We evaluate the CIs produced by each method by considering their *width* and *coverage*. The width measures how wide and thus how informative or specific the CI is, where a smaller width is better. The coverage measures how frequent the CI covers the true performance on the test-set over 500 independently repeated experiment runs, thus the higher the better. The target for all the methods is a coverage of 95% or higher and we set $\alpha = 0.05$ accordingly.

8 Results

8.1 Number of required human-annotations

Our main results are displayed in Figure 2. Here we see how the width and coverage of the different methods vary, as they are provided with n queries with human annotations sampled from the validation set. As expected, all methods provides better CIs when provided with larger portion human-annotated queries, i.e., as n increases coverage increases and width decreases.

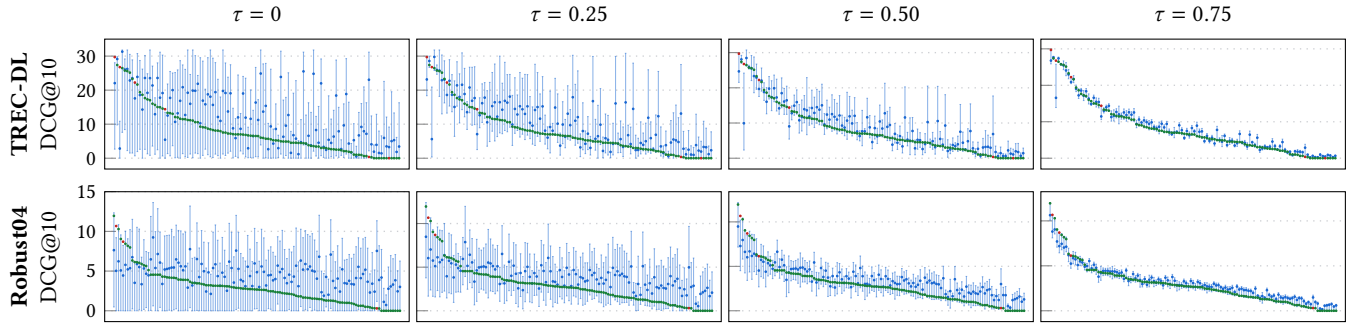


Figure 4: 95% CI produced per-query by CRC using LLM predicted relevance annotations ($\tau = 0$) and oracle-enhanced LLM annotations ($\tau > 0$). The queries are sorted by their true DCG performance (according to human-annotations), indicated by red and green dots. Green dots are covered by their CI whereas red dots are not. Blue dots indicate the predicted DCG performance (according to LLM-generated annotations). Clearly, the CIs shrink considerably as annotations become more accurate ($\tau \rightarrow 1$).

We start by considering the performance of the empirical bootstrap baseline. On the TREC-DL dataset, we see that it requires at least 40 labeled queries to achieve 95% coverage. Furthermore, on Robust04, with 100 labeled queries it almost reaches 95% coverage. However, the plotted prediction intervals around the reported coverage reveal that many of its runs did not reach 95% coverage.

In contrast, both our PPI and CRC approaches have stronger coverage with less queries: PPI needs less than 20 queries on TREC-DL and less than 40 on Robust04. Similarly, CRC needs less than 30 on TREC-DL and less than 50 on Robust04. In terms of width, CRC clearly provides the smallest width of all the methods, whilst PPI is worse than empirical bootstrap on TREC-DL and comparable on Robust04. This comparison is not entirely fair, i.e., there is generally a tradeoff between coverage and width, it appears PPI does better in terms of coverage but that results in wider CIs. Thus, PPI has a clear advantage over empirical bootstrap on Robust04 where it has the same width but much better coverage. Nevertheless, when CRC and PPI have the same coverage, CRC has smaller widths, with an especially large difference on TREC-DL. Therefore, it appears that CRC has the most informative CI, whilst PPI needs fewer queries to reach 95% coverage. Both methods provide substantial improvements over empirical bootstrapping.

Thus we answer RQ1 as follows: both PPI and CRC require as few as 30 human-labeled queries to produce informative and reliable confidence intervals. Whilst empirical bootstrapping requires significantly more human-labeled queries to achieve similar results.

