

# Best-Case Retrieval Evaluation: Improving the Sensitivity of Reciprocal Rank with Lexicographic Precision

Fernando Diaz\*  
 Carnegie Mellon University  
 Pittsburgh, PA, USA  
 diazf@acm.org

## ABSTRACT


Across a variety of ranking tasks, researchers use reciprocal rank to measure the effectiveness for users interested in exactly one relevant item. Despite its widespread use, evidence suggests that reciprocal rank is brittle when discriminating between systems. This brittleness, in turn, is compounded in modern evaluation settings where current, high-precision systems may be difficult to distinguish. We study the scenario where there is more than one relevant item and address the lack of sensitivity of reciprocal rank by introducing and connecting it to the concept of best-case retrieval, an evaluation method focusing on assessing the quality of a ranking for the most satisfied possible user across possible recall requirements. This perspective allows us to generalize reciprocal rank and define a new preference-based evaluation we call lexicographic precision or lexiprecision. By mathematical construction, we ensure that lexiprecision preserves differences detected by reciprocal rank, while empirically improving sensitivity and robustness across a broad set of retrieval and recommendation tasks.

## 1 INTRODUCTION

Evaluating ranking systems for users seeking exactly one relevant item has a long history in information retrieval. As early as 1968, Cooper [7] proposed *Type 1 expected search length* or  $ESL_1$ , defined as the rank position of the highest ranked relevant item. In the context of TREC-5, Kantar and Voorhees [13] proposed using the reciprocal of  $ESL_1$  in order to emphasize rank changes at the top of the ranked list and modeling the impatience of a searcher as they need to scan for a single item. Over the years, reciprocal rank (and less so  $ESL_1$ ) has established itself as a core metric for retrieval [5] and recommendation [4], adopted in situations where there is actually only one relevant item as well as in situations where there are multiple relevant items. Given two rankings, reciprocal rank and  $ESL_1$  always agree in terms of which ranking is better. Because of this, we refer to them collectively as the recall level 1 or  $RL_1$  metrics.

Despite the widespread use of reciprocal rank, recent evidence suggests that it may be brittle when it comes to discriminating between ranking systems [11, 21, 22]. In particular, the low number of unique values of reciprocal rank means that, especially when evaluating multiple highly-performing systems, we are likely to observe tied performance. Voorhees et al. [23] demonstrate that these conditions exist in many modern deep learning benchmarks.

We address these issues by theoretically interpreting  $RL_1$  as a population-level metric we refer to as *best-case retrieval evaluation*. This allows us to propose a generalization of the  $RL_1$  ordering based



**Figure 1: Hasse diagram of possible positions of relevant items. Each tuple represents the possible positions of five relevant items in a corpus of  $n$  items.  $RL_1$  metrics such as reciprocal rank (left) have  $n - 4$  unique values and, therefore, result in a *partial order* over all possible positions of relevant items. Lexicographic precision (right) is a *total order* over all possible positions of relevant items that preserves all strict orders in  $RL_1$  evaluation.**

on social choice theory [19] and preference-based evaluation [8]. This evaluation method, lexicographic precision or lexiprecision, preserves any strict ordering between rankings based on  $RL_1$  while also providing a theoretically-justified ordering when  $RL_1$  is tied.

We compare lexiprecision and  $RL_1$  orderings using Hasse diagrams in Figure 1. On the left, we show the partial order of all possible positions of five relevant items in a corpus of size  $n$ . Since reciprocal rank and  $ESL_1$  only consider the position of the first relevant item, we only have  $n$  different relevance levels. While this may not be an issue in general (since  $n$  is usually large), the number of rankings within each level can be very large and multiple highly effective systems can result in numerous ties. In contrast, lexiprecision has one relevance level for each unique arrangement of relevant items. That is, the number of relevance levels scales with the number relevant items and, by design, two rankings are tied *only* if they place relevant items in exactly the same positions.

In this paper, we contribute to the theoretical understanding of evaluation through a detailed study of  $RL_1$  metrics, best-case retrieval evaluation, and lexiprecision. In Section 2, we motivate our work by showing that  $RL_1$  has fundamental theoretical limits, especially in situations where there are multiple relevant items. In Section 3, we demonstrate that  $RL_1$  can be interpreted as best-case retrieval evaluation, allowing us to address its limitations by

\*work done at Google

using methods from social choice theory and generalizing it as lexicprecision. In Section 5, we then conduct extensive empirical analysis to show that lexicprecision is strongly correlated with  $RL_1$  metrics while substantially improving its discriminative power.<sup>1</sup>

## 2 MOTIVATION

Our work is based on the observation that ceiling effects are inherent in  $RL_1$  evaluation. Assume a standard ranking problem where, given a query with  $m$  associated relevant items, a system orders all  $n$  documents in the collection in decreasing order of predicted relevance. The set of all possible rankings of  $n$  is referred to as the symmetric group over  $n$  elements and is represented as  $S_n$ . For a given ranking  $\pi \in S_n$ , let  $p_i$  be the position of the  $i$ th highest-ranked relevant item. We can then define reciprocal rank as  $RR_1(\pi) = \frac{1}{p_1}$ . When no relevant document is retrieved (e.g. if no relevant items are in the system’s top  $k$  retrieval), we set  $RR_1 = 0$ . For two rankings, we define  $\delta RR_1(\pi, \pi') = RR_1(\pi) - RR_1(\pi')$ . For the remainder of this section, we will use reciprocal rank for clarity although the analysis applies to  $ESL_1$  as well.

Although we can easily see that there are  $n$  different values for  $RR_1(\pi)$ , we are interested in the distribution of ties amongst system rankings for these  $n$  values as predicted by theoretical properties of reciprocal rank. Specifically, we want to compute, for a given position of the first relevant item  $p_1$  and a *random* second ranking, the probability that we will observe a tie. For any  $p_1$ , there are  $\binom{n-p_1}{m-1}$  tied arrangements of positions of relevant items amongst all of the possible arrangements  $p'$  from a second system. If we sample an arrangement of relevant items uniformly at random, then the probability of a tie with  $\pi$  is  $Pr(p_1 = p'_1 | p_1) = \frac{\binom{n-p_1}{m-1}}{\binom{n}{m}}$ . We plot this probability in Figure 2. We can observe that, when we have few relevant items (i.e. small  $m$ ), we have a relatively small and uniform probability of ties across all values of  $p_1$ . However, as we increase the number of relevant items, the distribution begins to skew toward a higher probability of a tie as  $p_1$  is smaller. This means that, if we have a ranking where the first relevant item is close to the top, even if the second ranking is drawn uniformly at random, we will be more likely to find a tie than if the first relevant item were lower in the ranking.