8.2 Sensitivity to LLM accuracy

Our PPI and CRC methods can benefit from accurate LLM labels, but in order to be reliable, it is also important that they are robust to inaccurate labels. We investigate the effect of LLM accuracy by adding adversarial bias to the predicted relevance distributions, with $\beta \in [0, 1]$, change the predictions as follows: $\hat{P}_\beta(R = r | d, q) = \frac{1}{Z} \left((1 - \beta) \hat{P}(R = r | d, q) + \beta (1 - \hat{P}(R = r | d, q)) \right)$, where Z is a normalizing factor to ensure the result is a valid probability distribution. For $\beta = 0$ this leaves predictions unaltered, with $\beta = 0.5$ this is a uniform distribution and at $\beta = 1$ it produces the inverse of the original predictions.

Figure 3 shows the widths as β is varied ($n = 112$ on TREC-DL and $n = 125$ on Robust04). We do not report coverage as all methods obtain a mean coverage of at least 95%. When $\beta < 0.5$, CRC consistently provides better widths than empirical bootstrap, whilst PPI has inconsistent improvements. As expected, when $\beta > 0.5$ both methods do worse than empirical bootstrap in terms of width.

Thus, we can answer RQ2: the coverage of both PPI and CRC-bootstrap are robust to systematic mistakes made by the LLM, however, improvements in widths are dependent on LLM accuracy.

8.3 Potential from more accurate labels

We run additional experiments to understand how the CIs behave under an oracle LLM: one that can perfectly generate relevance labels. In Figure 3, we increasingly interpolate between the LLM-generated relevance labels and the true (human-annotated) relevance labels using a parameter $\tau \in [0, 1]$. As τ increases, the performance of the LLM labels becomes better. First, we note that all methods retain a perfect 100% coverage in these scenarios, so we omit the plots for coverage. The empirical bootstrap approach does not use the LLM-generated labels and its CI is thus not impacted by the increasingly stronger LLM labels. The PPI method is able to leverage the stronger LLM labels and is able to significantly outperform the empirical bootstrap method. The fact that its CI is placed around the overall performance (dataset-level), prevents it from further improving the width, as it is inherently limited by the number of queries. The CRC approaches are able to work around this limitation by efficiently identifying that the LLM-generated labels are more accurate as $\tau \rightarrow 1$ on the *per-document* level. Their per-query CIs correspondingly shrink and approach 0 as the LLM-generated labels become better. This answers RQ3: Both PPI and CRC benefit from improvements in label generation accuracy.

8.4 Query-performance confidence intervals

We plot the confidence intervals produced by CRC on individual queries in Figure 4. Each plot in the figure shows the true DCG (based on human-annotated relevance labels) and the predicted DCG (based on LLM-generated labels) of all queries in the test split. The queries are sorted by their true DCG, that is, queries where the ranker performs best appear on the left and progressively the query performance goes down. Furthermore, we plot the per-query

CI for varying values of τ , to indicate how the confidence intervals behave as the LLM-generated labels become more accurate, similar to Section 8.3. First, for all plots, we observe that the CIs vary per query: CRC captures the uncertainty throughout LLM-generated labels. Second, for the LLM-generated labels ($\tau = 0$), we observe that when the LLM predicts that the ranker performs poorly on a query, the bounds tend to be smaller for that query. Similarly, when the predicted performance of the ranker is large, the bounds tend to be wider. This indicates that the LLM-generated labels are generally better at identifying queries with poor ranking performance. Third, as $\tau \rightarrow 1$, we see that CRC is able to identify that the labels are more accurate and its per-query CIs become significantly tighter. This shows that CRC is not only able to vary its CI per query, but is also able to establish better per-query CIs as LLM labels become more accurate. This is especially noticeable in the $\tau = 0.75$ plot for TREC-DL (top-right plot in Figure 4). In this plot there is a single outlier query on the left where the LLM is wrong and its predicted labels are uncertain. As a result the CRC method correctly places a very wide CI around this particular query, while keeping the CIs on other queries tight. Finally, on both datasets the empirical coverage of 95% is reached, indicating the CIs are reliable. Thus, we answer RQ4 positively: CRC is able to construct CIs on a *per-query* basis.

9 Conclusion

In this paper we study reliable evaluation of IR systems using LLM-generated relevance labels. Obtaining human-annotated relevance labels is costly, especially in low-resource settings. While LLMs can help generate relevance labels at scale, they are prone to make systematic errors and may be unreliable. We resolve this by introducing two methods that construct confidence intervals (CIs) around ranking metrics produced by LLM-generated relevance labels: PPI and CRC. These approaches require a small amount of reliable ground truth annotations to statistically analyze the distribution of errors and correct those errors.

Our results demonstrate that the proposed methods can correct errors in LLM-generated labels and produce reliable CIs. Compared to other CI approaches, we can produce CIs of superior coverage with tighter bounds, leading to more informative evaluation. Furthermore, the CIs produced by CRC can be computed per-query, providing further insights into low or high performing queries.

Our work is not without limitations. First, we note that our methods require an LLM with scoring-mode to produce a distribution over LLM labels. For LLMs without scoring-mode one could generate multiple labels stochastically to approximate a predicted distribution. Second, our results suggest that applying some smoothing to the LLM-generated label distribution is beneficial to the resulting CIs. How to systematically optimize the amount of smoothing is an open question. Similarly, fine-tuning or prompt-engineering could also lead to distributions better suited for CI construction. Third, we only use the Flan-UL2 as an LLM labeler. Our work can be extended to use different and potentially more powerful LLMs. Future work could explore all of these directions further.

Acknowledgements

This research was supported by the Google Visiting Researcher program. Any opinions, findings and recommendations expressed

in this work are those of the authors and are not necessarily shared or endorsed by their respective employers or sponsors.