While our analysis indicates a lack of sensitivity of reciprocal rank for  $p'$  drawn uniformly at random as  $m$  increases, we are also interested in the probability of ties when  $p'$  is drawn from rankings produced by real systems. We collected runs associated with multiple public benchmarks (see Section 4.1 for details) and computed the the empirical distribution of ties conditioned  $p_1$  (Figure 3). Because of the highly skewed distribution, we plot the logarithmic transform of the probability of a rank position. As we can see, across both older and newer benchmarks, the probability of a tie for rankings when the top-ranked relevant item is at position 1 is substantially larger than if we assume  $p'$  is drawn uniformly at random. The 2021 TREC Deep Learning track data in particular demonstrates higher skew than others, confirming observations previously made about saturation at top rank positions [23].

<sup>1</sup>In lieu of an isolated ‘Related Work’ section, we have included discussion of relevant literature when necessary. This helps make connections explicit to our work.




Figure 2: Given a ranking where the highest-ranked relevant item is at position  $p_1$ , the probability of a tie with a second ranking sampled uniformly from all arrangements of relevant items for a corpus size of  $n = 50000$ . This figure and others are best rendered or printed in color.




Figure 3: Empirical probability of a tie with a second ranking for several benchmarks (see Section 4.1 for details). Horizontal and vertical axes are on a logarithmic scale for clarity.

Taken together, these results demonstrate a fundamental limitation of  $RL_1$  metrics (i.e., reciprocal rank and  $ESL_1$ ) for evaluation. As retrieval and other scenarios where reciprocal rank is used begin to attract highly performant systems, we need to extend our evaluation approaches to address these issues.

## 3 LEXICOGRAPHIC PRECISION

$RL_1$  evaluation emphasizes precision by considering the position of the top-ranked relevant item and ignoring the positions of other relevant items. However, *only* ever looking at the position of the top-ranked relevant item results in the ceiling effects described in the previous section. Our goal is to develop an evaluation method that

preserves the ordering of a pair of rankings by  $RL_1$  (i.e., agree with  $RR_1$  when  $RR_1(\pi) \neq RR_1(\pi')$ ) and provides a justified ordering of a pair of rankings when  $RL_1$  is tied (i.e., generate a sensible order when  $RR_1(\pi) = RR_1(\pi')$ ). Although metrics like expected reciprocal rank [6] and average precision include reciprocal rank as a component in their computation, they are not *guaranteed* to preserve the ordering of reciprocal rank when there is one.<sup>2</sup> In this section, we will interpret  $RL_1$  metrics as best-case retrieval evaluation, allowing us to derive a preference-based evaluation method based on social choice theory.

### 3.1 Best-Case Retrieval Evaluation

When a user approaches a retrieval system, there is a great deal of uncertainty about their information need. While a request such as a text query provides information about which items in the corpus might be relevant, it says much less about the user’s appetite for relevant information. As a result, there is an implicit population of possible user types issuing any particular request, each of whom may have a different utility for any particular ranking. In this section, we explore two types of uncertainty and demonstrate that, from both perspectives,  $RL_1$  evaluation represents the best-case utility over that population.

We first consider uncertainty over recall requirements. Robertson [17] presented a model for evaluating rankings based on the diverse set of recall requirements that a user might have. Given a request and its associated relevant items, users may be interested in one relevant item, a few relevant items, or the complete set of all relevant items. We can assess the quality of a ranking for any particular user recall requirement with what Cooper [7] refers to as the *Type 2 expected search length*: the number of items a user with requirement  $i$  has to scan before finding  $i$  relevant items. So, each information need has  $m$  recall levels and  $RL_1$  is the evaluation measure associated with users requiring exactly one relevant item. From this perspective, we can, for a specific ranking, look at how utility is distributed amongst possible user types, as represented by their recall levels. For example, we can ask how utility for users with high and low recall requirements compares; or what the average utility across these populations is. While previous work has looked at the average-case utility [8] and worst-case utility [10], in this work we suggest that  $RL_1$  represents the best-case performance over these possible user types. The proof is relatively simple. Because  $RR_i$  monotonically degrades in rank, the best-case utility over this representation of users is  $RR_1$  (equivalently  $ESL_1$ ). The next-best-case is  $RR_2$  and so forth until we reach  $RR_m$ , which we refer to as the worst-case. So, given two rankings  $\pi$  and  $\pi'$ , observing  $RR_1(\pi) > RR_1(\pi')$  implies that the best-case performance over possible user recall requirements is higher in  $\pi$  compared to  $\pi'$ .

Next, we consider uncertainty over psychologically relevant items. When evaluating a retrieval system, we often use relevance labels derived from human assessors or statistical models. But what if a specific user does not find the top-ranked item labeled relevant actually relevant to them? For example, a user may have already

seen a specific item or they may desire an item with a specific (missing) attribute. A judged relevant item might be inappropriate for any number of reasons not expressed in the request. The concept of *psychological relevance* [12] suggests that judging any item relevant *in general* (as is the case in many retrieval benchmarks, including those used in TREC) is a necessary but not sufficient criteria to determine an item’s psychological relevance to any particular user. From this perspective, there are  $2^m - 1$  possible non-empty sets of relevant items for a specific request, each representing psychological relevance to a possible user type. Nevertheless, amongst these possible user types, if they are interested in precisely one relevant item, there are  $m$  unique utilities. Again, since  $RL_1$  monotonically decreases in rank, the best-case utility is  $RR_1$ , followed by  $RR_2$  until we reach  $RR_m$ .

Both uncertainty over recall levels and over psychological relevance focus on possible populations of users. Because the utility to the user implies utility to the system designer (e.g., for objectives like retention), understanding the best-case performance is valuable in decision-making. From the perspective of social choice theory, best-case retrieval evaluation is inherently optimistic and represents risk-seeking decision-making.

### 3.2 Lexicographic Precision

The problem with evaluating for best-case retrieval (as shown in Section 2) is the tendency for multiple rankings to be tied, especially as (i) we increase the number of relevant items and (ii) systems optimize for retrieval metrics. We can address these ceiling effects by developing a best-case *preference-based evaluation* that focuses on measuring differences in performance instead of absolute performance [8]. While *metric-based evaluation* models the preference between rankings by first computing some evaluation metric for each ranking, preference-based evaluation explicitly models the preference between two rankings. Prior research has demonstrated that preference-based evaluation can be much more sensitive than metric-based evaluation [8], making it well-suited for addressing the ceiling effects described in Section 2.