References

- [1] Meysam Alizadeh, Maël Kubli, Zeynab Samei, Shirin Dehghani, Juan Diego Bermeo, Maria Korobeynikova, and Fabrizio Gilardi. 2023. Open-source large language models outperform crowd workers and approach ChatGPT in text-annotation tasks. *arXiv preprint arXiv:2307.02179* (2023).
- [2] Anastasios N Angelopoulos and Stephen Bates. 2021. A gentle introduction to conformal prediction and distribution-free uncertainty quantification. *arXiv preprint arXiv:2107.07511* (2021).
- [3] Anastasios N Angelopoulos, Stephen Bates, Clara Fannjiang, Michael I Jordan, and Tijana Zrnic. 2023. Prediction-powered inference. *arXiv preprint arXiv:2301.09633* (2023).
- [4] Anastasios N Angelopoulos, Stephen Bates, Adam Fisch, Lihua Lei, and Tal Schuster. 2022. Conformal risk control. *arXiv preprint arXiv:2208.02814* (2022).
- [5] Javed A Aslam, Virgil Pavlu, and Emine Yilmaz. 2006. A statistical method for system evaluation using incomplete judgments. In *Proceedings of the 29th annual international ACM SIGIR conference on Research and development in information retrieval*. 541–548.
- [6] 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*. 667–674.
- [7] Vineeth Balasubramanian, Shen-Shyang Ho, and Vladimir Vovk. 2014. *Conformal prediction for reliable machine learning: theory, adaptations and applications*. Newnes.
- [8] Emily M Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the dangers of stochastic parrots: Can language models be too big?. In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*. 610–623.
- [9] Luiz Bonifacio, Hugo Abonizio, Marzieh Fadaee, and Rodrigo Nogueira. 2022. Inpars: Unsupervised dataset generation for information retrieval. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 2387–2392.
- [10] F Jay Breidt and Jean D Opsomer. 2017. Model-assisted survey estimation with modern prediction techniques. (2017).
- [11] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems* 33 (2020), 1877–1901.
- [12] Michael Buckland and Fredric Gey. 1994. The relationship between recall and precision. *Journal of the American society for information science* 45, 1 (1994), 12–19.
- [13] Ben Carterette, James Allan, and Ramesh Sitaraman. 2006. Minimal test collections for retrieval evaluation. In *Proceedings of the 29th annual international ACM SIGIR conference on Research and development in information retrieval*. 268–275.
- [14] Olivier Chapelle and Yi Chang. 2011. Yahoo! learning to rank challenge overview. In *Proceedings of the learning to rank challenge*. PMLR, 1–24.
- [15] Charles LA Clarke, Gianluca Demartini, Laura Dietz, Guglielmo Faggioli, Matthias Hagen, Claudia Hauff, Noriko Kando, Evangelos Kanoulas, Martin Potthast, Ian Soboroff, et al. 2023. 4.2 HMC: A Spectrum of Human-Machine-Collaborative Relevance Judgment Frameworks. *Frontiers of Information Access Experimentation for Research and Education* (2023), 41.
- [16] Gordon V Cormack and Thomas R Lynam. 2006. Statistical precision of information retrieval evaluation. In *Proceedings of the 29th annual international ACM SIGIR conference on Research and development in information retrieval*. 533–540.
- [17] 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. In *Proceedings of the 44th international ACM SIGIR conference on research and development in information retrieval*. 2369–2375.
- [18] Antonia Creswell, Tom White, Vincent Dumoulin, Kai Arulkumaran, Biswa Sengupta, and Anil A Bharath. 2018. Generative adversarial networks: An overview. *IEEE signal processing magazine* 35, 1 (2018), 53–65.
- [19] Jia Cui, Brian Kingsbury, Bhuvana Ramabhadran, Abhinav Sethy, Kartik Audhkhasi, Xiaodong Cui, Ellen Kislal, Lidia Mangu, Markus Nussbaum-Thom, Michael Picheny, et al. 2015. Multilingual representations for low resource speech recognition and keyword search. In *2015 IEEE workshop on automatic speech recognition and understanding (ASRU)*. IEEE, 259–266.
- [20] Domenico Dato, Sean MacAvaney, Franco Maria Nardini, Raffaele Perego, and Nicola Tonello. 2022. The Istella22 Dataset: Bridging Traditional and Neural Learning to Rank Evaluation. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 3099–3107.
- [21] Thomas Demeester, Robin Aly, Djoerd Hiemstra, Dong Nguyen, and Chris Davelder. 2016. Predicting relevance based on assessor disagreement: analysis and practical applications for search evaluation. *Information Retrieval Journal* 19