Under best-case preference-based retrieval, we are interested in answering the question, ‘under the best possible scenario, which ranking would the user prefer?’ In this respect, it is a user-based evaluation method, but one based on preferences and measurement over a population of users. More formally, given an information need and two rankings  $\pi$  and  $\pi'$  associated with two systems, metric-based evaluation uses an evaluation metric  $\mu : S_n \rightarrow \mathfrak{R}$  (e.g. reciprocal rank or average precision) to compute a preference,

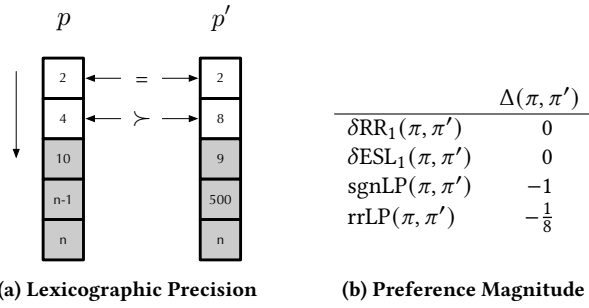
$$\mu(\pi) > \mu(\pi') \implies \pi > \pi'$$

where  $\pi > \pi'$  indicates that we prefer  $\pi$  to  $\pi'$ . Notice that, if  $\mu(\pi) = \mu(\pi')$ , then we cannot infer a preference between  $\pi$  and  $\pi'$ . We contrast this with preference-based evaluation, which directly models this relationship  $\Delta : S_n \times S_n \rightarrow \mathfrak{R}$ ,

$$\Delta(\pi, \pi') > 0 \implies \pi > \pi'$$

Our goal is to design a preference-based evaluation that preserves the best-case properties of  $RL_1$  metrics with much higher sensitivity. Consider the two position vectors  $p$  and  $p'$  in Figure 4 associated with the two rankings  $\pi$  and  $\pi'$ . These two vectors are tied in the best case (i.e.,  $p_1 = p'_1$ ). However, we can break this tie by

<sup>2</sup>Our emphasis of theoretical guarantees—as opposed to empirical or average behavior—follows from our related questions of justice and fairness over a population of users [9].



**Figure 4: Lexicographic precision between two rankings  $\pi$  and  $\pi'$  with  $m = 5$  relevant items in corpus of size  $n$ . (a) Using the sorted positions of relevant items, lexicoprecision returns a preference based on the highest-ranked difference in positions. (b) The magnitude of preference between  $\pi$  and  $\pi'$  under different schemes.**

looking at the next-best case (i.e.  $p_2$ ) where, because  $p_2 < p'_2$ , we say that  $\pi > \pi'$ . If we had observed a tie between the next-best case, we could compare  $p_3$ , and so forth. This is known as lexicographic sorting in the social choice literature [19] and reflects a generalization of best-case sorting. Given two sorted vectors of utilities, here reflected by the rank position, the lexicographic maximum begins by looking at utilities in the best-off positions (i.e.  $p_1$  and  $p'_1$ ) and iteratively inspects lower utility positions until we find an inequality. If we exhaust all  $m$  relevance levels, we indicate that there is not preference between the rankings. Note that a tie can only happen if two rankings have all relevant items in exactly the same positions. Lexicographic sorting generates a total ordering over all positions of relevant items, in contrast with just inspecting  $p_1$ , which compresses all arrangements onto  $n$  possible values. Because of its basis in lexicographic ordering, we refer to this as lexicographic precision or lexiprecision.

### 3.3 Number of Ties Under Lexicographic Precision

We can contrast the number of ties as  $m$  increases in  $RL_1$  metrics with the number of ties as  $m$  increases in lexiprecision. In the latter, we only observe ties when the positions of the relevant items for two rankings are the same and, therefore, we have  $\binom{n}{m}$  possible ‘values’ and the number of ties given a fixed ranking is constant. If we add  $k$  relevant items, the number of ‘values’ increases, resulting in an increase in discriminative power. Specifically, if we add  $k$  relevant items to  $m$ , then the number of possible values scales exponentially in  $k$ .

$$\frac{\binom{n}{m+k}}{\binom{n}{m}} = \prod_{i=m+1}^{m+k} \frac{n+1-i}{i} \quad (1)$$

By contrast, for  $RL_1$  metrics, this increase in the number of unique position vectors needs to be allocated to a fixed  $n$  values, resulting in collisions, as suggested by the pigeonhole principle. Moreover, these collisions will tend to increasingly occur at values associated with position vectors where  $p_1$  is small (Section 2).

### 3.4 Best-Case Retrieval Evaluation Revisited

In Section 3.1, we described two dimensions of uncertainty in retrieval evaluation: recall level and psychological relevance. In both cases, we saw that the best-case utility was represented by  $RL_1$ . In terms of preference-based evaluation, we would like to show that, for both recall level uncertainty and psychological relevance uncertainty, the highest ranked difference in utility will be  $\delta RR_{i^*}$ , where  $i^* = \operatorname{argmin}_{j \in [1, m]} \delta RR_j(\pi, \pi') \neq 0$ . This is clear for recall level uncertainty because the population of possible user types exactly matches the recall levels defining  $i^*$ .

However, for psychological relevance uncertainty, we have  $2^m - 1$  possible user types. That said, there are only  $m$  possible  $RL_1$  metric values. Moreover, the number of possible user types tied at the first recall level is  $2^{m-1}$ ; at the second recall level is  $2^{m-2}$ ; down to the final recall level where there is a single possible user type. This arrangement of ties is the same regardless of the exact positions of the relevant items. Therefore, if we observe  $\delta RR_1 = 0$ , we will observe  $2^{m-1}$  ties amongst the possible psychological relevance states where the first relevant item is at position  $p_1$ . The next highest utility is, by the monotonicity of  $RL_1$  metrics, associated with the second recall level. We can continue this procedure until we observe an inequality, which will occur exactly at the first  $i$  such that  $\delta RR_i(\pi, \pi') \neq 0$ . In other words,  $i^*$ .

These observations are important since they demonstrate that lexiprecision generalizes  $RL_1$  evaluation and best-case performance across two types of uncertainty.

### 3.5 Quantifying Preferences

Although lexiprecision provides an ordering over a pair of rankings, it does not quantify the magnitude of the preference (i.e. the value of  $\Delta(\pi, \pi')$ ). Defining a magnitude allows us to measure the degree of preference, which can then be averaged over multiple requests.

We can define the magnitude directly as the value of  $\delta RR_i$  and, therefore, defining  $\Delta(\pi, \pi')$  as,

$$\text{rrLP}(\pi, \pi') = \delta RR_{i^*}(\pi, \pi') \quad (2)$$

where  $i^*$  is defined in Section 3.4. This has the advantage of, when  $i^* = 1$ , reproducing the difference in reciprocal rank. Under this definition, the magnitude of preferences for higher recall levels will tend to be smaller due to the aggressive discounting in reciprocal rank.

Alternatively, we can be more conservative in our quantification and just return a constant value based on the preference, defining  $\Delta(\pi, \pi')$  as,

$$\text{sgnLP}(\pi, \pi') = \text{sgn}(\delta RR_{i^*}(\pi, \pi')) \quad (3)$$

where  $i^*$  is defined as above. Although the direction of the preference agrees with  $\text{rrLP}$ , we discard its magnitude and, as a result, differences at lower ranks are equal to those at higher ranks. Prior work found that looking at unweighted preference information alone can help with preference sensitivity [8].

### 3.6 Lexicographic Precision as Modeling $\delta RR_1$

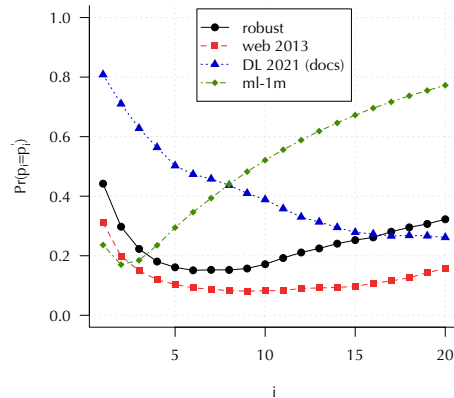
A different way to interpret lexiprecision is as a method to estimate a high-precision preference between rankings. Assume that we have some latent preference between two rankings,  $\hat{\Delta}(\pi, \pi')$ , that

we know to be ‘high-precision’. That is, users prefer finding *some* relevant items quickly than *all* relevant items quickly.


One way to model this preference is to inspect the positions of relevant items in  $\pi$  and  $\pi'$ . From the perspective of ‘very high precision’, observing  $\delta\text{RR}_1(\pi, \pi') > 0$  provides significant evidence that  $\hat{\Delta}(\pi, \pi') > 0$ . What if we do not observe a preference at the first recall level? Inspired by Katz’s back-off model [14], we inspect the second recall level for evidence of the value of  $\hat{\Delta}(\pi, \pi')$ . If we do not observe a preference, we can progressively back off to higher and higher recall levels.

While Section 2 demonstrated that  $\delta\text{RR}_1(\pi, \pi') = 0$  with high probability, backing off our estimates works best if, for  $i > 1$ , we expect  $\delta\text{RR}_i(\pi, \pi') = 0$  with lower probability. Using the runs associated with several public benchmarks, we computed  $\delta\text{RR}_i$  for all pairs of rankings generated by multiple systems for the same query. We show the probability of a tie for the first twenty recall levels in Figure 5. We can see that the number of ties at  $\delta\text{RR}_1$  are high, ranging from roughly 20% (ml-1m) to 80% (DL 2021). Across all datasets, the probability of a tie drops at  $\delta\text{RR}_2$ . With the exception of ml-1m, the probability of a tie continues to drop or plateaus until  $\delta\text{RR}_{10}$ , when, depending on the dataset, the number of ties increases, plateaus, or continues to decrease. The increase in the number of ties with recall level may seem surprising since, at higher recall levels, there are more rank positions and, therefore, ‘more ways to be different’. However, retrieving relevant documents for higher recall levels is harder since all systems in our dataset return top  $k$  rankings. This means that there will be unretrieved—and therefore unranked—items, especially for higher recall levels. So, the growth in ties after a certain recall level can be accounted for by the growth in recall failure which then results in ties, since two systems that fail to retrieve an  $i$ th ranked relevant item will have  $\delta\text{RR}_i = 0$ . Inspecting the number of relevant items retrieved confirms this. The DL 2021 submissions had  $38.74 \pm 21.75$  relevant items in their retrievals, compared to web with  $53.14 \pm 47.06$ . Meanwhile, robust submissions had  $40.51 \pm 41.49$  relevant items retrieval, suggesting much higher variance and ml-1m with  $7.46 \pm 8.57$  relevant items retrieved and much higher variance, leading to more more ties at higher recall levels.

Given that different benchmarks observed different behaviors for ties amongst recall levels, we need to understand how many recall levels we need to visit before finding evidence for  $\hat{\Delta}$ . If a benchmark needs many recall levels but observes many ties at high recall levels, then our model of  $\hat{\Delta}$  may be less reliable. We computed the number of recall levels needed,  $i^*$ , for each benchmark and plotted the empirical cumulative distribution function in Figure 6. We find that we need fewer than ten recall levels to capture 90% of the differences for most benchmarks. The DL 2021 dataset is an exception, needing close to 20 recall levels to detect the same number of differences. This is somewhat expected since the DL 2021 submissions tended to all be high-performing runs with many ties at top positions. One advantage of lexicoprecision as a model of  $\hat{\Delta}$ , then, is that it can, unlike a rank cutoff (e.g. precision at  $k$ ) or a recall level cutoff (e.g. reciprocal rank), adapt its depth according to the behavior of queries, relevant items, and systems as needed to distinguish rankings.



**Figure 5: Empirical probability of a tie in position by recall level. Note that, while Figure 2 measures the probability of a tie for different positions of the highest ranked relevant item (i.e.  $p_1$ ), this figure measures the probability of a tie for different recall levels.**



**Figure 6: Empirical cumulative distribution function of recall level needed to distinguish systems.**

Although our preceding analysis demonstrates that a backoff model of  $\hat{\Delta}$  based on lexicoprecision will terminate at a reasonable depth, we still need to show that there is locality amongst  $\delta\text{RR}_i$ . This means that we ask, if we observe  $\delta\text{RR}_1(\pi, \pi') > 0$ , how likely is it that  $\delta\text{RR}_2(\pi, \pi') > 0$ ?  $\delta\text{RR}_3(\pi, \pi') > 0$ ? If there is high locality amongst  $\delta\text{RR}_i$ , then information from  $\delta\text{RR}_{i+1}$  can help in predicting the true value of  $\delta\text{RR}_i$  when it is missing or tied. Note that, if we observe  $\delta\text{RR}_i > 0$  and  $n$  is large, there is absolutely no guarantee that  $\delta\text{RR}_{i+1} > 0$  since the next ranked relevant items could, in theory, occur anywhere in the range  $[p_i + 1, n]$  and  $[p'_i + 1, n]$ . That said, given the number of ties at recall level 1, we are interested in understanding whether information at other rank positions can provide a way to distinguish tied rankings. In Figure 7a, we computed the Pearson correlation amongst all pairs of  $\delta\text{RR}_i$  for  $i \in [1, 8]$  for the Robust 2004 benchmark. The fact that correlation between  $\delta\text{RR}_i$

and  $\delta RR_{i+j}$  degrades as  $j$  increases from 1 demonstrates that there is indeed high locality. The implication justifies the use of backoff modeling of  $\hat{\Delta}$ .