- (2016), 284–312.
- [22] Thomas J DiCiccio and Bradley Efron. 1996. Bootstrap confidence intervals. *Statistical science* 11, 3 (1996), 189–228.
- [23] Thomas J Diccio and Joseph P Romano. 1988. A review of bootstrap confidence intervals. *Journal of the Royal Statistical Society Series B: Statistical Methodology* 50, 3 (1988), 338–354.
- [24] Bradley Efron. 1987. Better bootstrap confidence intervals. *Journal of the American statistical Association* 82, 397 (1987), 171–185.
- [25] Guglielmo Faggioli, Laura Dietz, Charles LA Clarke, Gianluca Demartini, Matthias Hagen, Claudia Hauff, Noriko Kando, Evangelos Kanoulas, Martin Potthast, Benno Stein, et al. 2023. Perspectives on large language models for relevance judgment. In *Proceedings of the 2023 ACM SIGIR International Conference on Theory of Information Retrieval*. 39–50.
- [26] Norbert Fuhr. 2018. Some common mistakes in IR evaluation, and how they can be avoided. In *Acm sigir forum*, Vol. 51. ACM New York, NY, USA, 32–41.
- [27] Fabrizio Gilardi, Meysam Alizadeh, and Maël Kubli. 2023. Chatgpt outperforms crowd-workers for text-annotation tasks. *arXiv preprint arXiv:2303.15056* (2023).
- [28] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2020. Generative adversarial networks. *Commun. ACM* 63, 11 (2020), 139–144.
- [29] Jiafeng Guo, Yixing Fan, Liang Pang, Liu Yang, Qingyao Ai, Hamed Zamani, Chen Wu, W Bruce Croft, and Xueqi Cheng. 2020. A deep look into neural ranking models for information retrieval. *Information Processing & Management* 57, 6 (2020), 102067.
- [30] Donna Harman. 2011. *Information retrieval evaluation*. Morgan & Claypool Publishers.
- [31] Donna K Harman. 2005. The TREC test collections. (2005).
- [32] Tim Hesterberg. 2011. Bootstrap. *Wiley Interdisciplinary Reviews: Computational Statistics* 3, 6 (2011), 497–526.
- [33] George Hripcsak and Adam S Rothschild. 2005. Agreement, the f-measure, and reliability in information retrieval. *Journal of the American medical informatics association* 12, 3 (2005), 296–298.
- [34] Zhiqi Huang, Puxuan Yu, and James Allan. 2023. Improving Cross-lingual Information Retrieval on Low-Resource Languages via Optimal Transport Distillation. In *Proceedings of the Sixteenth ACM International Conference on Web Search and Data Mining*. 1048–1056.
- [35] Kalervo Järvelin and Jaana Kekäläinen. 2002. Cumulated gain-based evaluation of IR techniques. *ACM Transactions on Information Systems (TOIS)* 20, 4 (2002), 422–446.
- [36] Kalervo Järvelin and Jaana Kekäläinen. 2017. IR evaluation methods for retrieving highly relevant documents. In *ACM SIGIR Forum*, Vol. 51. ACM New York, NY, USA, 243–250.
- [37] Mladan Jovanovic and Mark Campbell. 2022. Generative artificial intelligence: Trends and prospects. *Computer* 55, 10 (2022), 107–112.
- [38] Noriko Kando, Kazuko Kuriyama, Toshihiko Nozue, Koji Eguchi, Hiroyuki Kato, and Souichiro Hidaka. 1999. Overview of IR tasks at the first NTCIR workshop. In *Proceedings of the first NTCIR workshop on research in Japanese text retrieval and term recognition*. 11–44.
- [39] Jaana Kekäläinen and Kalervo Järvelin. 2002. Using graded relevance assessments in IR evaluation. *Journal of the American Society for Information Science and Technology* 53, 13 (2002), 1120–1129.
- [40] Michael E Lesk and Gerard Salton. 1968. Relevance assessments and retrieval system evaluation. *Information storage and retrieval* 4, 4 (1968), 343–359.
- [41] Yue Liu, Zhengwei Yang, Zhenyao Yu, Zitu Liu, Dahui Liu, Hailong Lin, Mingqing Li, Shuchang Ma, Maxim Avdeev, and Siqi Shi. 2023. Generative artificial intelligence and its applications in materials science: Current situation and future perspectives. *Journal of Materiomics* (2023).
- [42] Shayne Longpre, Le Hou, Tu Vu, Albert Webson, Hyung Won Chung, Yi Tay, Denny Zhou, Quoc V Le, Barret Zoph, Jason Wei, et al. 2023. The flan collection: Designing data and methods for effective instruction tuning. *arXiv preprint arXiv:2301.13688* (2023).
- [43] Claudio Lucchese, Franco Maria Nardini, Raffaele Perego, Salvatore Orlando, and Salvatore Trani. 2018. Selective gradient boosting for effective learning to rank. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*. 155–164.
- [44] Sean MacAvaney and Luca Soldaini. 2023. One-Shot Labeling for Automatic Relevance Estimation. *arXiv preprint arXiv:2302.11266* (2023).
- [45] Bhaskar Mitra and Nick Craswell. 2017. Neural models for information retrieval. *arXiv preprint arXiv:1705.01509* (2017).
- [46] Bhaskar Mitra, Nick Craswell, et al. 2018. An introduction to neural information retrieval. *Foundations and Trends® in Information Retrieval* 13, 1 (2018), 1–126.
- [47] 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. *choice* 2640 (2016), 660.
- [48] Kezban Dilek Onal, Ye Zhang, Ismail Sengor Altıngövdü, Md Mustafizur Rahman, Pinar Karagoz, Alex Braylan, Brandon Dang, Heng-Lu Chang, Henna Kim, Quinten McNamara, et al. 2018. Neural information retrieval: at the end of the early years. *Information Retrieval Journal* 21 (2018), 111–182.
- [49] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems* 35 (2022), 27730–27744.
- [50] Harris Papadopoulos. 2008. Inductive conformal prediction: Theory and application to neural networks. In *Tools in artificial intelligence*. Citeseer.
- [51] John V Pavlik. 2023. Collaborating with ChatGPT: Considering the implications of generative artificial intelligence for journalism and media education. *Journalism & Mass Communication Educator* 78, 1 (2023), 84–93.
- [52] Carol Peters. 2001. *Cross-Language Information Retrieval and Evaluation: Workshop of Cross-Language Evaluation Forum, CLEF 2000, Lisbon, Portugal, September 21–22, 2000, Revised Papers*. Vol. 2069. Springer Science & Business Media.
- [53] Tao Qin and Tie-Yan Liu. 2013. Introducing LETOR 4.0 datasets. *arXiv preprint arXiv:1306.2597* (2013).
- [54] Tao Qin, Tie-Yan Liu, Jun Xu, and Hang Li. 2010. LETOR: A benchmark collection for research on learning to rank for information retrieval. *Information Retrieval* 13 (2010), 346–374.
- [55] Zhen Qin, Rolf Jagerman, Kai Hui, Honglei Zhuang, Junru Wu, Jiaming Shen, Tianqi Liu, Jialu Liu, Donald Metzler, Xuanhui Wang, et al. 2023. Large language models are effective text rankers with pairwise ranking prompting. *arXiv preprint arXiv:2306.17563* (2023).
- [56] Tetsuya Sakai. 2014. Statistical reform in information retrieval?. In *ACM SIGIR Forum*, Vol. 48. ACM New York, NY, USA, 3–12.
- [57] Mark Sanderson et al. 2010. Test collection based evaluation of information retrieval systems. *Foundations and Trends® in Information Retrieval* 4, 4 (2010), 247–375.
- [58] Mark Sanderson and Justin Zobel. 2005. Information retrieval system evaluation: effort, sensitivity, and reliability. In *Proceedings of the 28th annual international ACM SIGIR conference on Research and development in information retrieval*. 162–169.
- [59] Mark D Smucker, James Allan, and Ben Carterette. 2007. A comparison of statistical significance tests for information retrieval evaluation. In *Proceedings of the sixteenth ACM conference on Conference on information and knowledge management*. 623–632.
- [60] Yi Tay, Mostafa Dehghani, Vinh Q Tran, Xavier Garcia, Dara Bahri, Tal Schuster, Huaixiu Steven Zheng, Neil Houlsby, and Donald Metzler. 2022. Unifying language learning paradigms. *arXiv preprint arXiv:2205.05131* (2022).
- [61] Gemini Team, Rohan Anil, Sebastian Borgeaud, Yonghui Wu, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, et al. 2023. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805* (2023).
- [62] Nandan Thakur, Nils Reimers, Andreas Rücklé, Abhishek Srivastava, and Iryna Gurevych. 2021. Beir: A heterogenous benchmark for zero-shot evaluation of information retrieval models. *arXiv preprint arXiv:2104.08663* (2021).
- [63] Paul Thomas, Seth Spielman, Nick Craswell, and Bhaskar Mitra. 2023. Large language models can accurately predict searcher preferences. *arXiv preprint arXiv:2309.10621* (2023).
- [64] Petter Törnberg. 2023. Chatgpt-4 outperforms experts and crowd workers in annotating political twitter messages with zero-shot learning. *arXiv preprint arXiv:2304.06588* (2023).
- [65] George Tsatsaronis, Georgios Balikas, Prodomos Malakasiotis, Ioannis Patalas, Matthias Zschunke, Michael R Alvers, Dirk Weissenborn, Anastasia Krithara, Sergios Petridis, Dimitris Polychronopoulos, et al. 2015. An overview of the BIOASQ large-scale biomedical semantic indexing and question answering competition. *BMC bioinformatics* 16, 1 (2015), 1–28.
- [66] Julián Urbano, Harley Lima, and Alan Hanjalic. 2019. Statistical significance testing in information retrieval: an empirical analysis of type I, type II and type III errors. In *Proceedings of the 42nd International ACM SIGIR conference on Research and development in information retrieval*. 505–514.
- [67] Cornelis Joost Van Rijsbergen and W Bruce Croft. 1975. Document clustering: An evaluation of some experiments with the Cranfield 1400 collection. *Information Processing & Management* 11, 5-7 (1975), 171–182.
- [68] Ellen Voorhees, Tasmeer Alam, Steven Bedrick, Dina Demner-Fushman, William R Hersh, Kyle Lo, Kirk Roberts, Ian Soboroff, and Lucy Lu Wang. 2021. TREC-COVID: constructing a pandemic information retrieval test collection. In *ACM SIGIR Forum*, Vol. 54. ACM New York, NY, USA, 1–12.
- [69] Ellen M Voorhees. 2019. The evolution of cranfield. *Information Retrieval Evaluation in a Changing World: Lessons Learned from 20 Years of CLEF* (2019), 45–69.
- [70] Ellen M Voorhees et al. 2003. Overview of the TREC 2003 robust retrieval track. In *Trec*. 69–77.
- [71] Ellen M Voorhees, Donna K Harman, et al. 2005. *TREC: Experiment and evaluation in information retrieval*. Vol. 63. MIT press Cambridge.
- [72] William Webber. 2013. Approximate recall confidence intervals. *ACM Transactions on Information Systems (TOIS)* 31, 1 (2013), 1–33.
- [73] William Edward Webber. 2010. *Measurement in information retrieval evaluation*. Ph.D. Dissertation. University of Melbourne, Department of Computer Science and Software Engineering.