To test this hypothesis explicitly, we fit a linear model of  $\delta RR_1$  using  $\delta RR_2, \dots, \delta RR_4$  as independent variables. We plot the coefficients of the linear regression in the solid line in Figure 7b. The substantially larger coefficient on  $\delta RR_2$  indicates that the majority of the predictive power can be found at recall level 2 ( $j = 1$ ). Higher recall levels ( $j > 1$ ) are associated with much smaller coefficients. The actual contributions of higher recall levels are much smaller than this suggests since, because we are operating with reciprocals, the magnitude of  $\delta RR_i$  shrinks as  $i$  grows. While the colinearity in Figure 7a might explain some of this disparity in weights, the locality of individual Pearson correlations and high predictive accuracy means, from a modeling perspective, that a backoff model is justified. We repeated this analysis for predicting  $\delta RR_2$  from  $\delta RR_3, \dots, \delta RR_6$  and similarly for  $\delta RR_3$  and  $\delta RR_4$ . Similar to our observation when modeling  $\delta RR_1$ , these results suggest that the next higher recall level is the most valuable predictor when modeling  $\delta RR_i$  for any specific recall level.

We repeated this regression analysis for explicitly cascaded data (i.e. only modeling cases when there is a tie at positions  $i' < i$ ) as well as for regressing against the sign of the preference and observed identical findings. Although we omit those plots due to space constraints, they further support a backoff model interpretation of lexicprecision.

## 4 METHODS

In previous sections, we theoretically and conceptually connected  $RL_1$  to the notion of best-case retrieval evaluation, with a few illustrative empirical results. In order to rigorously test the viability of lexicprecision, we conducted a series of empirical analyses based on publicly available benchmarking data.<sup>3</sup>


### 4.1 Data

We analyzed the performance of lexicprecision across a variety of retrieval and recommendation tasks. Specifically, we collected runs submitted to TREC news (Robust 2004, Core 2017 and 2018), web (Web 2009-2014), and deep learning (Deep Learning 2019-2021) tracks as well as several public recommendation tasks [22].

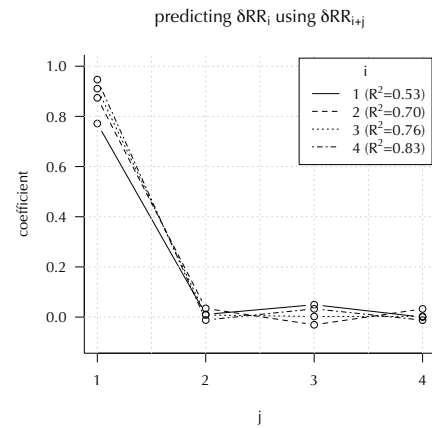
### 4.2 Analyses

Our empirical analyses were founded on two core questions, (i) how empirically correlated are lexicprecision and  $RL_1$  metrics, and (ii) how much more robust is lexicprecision than  $RL_1$  metrics. Because of its widespread adoption in the research community, we will use reciprocal rank for analyses. In order to answer the first question, we conducted experiments designed to predict the agreement between lexicprecision and  $RL_1$  metrics under different conditions. We considered two types of agreement. *Agreement in ranking preference* tests whether  $\pi >_{LP} \pi'$  agrees with  $\pi >_{RR} \pi'$ . Because lexicprecision is substantially more sensitive than  $RL_1$  metrics, we only consider situations where  $\delta RR_1(\pi, \pi') \neq 0$ . Because sgnLP and rrLP always agree in sign, we will only show results for one of the metrics when computing ranking agreement. *Agreement in system*

<sup>3</sup>Code for computing lexicprecision can be found at [https://github.com/diaz/pref\\_eval](https://github.com/diaz/pref_eval)



(a) Correlation between  $\delta RR_i$



(b) Regression of  $\delta RR_i$  using  $\delta RR_{i+1}, \dots, \delta RR_{i+4}$

**Figure 7: Locality of  $\delta RR_i$ . Relationship between the difference in reciprocal rank across recall levels using Robust 2004 runs. (a) Pearson and linear fit between all pairs of  $\delta RR_i$ . (b) Linear regression of  $\delta RR_i$  using  $\delta RR_{i+1:i+4}$  as independent variables. Regression shown for  $i = \{1, 2, 3, 4\}$ .**

*preference* tests whether  $\mathbb{E}_{q \sim Q} [\Delta_{LP}(\pi_q, \pi'_q)]$  agrees in sign with  $\mathbb{E}_{q \sim Q} [\Delta_{RR}(\pi_q, \pi'_q)]$ . This measures whether our choice of rrLP or sgnLP affects its correlation with reciprocal rank. Agreement is measured as a percentage of preferences agreed upon.

In order to assess the robustness of lexicprecision, we measure the number of ties observed amongst pairs of rankings and discriminative power. We claim that a robust approach has fewer ties and higher discriminative power. For discriminative power, we adopt Sakai’s approach of measuring the number of statistically significant differences between runs [18], using both Tukey’s honestly significant difference (HSD) test [3] and classic paired test to compute  $p$ -values. The paired test uses the Student’s  $t$ -test for reciprocal rank and rrLP [20]; and the binomial test for sgnLP.

Best-Case Retrieval Evaluation: Improving the Sensitivity of Reciprocal Rank with Lexicographic Precision

**Table 1: Ranking agreement between  $\delta\text{RR}_1$  and preferences based on the positions of the last  $m - 1$  relevant items. The computation of  $\text{sgnLP}$  in the table is based on the  $m - 1$  positions of relevant items after the top-ranked relevant item.**

	$\text{sgnLP}$	$\delta\text{RR}_2$
<b>news</b>		
robust (2004)	<b>85.78</b>	83.44
core (2017)	<b>89.23</b>	87.30
core (2018)	<b>88.01</b>	86.58
<b>web</b>		
web (2009)	<b>85.87</b>	84.79
web (2010)	<b>87.29</b>	85.41
web (2011)	<b>88.91</b>	87.54
web (2012)	<b>87.22</b>	85.45
web (2013)	<b>86.51</b>	84.45
web (2014)	<b>88.02</b>	85.82
<b>deep</b>		
deep-docs (2019)	<b>86.56</b>	83.10
deep-docs (2020)	<b>83.73</b>	79.34
deep-docs (2021)	<b>92.41</b>	89.78
deep-pass (2019)	<b>90.45</b>	88.87
deep-pass (2020)	<b>92.86</b>	91.08
deep-pass (2021)	<b>91.97</b>	90.14
<b>recsys</b>		
ml-1M (2018)	<b>78.90</b>	77.56
libraryThing (2018)	<b>66.50</b>	66.08
beerAdvocate (2018)	<b>58.84</b>	58.25

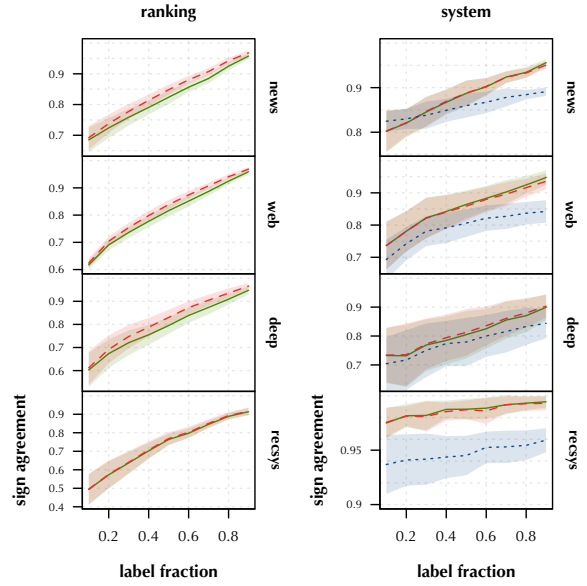
## 5 RESULTS

### 5.1 Correlation with Reciprocal Rank


By construction, we know that  $\delta\text{RR}_1(\pi, \pi') > 0 \implies \text{LP}(\pi, \pi') > 0$  and, so, the correlation between the two will be high. We can further test this by comparing how well lexicprecision predicts a ground truth preference between rankings based on  $\delta\text{RR}_1$ .

In our first analysis, given an observed  $\delta\text{RR}_1(\pi, \pi') \neq 0$ , we measure the ability of lexicprecision and reciprocal *based only on  $m - 1$  subsequent recall levels* to predict the sign of  $\delta\text{RR}_1(\pi, \pi')$ . That is, we use  $\delta\text{RR}_1(\pi, \pi')$  as a target value and compute  $\delta\text{RR}_1$  and LP using suffixes  $p_{2:m}$  and  $p'_{2:m}$ . Although artificial, this analysis provides an indication of the predictive value gained through cascaded modeling (as opposed to just looking at the top-ranked relevant item). We present the results in Table 1. As we can see, lexicprecision consistently agrees more with the target (masked)  $\delta\text{RR}_1$  than  $\delta\text{RR}_1$  of the suffix across all datasets, indicating that the additional information in higher recall levels can be used to predict the target (masked)  $\delta\text{RR}_1$ . This agrees with our preliminary analysis in Section 3.6.

We can also test the relationship between reciprocal rank and lexicprecision by measuring the agreement under incomplete information. Specifically, we consider removing either labels (treating unlabeled items as non-relevant) or requests (i.e. queries or users). We then measure the agreement between preferences with incomplete data and  $\delta\text{RR}_1$  on complete data (i.e. all requests and labels). Methods that agree more with reciprocal rank on complete data are considered more correlated. We present results for ranking and system agreement when removing labels (Figure 8a) and queries



(a) Removing labels.



(b) Removing queries.

**Figure 8: Preference agreement with  $\delta\text{RR}_1$  with full data. Labels and requests removed randomly. Results averaged across ten samples. Solid green lines:  $\delta\text{RR}_1$  with incomplete information. Dashed red lines: rrLP with incomplete information. Dotted blue lines:  $\text{sgnLP}$  with incomplete information. Shaded areas: one standard deviation across samples. Ranking agreement with incomplete labels for  $\text{sgnLP}$  is identical to rrLP and omitted for clarity.**

(Figure 8b). Across all conditions, we observe that the rrLP has as

**Table 2: Percentage of ties between pairs of rankings from two systems for the same request. We collapse rrLP and sgnLP for clarity.**


	rrLP, sgnLP	$\delta RR_1$
<b>news</b>		
robust (2004)	<b>0.39</b>	44.22
core (2017)	<b>0.23</b>	48.50
core (2018)	<b>1.72</b>	31.43
<b>web</b>		
web (2009)	<b>4.93</b>	15.13
web (2010)	<b>0.61</b>	25.85
web (2011)	<b>1.02</b>	41.99
web (2012)	<b>0.34</b>	34.01
web (2013)	<b>0.83</b>	31.09
web (2014)	<b>0.64</b>	41.93
<b>deep</b>		
deep-docs (2019)	<b>1.06</b>	68.45
deep-docs (2020)	<b>2.43</b>	73.99
deep-docs (2021)	<b>0.23</b>	80.84
deep-pass (2019)	<b>2.63</b>	56.89
deep-pass (2020)	<b>2.58</b>	50.30
deep-pass (2021)	<b>1.32</b>	47.41
<b>recsys</b>		
ml-1M (2018)	<b>3.38</b>	21.39
libraryThing (2018)	<b>16.48</b>	25.85
beerAdvocate (2018)	<b>41.73</b>	45.72

high or slightly higher agreement with  $\delta RR_1$  with complete information than  $\delta RR_1$  with incomplete information. This means that rrLP can accurately predict  $\delta RR_1$  with complete information as well or better than using reciprocal rank. Moreover, we observed that sgnLP shows weaker system agreement which occurs because its magnitude does not decay with rank position and, therefore, resulting averages are inconsistent with averages of position-discounted reciprocal rank values.

### 5.2 Sensitivity

In Section 2, we motivated our work by showing that  $RL_1$  metrics theoretically and empirically suffer from ceiling effects. The primary instrument we used to determine this was the probability of ties between rankings. In Table 2, we present the percentage of tied rankings from different systems for the same request. As predicted by our analysis in Section 3.3, lexiprecision has substantially fewer ties because this only happens when two rankings place relevant items in exactly the same positions.

In Section 3.3, we showed that lexiprecision implicitly and exponentially increased its fidelity as the number of relevant items  $m$  increased, while  $RL_1$  would quickly suffer from ties. In Figure 9, we show the number of tied rankings as a function of incomplete labels. This allows us to see trends with respect to  $m$ . Across our three retrieval benchmark sets, we see the growth in number of ties for  $RL_1$  as  $m$  increases; meanwhile, they shrink for lexiprecision. The drop in ties for recommender systems benchmarks suggests that, as described in Section 3.6, rankings contain very few relevant items and, as a result, removing labels will result in no relevant items present and increasingly tied rankings.



**Figure 9: Number of ties as labels are removed randomly. Results are averaged across ten samples. Solid green lines:  $\delta RR_1$  with incomplete information. Dashed red lines: rrLP with incomplete information. Shaded areas: one standard deviation across samples. Number of ties with incomplete labels for sgnLP is identical to rrLP and omitted for clarity.**

While the number of ties indicates that  $RL_1$  might not be able to distinguish systems, for a large enough sample of requests, a metric might still be good enough to distinguish systems. A different approach to measuring the discriminative power of an evaluation method is to count the number of differences that are statistically significant [18]. When we compare the percentage of pairs registering a statistically significant difference (Table 3), both rrLP and sgnLP outperform reciprocal rank, often by a very large margin. This indicates that the number of ties indeed hurts the ability of reciprocal rank to detect significant differences, while both variants of lexiprecision are much more sensitive.

## 6 DISCUSSION

Our results demonstrate that our lexiprecision variants capture the properties of  $RL_1$  while substantially increasing the ability to distinguish systems under the same best-case evaluation assumptions.

Practitioners and evaluators need to assess whether the assumptions behind  $RL_1$  metrics, including reciprocal rank, or lexiprecision or any other evaluation scheme are aligned with the use case. If a retrieval environment supports the assumptions behind  $RL_1$  metrics, including ties, then, by all means, they should be used to assess performance. However, in Section 3.1, we raised several reasons why uncertainty over recall requirements and psychological relevance suggest that  $RL_1$  metrics make quite strong assumptions not realized in most retrieval settings. We designed lexiprecision



**Table 3: Percentage of run differences detected at  $p < 0.05$ . Red: better than reciprocal rank. Bold: best for an evaluation setting.**

(a) Tukey’s HSD test				(b) Paired test with Bonferroni correction			
	rrLP	sgnLP	RR		rrLP	sgnLP	RR
<b>news</b>				<b>news</b>			
robust (2004)	<b>27.42</b>	27.34	23.55	robust (2004)	26.22	<b>27.36</b>	21.45
core (2017)	<b>17.41</b>	14.67	15.03	core (2017)	<b>16.22</b>	11.35	11.53
core (2018)	<b>28.60</b>	<b>31.42</b>	27.39	core (2018)	<b>29.30</b>	<b>31.73</b>	27.03
<b>web</b>				<b>web</b>			
web (2009)	23.85	<b>28.28</b>	24.11	web (2009)	<b>23.76</b>	<b>25.18</b>	23.49
web (2010)	<b>18.95</b>	13.51	18.35	web (2010)	<b>18.55</b>	9.27	17.74
web (2011)	<b>14.70</b>	10.22	13.83	web (2011)	<b>12.30</b>	6.94	9.73
web (2012)	<b>13.39</b>	11.61	13.21	web (2012)	<b>11.97</b>	10.11	11.35
web (2013)	5.85	5.79	<b>6.07</b>	web (2013)	<b>4.75</b>	<b>4.64</b>	4.32
web (2014)	<b>20.00</b>	11.72	18.85	web (2014)	<b>15.86</b>	7.13	14.02
<b>deep</b>				<b>deep</b>			
deep-docs (2019)	<b>8.25</b>	<b>19.20</b>	6.97	deep-docs (2019)	<b>11.66</b>	<b>16.36</b>	5.69
deep-docs (2020)	<b>5.26</b>	<b>3.47</b>	2.88	deep-docs (2020)	<b>2.33</b>	<b>1.79</b>	0.60
deep-docs (2021)	<b>6.39</b>	<b>11.19</b>	4.48	deep-docs (2021)	<b>3.73</b>	<b>9.14</b>	3.03
deep-pass (2019)	<b>16.52</b>	<b>18.47</b>	13.21	deep-pass (2019)	<b>15.02</b>	<b>17.42</b>	10.36
deep-pass (2020)	<b>37.46</b>	<b>40.91</b>	28.35	deep-pass (2020)	<b>39.04</b>	<b>39.45</b>	28.00
deep-pass (2021)	<b>24.07</b>	<b>24.78</b>	20.38	deep-pass (2021)	<b>23.55</b>	<b>20.99</b>	16.79
<b>recsys</b>				<b>recsys</b>			
ml-1M (2018)	<b>81.43</b>	<b>90.95</b>	80.00	ml-1M (2018)	90.00	<b>92.38</b>	90.48
libraryThing (2018)	93.81	<b>96.67</b>	93.81	libraryThing (2018)	<b>97.14</b>	<b>97.62</b>	96.67
beerAdvocate (2018)	<b>92.38</b>	<b>96.19</b>	90.95	beerAdvocate (2018)	94.76	<b>96.67</b>	94.76

to operate as conservatively as possible, preserving any preference from  $RL_1$  metrics and only acting to break ties.

Although  $RL_1$  metrics and lexiprecision agree perfectly when there is only one relevant item, this does not mean that all situations where we have a single judged relevant item should adopt a metric like reciprocal rank. For example, the MSMARCO dataset [16] includes requests and very sparse labels; the majority of requests have one *judged* relevant item. One might be tempted to use reciprocal rank but Arabzadeh et al. [1] demonstrate that this would obscure the multitude of *unjudged* relevant items (of which there are many). This hurts efficacy of best-case retrieval evaluation *including reciprocal rank*, as shown in Figures 8a and 9. Recommendation tasks have similar issues with sparsity due in part to it being more difficult for a third party to assess the relevance of personalized content and to the difficulty in gathering explicit feedback. Labels derived from behavioral feedback in general suffer from similar sparsity [2]. In this respect, we echo the call from Arabzadeh et al. [1] to make labeling practices across all of these domains much more robust. Given the observation of Voorhees et al. [23] that better labeling can result in less informative evaluation, we need to also develop more sensitive evaluation schemes such as lexiprecision.

This study has introduced a new preference-based evaluation method for  $RL_1$  metrics. As such, our focus has been on developing an understanding for comparing pairs of rankings and systems. We do not claim that lexiprecision itself is a metric and emphasize that we use it for comparing two rankings or systems. As such, although we address some concerns with reciprocal rank raised by Ferrante et al. [11], we do not make claims about lexiprecision being an interval measure. That said, the total ordering shown in Figure 1 suggests that there may be version of lexiprecision that can indeed be represented as an interval measure.

Finally, one can construct a metric  $\mu_{LP}$  such that  $\mu_{LP}(\pi) > \mu_{LP}(\pi') \rightarrow \text{sgnLP}(\pi, \pi') > 0$ . To do so, we can interpret  $\pi$  as a binary string based on item relevance and compute the integer representation,  $\mu_{LP}(\pi) = \sum_{i=1}^n \pi_i 2^{n-i}$ , where  $\pi_i$  is the binary relevance of the  $i$ ’th ranked item in  $\pi$ . This is proportional to  $\sum_{i=1}^n \pi_i \frac{1}{2}^{i-1}$ , which is equal to rank-biased precision [15] with a ‘patience parameter’  $\gamma = \frac{1}{2}$ . In fact,  $\mu_{LP}$  can be RBP with  $\gamma \in (0, \frac{1}{2}]$ . That said, there are a few differences between RBP and LP. First, the difference in RBP values will be sign-equivalent with LP but not equivalent in magnitude, which has implications for statistical testing. Second, lexicographic and preference-based evaluation has a clear conceptual grounding in population-based metrics [9], allowing connection to different fairness/robustness notions, which is where leximin and leximax originated. More importantly, lexiprecision is meant to complete the conceptual model covering average case [?] and worst case [9] preferences amongst a population of user groups defined by desired recall level [9, Section 4].

## 7 CONCLUSION

Motivated by ceiling effects in evaluation with reciprocal rank, we have attempted to increase our understanding of the metric and designed a well-grounded mitigation to conducting best-case retrieval evaluation. We have shown that lexiprecision can effectively address the limitations of reciprocal rank in retrieval evaluation. Our results highlight the importance of considering the effects of tie-breaking in the evaluation process and provide a method for conducting more reliable best-case retrieval evaluation. Given the use of retrieval metrics—including reciprocal rank—outside of information retrieval contexts, we believe these contributions will be relevant to a researchers in the broader research community.

## REFERENCES

- [1] Negar Arabzadeh, Alexandra Vtyurina, Xinyi Yan, and Charles L. A. Clarke. 2022. Shallow Pooling for Sparse Labels. *Inf. Retr.* 25, 4 (dec 2022), 365–385. <https://doi.org/10.1007/s10791-022-09411-0>
- [2] Michael Bendersky, Xuanhui Wang, Marc Najork, and Donald Metzler. 2018. Learning with Sparse and Biased Feedback for Personal Search. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*. International Joint Conferences on Artificial Intelligence Organization, 5219–5223. <https://doi.org/10.24963/ijcai.2018/725>
- [3] Benjamin A. Carterette. 2012. Multiple testing in statistical analysis of systems-based information retrieval experiments. *ACM Trans. Inf. Syst.* 30, 1, Article 4 (March 2012), 34 pages. <https://doi.org/10.1145/2094072.2094076>
- [4] Pablo Castells and Alistair Moffat. 2022. Offline recommender system evaluation: Challenges and new directions. *AI Magazine* 43, 2 (2022), 225–238. <https://doi.org/10.1002/aaai.12051> arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1002/aaai.12051>
- [5] Praveen Chandar, Fernando Diaz, and Brian St. Thomas. 2020. Beyond Accuracy: Grounding Evaluation Metrics for Human-Machine Learning Systems. <https://github.com/pchandar/beyond-accuracy-tutorial>. In *Advances in Neural Information Processing Systems*.
- [6] Olivier Chapelle, Donald Metzler, Ya Zhang, and Pierre Grinspan. 2009. Expected reciprocal rank for graded relevance. In *Proceedings of the 18th ACM conference on Information and knowledge management (Hong Kong, China) (CIKM '09)*. ACM, New York, NY, USA, 621–630. <https://doi.org/10.1145/1645953.1646033>
- [7] William S. Cooper. 1968. Expected search length: A single measure of retrieval effectiveness based on the weak ordering action of retrieval systems. *American Documentation* 19, 1 (1968), 30–41. <https://doi.org/10.1002/asi.5090190108>
- [8] Fernando Diaz and Andres Ferraro. 2022. Offline Retrieval Evaluation Without Evaluation Metrics. In *Proceedings of the 45th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*.
- [9] Fernando Diaz and Bhaskar Mitra. 2023. Recall, Robustness, and Lexicographic Evaluation. arXiv:2302.11370 [cs.IR]
- [10] Fernando Diaz and Bhaskar Mitra. 2023. Recall, Robustness, and Lexicographic Evaluation. (2023). under review for ACM Transactions on Information Systems.
- [11] Marco Ferrante, Nicola Ferro, and Norbert Fuhr. 2021. Towards Meaningful Statements in IR Evaluation: Mapping Evaluation Measures to Interval Scales. *IEEE Access* 9 (2021), 136182–136216. <https://doi.org/10.1109/ACCESS.2021.3116857>
- [12] Stephen P. Harter. 1992. Psychological relevance and information science. *Journal of the American Society for Information Science* 43, 9 (1992), 602–615.
- [13] Paul B Kantor and Ellen Voorhees. 1997. Report on the TREC Confusion Track. In *Proceedings of The Fifth Text REtrieval Conference (TREC-5)*.
- [14] S. Katz. 1987. Estimation of probabilities from sparse data for the language model component of a speech recognizer. *IEEE Transactions on Acoustics, Speech, and Signal Processing* 35, 3 (1987), 400–401. <https://doi.org/10.1109/TASSP.1987.1165125>
- [15] Alistair Moffat and Justin Zobel. 2008. Rank-biased Precision for Measurement of Retrieval Effectiveness. *ACM Trans. Inf. Syst.* 27, 1, Article 2 (Dec. 2008), 27 pages. <https://doi.org/10.1145/1416950.1416952>
- [16] 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. (November 2016). <https://www.microsoft.com/en-us/research/publication/ms-marco-human-generated-machine-reading-comprehension-dataset/>
- [17] Stephen Robertson. 2008. A New Interpretation of Average Precision. In *Proceedings of the 31st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (Singapore, Singapore) (SIGIR '08)*. Association for Computing Machinery, New York, NY, USA, 689–690. <https://doi.org/10.1145/1390334.1390453>
- [18] Tetsuya Sakai. 2014. Metrics, statistics, tests. In *Bridging Between Information Retrieval and Databases - PROMISE Winter School 2013, Revised Tutorial Lectures (Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics))*. Springer Verlag, 116–163. [https://doi.org/10.1007/978-3-642-54798-0\\_6](https://doi.org/10.1007/978-3-642-54798-0_6) 2013 PROMISE Winter School: Bridging Between Information Retrieval and Databases ; Conference date: 04-02-2013 Through 08-02-2013.
- [19] Amartya Sen. 1970. *Collective Choice and Social Welfare*. Holden-Day.
- [20] 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 (Lisbon, Portugal) (CIKM '07)*. Association for Computing Machinery, New York, NY, USA, 623–632. <https://doi.org/10.1145/1321440.1321528>
- [21] Daniel Valcarce, Alejandro Bellogín, Javier Parapar, and Pablo Castells. 2018. On the Robustness and Discriminative Power of Information Retrieval Metrics for Top-N Recommendation. In *Proceedings of the 12th ACM Conference on Recommender Systems (Vancouver, British Columbia, Canada) (RecSys '18)*. Association for Computing Machinery, New York, NY, USA, 260–268. <https://doi.org/10.1145/3240323.3240347>
- [22] Daniel Valcarce, Alejandro Bellogín, Javier Parapar, and Pablo Castells. 2020. Assessing ranking metrics in top-N recommendation. *Information Retrieval Journal* 23, 4 (2020), 411–448. <https://doi.org/10.1007/s10791-020-09377-x>
- [23] Ellen M. Voorhees, Nick Craswell, and Jimmy Lin. 2022. Too Many Relevant: Whither Cranfield Test Collections?. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval (Madrid, Spain) (SIGIR '22)*. Association for Computing Machinery, New York, NY, USA, 2970–2980. <https://doi.org/10.1145/3477495.3531728>