- [74] Mahsa Yarmohammadi, Xutai Ma, Sorami Hisamoto, Muhammad Rahman, Yiming Wang, Hainan Xu, Daniel Povey, Philipp Koehn, and Kevin Duh. 2019. Robust document representations for cross-lingual information retrieval in low-resource settings. In *Proceedings of Machine Translation Summit XVII: Research Track*. 12–20.
- [75] Emine Yilmaz, Evangelos Kanoulas, and Javed A Aslam. 2008. A simple and efficient sampling method for estimating AP and NDCG. In *Proceedings of the 31st annual international ACM SIGIR conference on Research and development in information retrieval*. 603–610.
- [76] Honglei Zhuang, Zhen Qin, Kai Hui, Junru Wu, Le Yan, Xuanhui Wang, and Michael Berdersky. 2023. Beyond yes and no: Improving zero-shot llm rankers via scoring fine-grained relevance labels. *arXiv preprint arXiv:2310.14122* (2023).

A Prompts

The exact prompts used in our experiments are listed here. We note that these prompts are tailored towards each dataset and use the relevance label definitions that human labelers used for each dataset. The {query} and {passage}/{document} are placeholders that are formatted with the actual query and passage/document during inference.

We observed that the model is sensitive to the particular prompt and dataset during scoring mode. For TREC-DL we score the suffixes "0", "1", "2" and "3". For Robust04 we found that scoring the suffixes with brackets is more effective: "[0]", "[1]" and "[2]".

Listing 1: Prompt for TREC-DL.

Assess the relevance of the passage to the query on a four-point scale:

- [0] Irrelevant: The passage has nothing to do with the query.
- [1] Related: The passage seems related to the query but does not answer it.
- [2] Highly relevant: The passage has some answer for the query, but the answer may be a bit unclear, or hidden amongst extraneous information.
- [3] Perfectly relevant: The passage is dedicated to the query and contains the exact answer.

Query: {query}
 Passage: {passage}
 Relevance:

Listing 2: Prompt for Robust04.

Assess the relevance of the document to the query on a three-point scale:

- [0] Not relevant: The document is not relevant to the query.
- [1] Relevant: Parts of the document may be relevant to the query.
- [2] Highly Relevant: The document is highly relevant to the query.

Query: {query}
 Document: {document}
 Relevance: