

Towards Fair RAG: On the Impact of Fair Ranking in Retrieval-Augmented Generation

To Eun Kim
Carnegie Mellon University
Pittsburgh, PA, USA
toeunk@cs.cmu.edu

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

Abstract

Despite the central role of retrieval in retrieval-augmented generation (RAG) systems, much of the existing research on RAG overlooks the well-established field of fair ranking and fails to account for the interests of all stakeholders involved. In this paper, we conduct the first systematic evaluation of RAG systems that integrate fairness-aware rankings, addressing both ranking fairness and attribution fairness, which ensures equitable exposure of the sources cited in the generated content. Our evaluation focuses on measuring item-side fairness, specifically the fair exposure of relevant items retrieved by RAG systems, and investigates how this fairness impacts both the effectiveness of the systems and the attribution of sources in the generated output that users ultimately see. By experimenting with twelve RAG models across seven distinct tasks, we show that incorporating fairness-aware retrieval often maintains or even enhances both ranking quality and generation quality, countering the common belief that fairness compromises system performance. Additionally, we demonstrate that fair retrieval practices lead to more balanced attribution in the final responses, ensuring that the generator fairly cites the sources it relies on. Our findings underscore the importance of item-side fairness in retrieval and generation, laying the foundation for responsible and equitable RAG systems and guiding future research in fair ranking and attribution.

CCS Concepts

- Computing methodologies → Natural language generation;
- Information systems → Evaluation of retrieval results.

Keywords

Fair Ranking, Fair Attribution, Retrieval-Augmented Generation

ACM Reference Format:

To Eun Kim and Fernando Diaz. 2025. Towards Fair RAG: On the Impact of Fair Ranking in Retrieval-Augmented Generation. In *Proceedings of the 2025 International ACM SIGIR Conference on Innovative Concepts and Theories in Information Retrieval (ICTIR '25)*, July 18, 2025, Padua, Italy. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3731120.3744599>

1 Introduction

In recent years, the concept of fair ranking has emerged as a critical concern in modern information access systems [13]. However,




Figure 1: Fairness concerns in RAG. A simplified example of how RAG models can ignore equally relevant items (d_3 and d_4) and always consume the fixed top-scoring items (d_1 and d_2) with the same order of ranking over the multiple user requests. This is due to the deterministic nature of the retrieval process and a short context-length of a language model that necessitates the top- k truncation of a ranked list.

despite its significance, fair ranking has yet to be thoroughly examined in the context of retrieval-augmented generation (RAG) [1, 33], a rapidly advancing trend in natural language processing (NLP) systems [32]. To understand why this is important, consider the RAG system in Figure 1, where a user asks a question about running shoes. A classic retrieval system might return several documents containing information from various running shoe companies. If the RAG system only selects the top two documents, then information from the remaining two relevant companies will not be relayed to the predictive model and will likely be omitted from its answer. The fair ranking literature refers to this situation as unfair because some relevant companies (i.e., in documents at position 3 and 4) receive less or no exposure compared to equally relevant company in the top position [13].

Understanding the effect of fair ranking in RAG is fundamental to ensuring responsible and equitable NLP systems. Since retrieval results in RAG often underlie response attribution [17], *unfair exposure of content* to the RAG system can result in incomplete evidence in responses (thus compromising recall of potentially relevant information for users) or downstream representational harms (thus creating or reinforcing biases across the set of relevant entities). In situations where content providers are compensated for contributions to inference, there can be financial implications for the unfairness [4, 22, 36]. Indeed, the fair ranking literature indicates that these are precisely the harms that emerge when *people* are searchers [13], much less RAG systems, where the searchers are *machines*. RAG complicates these challenges since it often truncates rankings to much shorter lengths to fit the generator's limited context size [2, 23, 32], making equal exposure of relevant items even harder.



This work is licensed under a Creative Commons Attribution 4.0 International License. *ICTIR '25, Padua, Italy*
© 2025 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-1861-8/2025/07
<https://doi.org/10.1145/3731120.3744599>

Moreover, the fact that machines are now the searchers necessitates a different notion of item-worthiness (how deserving an item is to be included in a ranked list). Traditionally, ranking quality has been assessed based on relevance labels, which are created according to how relevant an item is to the user's query [53]. However, with RAG systems, where the consumer is a language model, there is a growing shift towards evaluating ranking quality based on *utility labels*, which are determined by the usefulness of an item in aiding the model's task performance, rather than its relevance to the query [51, 62].

This shift from relevance to utility in the concept of item-worthiness can significantly alter our understanding of the relationship between fairness and ranking quality [3]—particularly the tradeoffs that are well-known in the fair ranking literature [5, 11, 56]. Since previous fair ranking studies were conducted based on relevance judgments, they may need to be reexamined in light of utility-based judgments within the context of RAG.

However, purely focusing on how often certain items appear in the top- k positions can neglect the fact that not all retrieved items are necessarily *attributed* in the final generated response. If an item is retrieved but never actually influences the RAG model's output, one cannot fully gauge whether it truly received exposure from the standpoint of the final generation. This reveals a subtle yet important gap: *fair retrieval* may not directly translate to *fair consumption*, especially when some retrieved items might be overshadowed in the generation step. Measuring how exposure is ultimately distributed across the attributed sources in the final response offers a more complete picture of exposure-based fairness in the context of RAG systems.

Our research aims to bridge the gap between traditional fair ranking studies and the emerging changes posed by RAG systems, ultimately enhancing our understanding of the interplay between fairness, ranking quality, and the effectiveness of RAG systems. We do this by evaluating RAG systems with a fairness-aware retriever across seven different tasks, experimenting with varying levels of retrieval fairness to observe changes in ranking quality and generation quality (utility)¹, as well as the fairness of attributed sources.

Our empirical results show that, in the context of machine users, there also exists an overall trend of fairness-quality tradeoff with respect to both retrieval and generation quality. However, the magnitude of this tradeoff is not particularly severe. In fact, we find that RAG models equipped with a fair ranker can often preserve a significant level of retrieval and generation quality, and in some cases, even surpass the quality achieved by the traditional RAG setup with a deterministic ranker that lacks fairness considerations. Moreover, while the fraction of retrieved sources that actually appear in the final response may vary, equitable retrieval frequently leads to more equitable usage of those sources by the generator.

This surprising finding offers significant insight into the potential of RAG-based applications, suggesting that fair treatment of individual content providers can be achieved without sacrificing much of the high-quality service delivered to end-users. This challenges the conventional assumption of an inevitable tradeoff

¹Throughout this paper, we use "utility" and "generation quality" interchangeably to refer to the downstream effectiveness of RAG models, measured by arbitrary string utility metrics.

between fairness and quality, opening new avenues for developing more equitable and effective RAG systems.

2 Background & Related Work

2.1 Retrieval-Augmented Generation

RAG, a specific type of retrieval-enhanced machine learning (REML) [32, 60], has been widely adopted in various domains, including language modeling [31], question-answering [27], and personalization [40, 49, 50]. Studies on the evaluation of RAG models have primarily focused on their effectiveness, including end-to-end performance [20, 27, 33] and the assessment of individual components [14, 47, 48, 51], such as retrieval relevance and model faithfulness. Furthermore, recent efforts have explored attribution mechanisms for ensuring the trustworthiness of RAG responses [16, 17], examining how thoroughly models reference the source text and thus promoting a more faithful generation process.

However, little research has focused on evaluating fairness in retrieval-enhanced generation models, with the exception of recent work [54], which improved demographic diversity in human image generation by conditioning a generative model with externally retrieved images that help debias the generation process.

2.2 Fairness in Ranking

Fair ranking has been approached through various definitions based on normative concerns, primarily with distinctions made according to the stakeholders we prioritize. These include consumer-side fairness [12, 38], which focuses on how fairly a system delivers satisfaction to users; provider-side fairness [28, 52], which addresses how fairly item providers receive monetary or reputational rewards; and item-side fairness [29], which examines how fairly items are treated in terms of representation or exposure. The motivation of item-side fairness is closely linked to provider-side fairness, as unfair treatment of items can lead to unfair compensation for providers. These fairness concerns can be further categorized by the scope of stakeholders, encompassing individual fairness—ensuring similar treatment for similar individuals—and group fairness—ensuring equitable outcomes across different groups [8, 13]. Previous studies have focused on developing metrics to measure fairness [44] and optimizing fair retrievers within a single [52, 58, 61] or multiple rankings [5, 11, 55, 56]. In the context of provider- and item-side fairness, ensuring equal exposure of similar items across multiple rankings has gained significant attention [13]. To achieve this, researchers have used stochastic rankers that return a distribution of rankings, in contrast to deterministic rankers commonly found in areas like RAG, which produce a fixed ranking. This approach ensures that, in expectation, similar items receive equal exposure across multiple user requests, with the distributions typically based on the merit of the rankings, such as an item's relevance [11, 56].

In this research, we employ a stochastic ranker in RAG to enhance *individual item-side fairness*, aiming to ensure equal expected exposure for items that offer similar merits.

3 Experimental Methodology

In traditional RAG systems, a user input is used to query a retrieval system for recommended items from some corpus, which are then used for generation. Given user input x , a query q generated by the

query generation function $\phi_q(x)$, and a corpus of documents C , a *deterministic retriever* $\mathcal{R}(q, C)$ returns a fixed ranked list L every time q is seen. Retrieval is followed by a top k truncation which is passed to a prompt generation function $\phi_p(x, L_{1:k})$ that returns a final prompt \bar{x} , which is subsequently passed to the language model $\mathcal{G}(\bar{x})$. Because deterministic retrievers allocate exposure to the same item over repeated samples, RAG systems with deterministic retrievers present a challenge to ensuring equal exposure of relevant items to the generator.

To address the issue of unfairness in the rankings passed to the generator, we can convert a deterministic retriever into a stochastic retriever, which can, in expectation, provide fair rankings [11]. By sampling a ranking based on its quality to users—in this case generators—the expected exposure of different relevant items becomes similar and, therefore, fairer. Because decisions are stochastic, the fairness and quality of stochastic retrieval is evaluated based on a sample of rankings. Likewise, since each sampled ranking is processed by a generator, we measure the expected generator effectiveness over these sampled rankings. We also consider how fairly exposure is allocated among the items that end up being attributed by the final output, acknowledging that some retrieved items may not be credited in the generated text. The complete evaluation pipeline of a RAG system with a stochastic retriever is illustrated in Figure 2.

The following sections describe how we construct a test collection with utility labels (§3.1), how we stochastically sample multiple rankings (§3.2), how we evaluate the fairness and ranking quality of the sampled rankings (§3.3.1), and how we assess a RAG system’s performance over sampled rankings (§3.3.2). We then explain how we measure the expected number of attributed items (§3.3.3) and evaluate the fairness of these attributed sources (§3.3.4).

3.1 Construction of a Test Collection with Utility Labels

Setting an appropriate proxy for measuring item-worthiness is crucial in the evaluation of fairness [3]. Drawing on the insight that utility-based judgments are more suitable than relevance judgments in the context of RAG [51, 62], we annotate item-level utility labels for all items in the corpus.

We define an item’s worthiness by the marginal gain in utility (utility-gain) it provides to a language model (specifically, the generator in a RAG system) when used to solve a specific task as part of the augmentation process. To assess this utility-gain, each item in the corpus is individually supplied to the generator along with an input question. The utility-gain is then calculated as the difference between the utility of the augmented generator and that of a baseline language model without any information about the item. Formally, let u_i denote the baseline string utility score from the vanilla language model prompted only with the input question, and let u_j represent the utility score from the language model with a prompt augmented by the j ’th item d_j in the corpus. The item d_j is considered useful if the utility-gain $\delta_j = u_j - u_i$ is positive, and not useful otherwise.

Therefore, the item-level utility labels are designed to be both task- and generator-dependent, as the utility of each item varies depending on the task and the language model used. This labeling

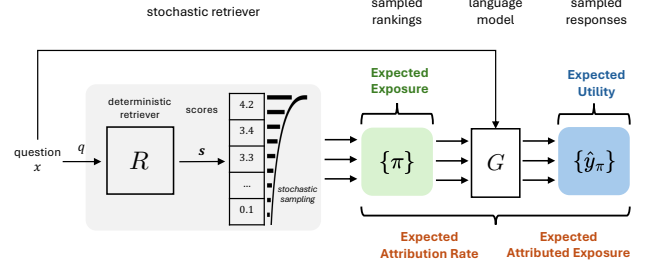


Figure 2: Overview of our experimental design for examining how item-fairness in retrieval impacts both ranking and generation quality, as well as the fairness of attributed items in retrieval-augmented generation. To evaluate system performance across multiple identical user requests, we sample N rankings from a stochastic retriever. Then, we measure the fairness and quality of these rankings (*Expected Exposure* §3.3.1), and then assess the system’s overall performance (*Expected Utility* §3.3.2), the usage of retrieved evidence (*Expected Attribution Rate* §3.3.3), and the fairness of attributed sources (*Expected Attributed Exposure* §3.3.4). The query and prompt generators are omitted for brevity.

process also aligns with the principles of task-based information retrieval, where, in the context of *human* searchers, document utility may vary on how the user expects to use the document [30].

3.2 Fairness-Aware Stochastic Retriever

Stochastic retrievers have been used for various purposes, such as optimization of retrieval models [7, 18, 41, 59], as well as ensuring equitable exposure of items [11, 41, 42]. Many of these studies use Plackett-Luce sampling [43] to achieve the stochasticity of retrieval. We follow the line of research and formally define how we derive a fairness-aware stochastic retriever through Plackett-Luce sampling. To enhance sampling efficiency, we adopt the methodology of Oosterhuis [41], and for controllable randomization, we utilize the approach proposed by Diaz et al. [11].

Given n items in a corpus C , a vector of retrieval scores $\mathbf{s} \in \mathbb{R}^n$ can be obtained from $\mathcal{R}(q, C)$, which can be used to generate a ranked list L . We then min-max normalize retrieval scores to be in $[0, 1]$ in order to construct a multinomial distribution over items [5]. The probability of an item d being selected as the i ’th item in a new ranking π through Plackett-Luce sampling is given by

$$p(d|L_{1:i-1}) = \frac{\exp(\bar{s}_d) \mathbb{1}[d \notin L_{1:i-1}]}{\sum_{d' \in C \setminus L_{1:i-1}} \exp(\bar{s}_{d'})} \quad (1)$$

where $L_{1:i-1}$ is the partial ranking up to position $i - 1$, $\bar{\mathbf{s}}$ represents the normalized retrieval score vector, and \bar{s}_d is the normalized score of item d . Using this probability, we iteratively sample an item, set its probability to 0, renormalize the distribution, and repeat the process. The probability of generating a complete ranking is then given by the product of the placement probabilities for each item, i.e., $p(\pi|q) = \prod_{i=1}^n p(\pi_i|\pi_{1:i-1})$.

This repeated sampling and renormalization process can be efficiently managed using the Gumbel-Softmax trick [19, 37], which enables the sampling of rankings to be performed at the speed

of sorting [41]. To do so, for each sampling iteration, we draw $U_i \sim \text{Uniform}(0, 1)$, followed by generating a Gumbel noise $G_i = -\log(-\log(U_i))$. The probability of each sampled ranking is then obtained by sorting the items based on their perturbed scores $\tilde{s}_{d_i} = \bar{s}_{d_i} + G_i$.

3.2.1 Controlling the Level of Fairness. Adjusting the level of randomization directly controls the degree of item-fairness, aligning with our goal to observe how varying levels of fairness in rankings affect the ranking and generation quality of a RAG model. To obtain the controllability, we follow the work of Diaz et al. [11] and use a fairness control parameter α . We apply the scalar α to each value in the normalized score vector \bar{s} by raising each value to the power of α .² This process is done before the scores are passed to the sampling policy. Therefore, the modified sampling distribution is thus defined as:

$$p(d|L_{1:i-1}) = \frac{\exp(\bar{s}_d^\alpha) \mathbb{1}[d \notin L_{1:i-1}]}{\sum_{d' \in C \setminus L_{1:i-1}} \exp(\bar{s}_{d'}^\alpha)} \quad (2)$$

This implies that the sharpness of the sampling distribution is controlled by the α . A higher α amplifies the probability of items with higher retrieval scores being sampled. Therefore, if multiple rankings are sampled by the stochastic retriever with high α , it results in high disparity (i.e., item-side unfairness) of sampled rankings. At extreme, with considerably high α , the procedure results in the identical rankings which is the behavior of a deterministic ranker (i.e., maximum item-unfairness). On the other hand, a lower α reduces the disparity of sampled rankings, making the exposure distribution fairer. At extreme, when $\alpha = 0$, the sampling procedure becomes uniformly random and achieves the lowest disparity (i.e., maximum item-fairness) in the sampled rankings.

3.3 Evaluation

As mentioned in Section 3, because we are dealing with stochastic retrievers, we need to measure the *expected* behavior of the system. Let $\mathcal{S}(s, N, k)$ be the stochastic sampler that samples a set of N rankings $\sigma = \{\pi\}$, given the deterministic retrieval scores s , where each ranking π is truncated to the size of k . From each ranking, we can get an output $\hat{y}_\pi = \mathcal{G}(\phi_p(x, \pi))$. With an arbitrary fairness metric $\mu_f(\sigma)$ and a ranking quality metric $\mu_r(\sigma)$ that takes a set of rankings as an input, we can measure the degree of fairness and ranking quality of the sampled rankings. Similarly, an arbitrary string utility metric $\mu_u(y, \hat{y}_\pi)$, such as ROUGE, can be used to assess an expected effectiveness of a RAG system by calculating the average of the N metric scores.

In this paper, based on the empirical investigation done by Raj and Ekstrand [44], we use expected exposure disparity (EE-D) and expected exposure relevance (EE-R) [11] as μ_f and μ_r , respectively (§3.3.1). For μ_u , we select the metric depending on the task, and we get the expectation of the utility of a RAG model which we call an expected utility (EU) (§3.3.2). Beyond these metrics, we also measure the expected rate of attributed items (EAR) in §3.3.3, capturing how many retrieved items are ultimately used by the generator, and introduce the *expected attributed exposure* (EAE) in §3.3.4, which evaluates fairness specifically among the items that are actually

²We normalized the values to the range of $[1, 2]$ instead of $[0, 1]$. The addition of 1 effectively serves the same purpose as adjusting a real-numbered α . We chose this range to allow for an integer-valued α .

attributed in the final output that can ultimately be displayed to human users.

3.3.1 Expected Exposure in the Context of Machine Users. Expected Exposure (EE) [11] works by estimating the exposure of items across rankings (e.g., σ) created by a subject model, and comparing them with an optimal set of rankings that always satisfy the item-fairness. To represent the attention over n items in the corpus given by the consumer (generator in RAG), an $n \times 1$ system exposure vector ϵ is created. This is then compared with an $n \times 1$ target exposure vector ϵ^* , where it represents the exposure of items allocated by an oracle retriever that always rank useful items above non-useful ones [11].

With the system and target exposure vector $\epsilon \in \mathbb{R}^n$ and $\epsilon^* \in \mathbb{R}^n$, we can get the difference between the two by the squared l_2 distance:

$$\|\epsilon - \epsilon^*\|_2^2 = \|\epsilon\|_2^2 - 2\langle \epsilon, \epsilon^* \rangle + \|\epsilon^*\|_2^2 \quad (3)$$

This difference yields two metrics useful for fairness and ranking quality evaluation. $\|\epsilon\|_2^2$ can be a measure for disparity of rankings (EE-D), and $\langle \epsilon, \epsilon^* \rangle$ can be a measure of ranking quality (EE-R) by calculating the degree of alignment of system exposure to the target exposure (i.e., how much of the exposure is on useful items). Therefore, the higher the value of EE-D, the more unfair the set of rankings are, and the higher the value of EE-R, the closer the set of system rankings are to the optimal set of rankings with respect to the ranking quality.

The exposure of an item is calculated by modeling users' (e.g., generators in RAG) attention to each item in a ranking. For example, one can assume that the user is affected by position bias and gives attention following an exponential decay [39]. However, these browsing models were developed for human-users not for machine-users, so we need a different user behavior model for generators in RAG. For the simplicity of the metric, and given recent efforts to reduce position bias and promote more even attention in machine-user settings [21, 25], we assume that the machine user consumes all items passed to the context with equal attention but pays zero attention to items placed after the k 'th position due to top- k truncation. This makes the user browsing model a step function parameterized by k . In this work, a relevance-independent machine-user model (MU) is set to the step function that reflects the behavior of top- k truncation of RAG:

$$\text{MU}(i) = \begin{cases} 1 & \text{if } i \leq k \\ 0 & \text{otherwise} \end{cases} = \mathbb{1}[i \leq k] \quad (4)$$

Given this machine user browsing model and a mapping from item index to its rank denoted as $\tilde{\pi}_d$, a system exposure for each item d is calculated as

$$\epsilon_d = \sum_{\pi \in S_n} p(\pi|q) \text{MU}(\tilde{\pi}_d) \quad (5)$$

and target exposures for a useful item d and a useless item d^- are calculated as

$$\epsilon_d^* = \frac{1}{m} \sum_{i=1}^m \text{MU}(i) = \begin{cases} 1 & \text{if } m \leq k \\ \frac{k}{m} & \text{otherwise} \end{cases} \quad \epsilon_{d^-}^* = \begin{cases} \frac{k-m}{n-m} & \text{if } m \leq k \\ 0 & \text{otherwise} \end{cases} \quad (6)$$

where m is the number of useful items in the corpus of size n .

3.3.2 Expected Utility. Given the set of N sampled rankings σ , we individually augment the generator with each ranking $\pi \in \sigma$, resulting in N outputs from the generator. The utility of these outputs is then measured using an arbitrary string utility metric μ_u . To determine the anticipated utility of a RAG model with fair rankings—represented by the tuple of a stochastic ranking sampler \mathcal{S} and a generator \mathcal{G} —we calculate the expected utility (EU) of the RAG system given an instance x .

$$\begin{aligned} \text{EU}(\langle \mathcal{S}, \mathcal{G} \rangle | x) &= \mathbb{E}_{\pi \sim \mathcal{S}} [\mu_u(y, \hat{y}_\pi)] \\ &= \sum_{\pi \in S_n} p(\pi | q) \mu_u(y, \hat{y}_\pi) \approx \frac{1}{N} \sum_{\pi \in \sigma} \mu_u(y, \hat{y}_\pi) \end{aligned} \quad (7)$$

where \hat{y}_π is the prediction of a system given the ranking π , S_n is the symmetric group of a ranked list L from the deterministic retriever \mathcal{R} , and $\sum_{\pi \in S_n} p(\pi | q) = 1$.

3.3.3 Expected Attribution Rate. We use a natural language inference (NLI) model for attribution, reflecting its effectiveness in measuring faithfulness. Honovich et al. [24] find that NLI outperforms other metrics in a meta-evaluation, leading works such as automated citation evaluation using NLI [17], which strongly correlates with human judgments. Building on the human evaluation framework by Rashkin et al. [45], Gao et al. [16] introduced a new NLI-based metric to approximate human judgments, followed by Bohnet et al. [6]. These studies collectively show that NLI-based methods capture attribution quality in a manner closely aligned with human assessments, making them a reliable choice for our evaluation.

Given a ranking $\pi \in \sigma$ of k items and a predicted output \hat{y}_π , we therefore measure item *attribution rate* (AR) using

$$\mu_{\text{AR}}(\pi, \hat{y}_\pi) = \frac{1}{k} \sum_{d \in \pi} \text{NLI}(d, \hat{y}_\pi) \quad (8)$$

where $\text{NLI}(d, \hat{y}_\pi) = 1$ if item d entails the output \hat{y}_π , and 0 otherwise, by a natural language inference model.³

Analogous to expected utility, we define expected attribution rate (EAR) of a RAG system as

$$\begin{aligned} \text{EAR}(\langle \mathcal{S}, \mathcal{G} \rangle | x) &= \mathbb{E}_{\pi \sim \mathcal{S}} [\mu_{\text{AR}}(\pi, \hat{y}_\pi)] \\ &= \sum_{\pi \in S_n} p(\pi | q) \mu_{\text{AR}}(\pi, \hat{y}_\pi) \approx \frac{1}{N} \sum_{\pi \in \sigma} \mu_{\text{AR}}(\pi, \hat{y}_\pi). \end{aligned} \quad (9)$$

3.3.4 Expected Attributed Exposure. While EE measures retrieval fairness and relevance, it does not capture whether the retrieved items actually appear in the final generated output or how fairly exposure is allocated among those items once they surface after generation. To address this gap, we define the *attributed exposure* vector $\epsilon^a \in \mathbb{R}^n$, where an attributed exposure for each item d is calculated as

$$\epsilon_d^a = \sum_{\pi \in S_n} p(\pi | q) \text{NLI}(d, \hat{y}_\pi). \quad (10)$$

³This metric is similar to the context utilization metric from Ru et al. [47], which also measures how effectively the generator leverages retrieved context. However, Ru et al. [47] decompose the generated response into multiple claims and focus only on the proportion of relevant items, whereas our approach applies a single NLI check for each retrieved item.

Similar to EE, a disparity measure *expected attributed exposure* disparity (EAE-D) can be derived by measuring the squared l_2 distance of an attributed exposure vector, thus

$$\text{EAE-D}(\langle \mathcal{S}, \mathcal{G} \rangle | x) = \|\epsilon^a\|_2^2. \quad (11)$$

Hence, EAE-D measures how fairly exposure is allocated across the items that are explicitly used, addressing the shortcoming of EE by directly linking exposure to items that influence the final output of a RAG system.

3.3.5 Normalization of Metrics. Since EAR is an average over the percentage of entailed items for each sampled ranking, its values lie on a consistent scale with $[0, 1]$ range. However, the bounds of EE-D, EE-R, and EAE-D depend on the number of useful items in the corpus. Consequently, we apply normalization on a per-query basis by min-max scaling each metric according to its theoretical lower and upper bounds. We denote the normalized EE-D, EE-R, and EAE-D as $\overline{\text{EE-D}}$, $\overline{\text{EE-R}}$, and $\overline{\text{EAE-D}}$ respectively.

However, theoretically determining the bounds of the expected utility (EU) of a RAG model is challenging. To address this, we normalized the EU by the model’s empirical upper bound, the maximum observed utility across all runs of the experiment with the same generators. To approach the true upper bound, these runs include RAG models with an oracle retriever that consistently ranks useful items (i.e., those with positive utility labels) above non-useful ones, stochastically returning one of the $m!(n-m)!$ different rankings, where m represents the number of useful items in the corpus. We denote the normalized EU as $\overline{\text{EU}}$, which can be interpreted as the distance to the optimal utility.

4 Experiments

We choose the LaMP benchmark [50] for our dataset. It assesses the personalization capability of language models through retrieval-augmentation of users’ interaction history in a platform. LaMP includes various prediction tasks, such as classification, regression, and generation, and is well-suited for tasks where multiple items can be relevant/useful, unlike QA tasks with typically one or two provenance items. The retrieval items in LaMP have clear providers and consumers, aligning with our goal to ensure fairness for individual item providers. For example, in LaMP-1, retrieval items are academic papers, where exposure can increase citation counts for authors. In LaMP-4, retrieval items are news articles, where exposure can lead to monetary compensation for journalists.⁴ Due to the absence of a test set, we constructed a test collection as described in §3.1, using the first thousand entries of a user-based development set. Then, we discarded entries that have only one useful item in the corpus, as it is unnecessary to concern item-fairness in that case. We release the test collection along with dataset statistics.

We use BM25 (lexical retriever) [46], SPLADE (learned sparse retriever) [15], and Contriever (bi-encoder dense retriever) [26] as deterministic retrievers providing retrieval scores to base the sampling on. These models represent commonly used retrievers in the RAG literature [32]. We use a sampling size of $N = 100$ and a truncation size of $k = 5$.

⁴Throughout this paper, we use LaMP-1 to represent classification and LaMP-4 to represent text generation in visualizations. All LaMP tasks were used in our experiments, and similar trends were observed.




Figure 3: Effect of a fairness control parameter (α) on the disparity of rankings ($\overline{EE-D}$) in LaMP task 1 and 4. The RAG model is configured with the Contriever and Flan-UL2. Each data point represents the normalized EE-D of each run of the experiment (i.e., one query $\rightarrow N$ sampled rankings $\rightarrow \overline{EE-D}$ of the N rankings).

For generation models, we use Flan-T5-Small, Flan-T5-Base, Flan-T5-XXL [9], and Flan-UL2 [57]. For decoding strategy, beam size is set to 4, and no sampling strategy is used. This is to ensure that stochasticity is only introduced to the retriever for controlled experiments. With the three base retrievers and four generators, we configure twelve different RAG models and evaluate them on the seven LaMP tasks. Utility measurement of the generated strings follows the metrics used in the LaMP paper. For NLI computations, a RoBERTa large model [35] fine-tuned on the natural language inference task is used.⁵

We repeat the experiments with four different fairness control parameters $\alpha = 1, 2, 4, 8$, which allows us to assess the utility of the RAG models with different levels of item-fairness. From Figure 3, we observe how effectively α , described in the Equation 2, controls the disparity of rankings. For example, when α is set to 4, we usually obtain a set of sampled rankings with $\overline{EE-D}$ mostly in the range of $[0.5, 0.8]$, and when α is set to 8, we often get a set of sampled rankings with $\overline{EE-D} = 1$.

We conducted paired (per-query) t -tests for $\overline{EE-D}$ across the different α values and found statistically significant differences ($p < 0.01$) in all 84 experimental conditions (12 models \times 7 tasks).

5 Results

RQ1: Is there a tradeoff between ensuring item-fairness in rankings and maintaining high ranking quality when utility labels are used for evaluation?

By gathering all four repeated runs of the experiments with different α values, we can plot the trend of ranking quality ($\overline{EE-R}$) against item fairness ($\overline{EE-D}$), as shown in Figure 4a.

As shown in previous studies [11, 56], there is a well-known tradeoff between fairness and ranking quality for human users. Similarly, we observe a general tradeoff for machine users. However, unlike past findings, this tradeoff is not always strict. For instance, in Figure 4a, both SPLADE and Contriever maintain consistently high ranking quality while being considerably fairer, and for BM25, ranking quality even improves as fairness increases, up to a certain point.

⁵<https://huggingface.co/FacebookAI/roberta-large-mnli>

	slope↓	AUC↑		slope↓	AUC↑
BM25	0.1351	0.3603	BM25	0.0895	0.3921
SPLADE	0.1971	0.4166	SPLADE	0.1453	0.4205
Contriever	0.1741	0.3864	Contriever	0.1624	0.4287
(a) LaMP-1			(b) LaMP-4		

Table 1: Comparison of fairness-ranking quality tradeoff (slope) and ranking quality (AUC) metrics for LaMP-1 and LaMP-4 datasets.

At the rightmost side of the lines, where $\overline{EE-D} = 1$ (representing the performance of deterministic rankers), we observe that these rankers do not always deliver the highest ranking quality. This suggests that commonly used deterministic rankers in RAG systems may be suboptimal, and that ranking quality can be improved while ensuring item fairness. This becomes even clearer when examining the impact of fair ranking on the downstream performance of a RAG system.

The leftmost side of the lines, where $\overline{EE-D} = 0$, represents the performance of a uniformly random ranking policy. At this point, the measured ranking quality should approximate the proportion of positively labeled items in the corpus, which is 31% for LaMP task 4. This is notably higher than in non-RAG (human-user) settings, where the percentage of relevant documents is typically much smaller, resulting in a $\overline{EE-R}$ value near 0 [11].

To quantify the tradeoff, we fit a linear line to the experiment results, where a steeper slope reflects a stronger tradeoff between fairness and ranking quality (slope). We also quantify the performance of fair rankers, by calculating the area under the disparity-ranking quality curve (AUC; Figure 4a). As can be seen in Table 1, we observe that retriever that yields higher ranking quality has higher tradeoff in terms of retrieval fairness.

RQ2: Is there a tradeoff between ensuring item-fairness in ranking and maintaining high generation quality of a RAG model?

Before examining the relationship between fairness and RAG utility, Figure 4b shows an auxiliary result confirming a strong correlation between utility-based ranking quality and the effectiveness of RAG models. This is unsurprising, as item-worthiness judgments were based on the utility-gain provided by the generator. However, this correlation suggests that the tradeoff observed in the disparity-ranking quality curve (Figure 4a) is likely to manifest similarly due to this strong relationship.

In fact, as observed from the disparity-utility curve (Figure 4c), we see a global trend of a non-strict tradeoff (i.e., RAG models maintain high generation quality while being considerably fair, and often even achieve higher quality).

However, a closer look at the local trend offers a significant insight: *RAG systems with fair ranking can often achieve higher system-effectiveness compared to models with deterministic rankers.* In Table 2, we divided the fairness levels into five disparity intervals based on the normalized $\overline{EE-D}$. As shown in the table, improving fairness to the level of $\overline{EE-D} \in [0.8, 1.0]$, and even $\overline{EE-D} \in [0.6, 0.8]$, can often enhance the expected utility of many RAG models across LaMP tasks. For example, having $\overline{EE-D}$ in the range of $[0.8, 1.0]$ outperforms the baseline for all models in LaMP-4 and for eight out of twelve models in LaMP-1.

Significance testing results in Table 2 provide additional nuance. The range $\overline{EE-D} \in [0.0, 0.2)$ shows a statistically significant drop in




Figure 4: Relationships between retrieval fairness, ranking quality, and generation quality. We empirically show that there is a positive relationship between ranking quality and utility, but trade-offs between retrieval fairness and ranking quality, as well as retrieval fairness and utility in RAG models. The curves are plotted by interpolating the runs for each disparity interval and averaging across generators for 4a and across retrievers for 4b and 4c. RetDisparity refers to the disparity among retrieved items.

Model (baseline utility)	Disparity Intervals				
	[0.0, 0.2)	[0.2, 0.4)	[0.4, 0.6)	[0.6, 0.8)	[0.8, 1.0)
LaMP-1					
BM25+FlanT5Small (0.308)	-0.12	-0.13	-0.18	-0.02	-0.15
BM25+FlanT5Base (0.670)	-0.20*	-0.04	-0.08	-0.05	-0.02
BM25+FlanT5XXL (0.531)	-0.07*	+0.03	+0.02	+0.06	+0.11*
BM25+FlanUL2 (0.557)	-0.1*	-0.02	-0.02	+0.01	+0.03
SPLADE+FlanT5Small (0.241)	-0.03	-0.22	+0.19	-0.04	+0.14
SPLADE+FlanT5Base (0.646)	-0.15*	+0.06	+0.08	0.00	+0.03
SPLADE+FlanT5XXL (0.671)	-0.18*	-0.16	+0.05	+0.02	+0.01
SPLADE+FlanUL2 (0.632)	-0.15*	-0.08	+0.02	+0.02	+0.01
Contriever+FlanT5Small (0.286)	-0.08	-0.29	-0.06	+0.03	-0.14
Contriever+FlanT5Base (0.637)	-0.16*	+0.05	-0.06	+0.03	0.00
Contriever+FlanT5XXL (0.651)	-0.19*	-0.04	-0.11	+0.03	0.00
Contriever+FlanUL2 (0.620)	-0.16*	-0.09	-0.05	+0.02	-0.03
Model (baseline utility)	Disparity Intervals				
	[0.0, 0.2)	[0.2, 0.4)	[0.4, 0.6)	[0.6, 0.8)	[0.8, 1.0)
LaMP-4					
BM25+FlanT5Small (0.217)	-0.06*	0.00	+0.02	+0.01	0.00
BM25+FlanT5Base (0.223)	-0.06*	0.00	+0.03	+0.01	+0.02
BM25+FlanT5XXL (0.322)	-0.05*	+0.11*	+0.03	+0.03	+0.05*
BM25+FlanUL2 (0.381)	-0.07*	+0.06*	+0.07*	+0.01	+0.04
SPLADE+FlanT5Small (0.235)	-0.07*	-0.01	+0.02	+0.03*	+0.02
SPLADE+FlanT5Base (0.268)	-0.10*	-0.03	+0.02	0.00	+0.02
SPLADE+FlanT5XXL (0.342)	-0.06*	+0.09*	+0.05*	+0.03	+0.04*
SPLADE+FlanUL2 (0.429)	-0.12*	+0.04	+0.02	+0.01	+0.01
Contriever+FlanT5Small (0.254)	-0.09*	-0.02*	0.00	+0.01	0.00
Contriever+FlanT5Base (0.268)	-0.10*	-0.02	+0.01	0.00	+0.01
Contriever+FlanT5XXL (0.367)	-0.09*	+0.06*	+0.01	+0.01	+0.03
Contriever+FlanUL2 (0.449)	-0.15*	+0.01	-0.01	0.00	+0.02

Table 2: Each value in the table is the difference between the utility of a baseline (deterministic) RAG model and the average utility of a fairer RAG model at a specific retrieval disparity interval. Nonnegative differences are highlighted. All LaMP tasks shows similar trend. A superscript * denotes a statistically significant difference between the treatment performance and the baseline performance based on an unpaired Student’s t-test ($p < 0.05$).

utility compared to the baseline, indicating that pushing disparity to extremely low levels can indeed reduce performance. Beyond this small-disparity range ($\overline{EE-D} \in [0.2, 1.0)$), the differences in utility scores from the baseline either remain statistically insignificant or—when significant—reflect improved utility. This suggests that models can maintain a level of retrieval fairness close to the baseline’s utility performance without incurring a notable cost, and may even achieve higher effectiveness in many cases.

RQ3: What is the impact of item-fairness in ranking to the fairness of the attributed sources used in the final response?

To address this, we analyze three primary relationships: (1) how retrieval fairness ($\overline{EE-D}$) relates to the expected attribution rate (EAR), (2) how ranking quality ($\overline{EE-R}$) interacts with EAR, and (3) how retrieval fairness translates into consumption fairness ($\overline{EAE-D}$) in the final output.

Figures 5a and 5b together shed light on how many items are actually attributed by the generator, which is a vital step in understanding how exposure might become more or less equitable once the generator has consumed the retrieved items. Figure 5a shows that when retrieval becomes more fair (lower $\overline{EE-D}$), the generator tends to attribute fewer items overall, suggesting a tradeoff between distributing items equitably and actually getting them used in the final output.

Interestingly, the general shape of $\overline{EE-D}$ versus EAR in Figure 5a aligns closely with the shape of $\overline{EE-D}$ versus $\overline{EE-R}$ in Figure 4a. By looking at Figure 5b, it becomes clear that the generator is more likely to use multiple items if they are useful in solving the task, reflecting the idea that high-quality rankings encourage higher EAR.




Figure 5: Major relationships between retrieval fairness (RetDisparity), ranking quality, expected attribution rate, and consumption fairness (AttDisparity). We empirically show that there is a tradeoff between retrieval fairness and expected attribution rate as there is a positive relationship between ranking quality and attribution rate. Most importantly, we show that retrieval fairness does not necessarily directly propagate to the consumption fairness due to varying attribution rate of a generator. The same interpolation methods are used as Figure 4 and runs were averaged across generators. RetDisparity refers to the disparity among retrieved items, while AttDisparity refers to the disparity among attributed sources.

Figure 5c directly addresses the question of whether fair retrieval leads to fair consumption by comparing $\overline{EE-D}$ (retrieval fairness) to $\overline{EAE-D}$ (consumption fairness).⁶ The positive trend suggests that a more equitable ranking generally produces a more equitable distribution of attributed items in the final generation. However, the diagonal line ($y=x$) reveals subtle differences between fairness at the ranking stage and fairness at the consumption stage. Data points falling below this line imply that consumption is even fairer than the initial ranking might predict, while points above it suggest that the generator’s selective usage of items has introduced new disparity. This pattern shifts depending on how relevant and useful the retrieved items are, mirroring the earlier observation (Figure 5b) that the generator favors content that supports solving the task. In cases where many items have only marginal utility, the generator’s emphasis on a few high-utility sources can inadvertently raise disparity, placing the plotted values above the $y=x$ line.

6 Discussion

6.1 Higher System Utility with Fair Rankings

Although there is a general trend of a fairness-utility tradeoff, we observe that certain levels of fairness can actually improve the utility of a baseline RAG model. Recent line of research have uncovered relevant findings: 1) generators are not robust to changes in the position of useful information [34]; 2) items with high retrieval scores often include distracting content that can reduce the system-effectiveness [10, 47]; and 3) introducing some random documents can significantly boost the utility of RAG [10].

Building on these existing results, we find that perturbing the initial ranking through stochastic sampling often can impact the performance of certain inference decisions and lead to changes

in the system’s expected end-performance. In our experiments, we observe that the expected utility generally increases within the fairness interval of $[0.8, 1.0]$. This suggests that a fixed ranking from a deterministic ranker may be suboptimal for the generator, and that perturbing the ranking, along with the repositioning of items, not only improves expected end-performance but also enhances the fairness of the rankings.

Moreover, in fairness intervals where the system’s expected utility improves, it is possible that either fewer distracting items were included in the ranking passed to the generator or useful, previously overlooked items (which may have been considered random) were introduced due to the ranking perturbation. However, while higher utility paired with increased item-fairness (even within fairness intervals as low as $[0.4, 0.6]$) may seem advantageous, practitioners should exercise caution. This could result in compensating providers of items irrelevant to user requests, particularly in scenarios where content providers are rewarded for contributing to inference outcomes.

6.2 Attribution Rate and Consumption Fairness

RAG practitioners should consider both the expected utility (EU) of the generated text and the expected attributed exposure disparity ($\overline{EAE-D}$) to ensure a balance between utility and fairness in the system’s output. Achieving this balance requires careful adjustment of retrieval fairness to find an interval where the generated content remains both useful and fair in its exposure of sources. One important metric to monitor in this process is the attribution rate, which reflects the proportion of retrieved items that are actually cited or used in the final output.

Monitoring expected attribution rate (EAR) provides valuable insights for understanding how fairness manifests in the final output, offering an essential complement to $\overline{EAE-D}$. For example, even if a ranking appears fair at the retrieval stage, a low attribution rate

⁶Normalized EE-D and EAE-D scores are directly comparable, as they share the same dimensions in their exposure vectors and have the same l_1 norm.

indicates that the generator ultimately utilizes only a small subset of the retrieved items, which can amplify disparities among the attributed sources. In such cases, fairness concerns at the generation stage may outweigh those at the retrieval stage. Conversely, when the generator attributes a larger portion of the retrieved set, disparities among the attributed sources may become less pronounced, as the system distributes attention more evenly across a broader set of items.

By comparing expected attribution rate with consumption fairness, practitioners can better disentangle the root causes of imbalances in the final output. Specifically, they can determine whether disparities stem from the retrieval stage, such as an inherently skewed ranking, or from the generator’s behavior in selectively focusing on only a few documents. This comparison can help diagnose fairness issues more effectively, enabling practitioners to refine both retrieval and generation components of a RAG system to achieve more equitable outcomes in the generated content.

6.3 Measuring Consumption Fairness

Ensuring consumption fairness while maintaining high utility of the generated text is a fundamental objective in building a fair RAG system. To achieve this, we measure the exposure of the final attributed items, as these are the items that human users will actually encounter in the system’s output. This approach aligns with the real-world operation of RAG systems, particularly in applications like conversational QA or domain-specific assistants for general information or shopping, where the generated text is accompanied by explicitly cited documents or attributed items. In such cases, measuring the fair exposure of the attributed items is critical, as consumption fairness can have a more direct impact on users compared to retrieval fairness, given that human users only consume what is explicitly surfaced by the system.

A key factor in accurately evaluating exposure-based retrieval fairness is the development of sophisticated machine-user browsing models. The work of Liu et al. [34] provides valuable insights for designing such advanced browsing models. In our experiments, we employed encoder-decoder models with a short retrieval context to get close to the assumption that equal attention is given to all top- k retrieved items [34]. However, as the use of long-context becomes prevalent, accurately measuring exposure-based retrieval fairness becomes more challenging due to the generator’s tendency to allocate unequal attention to different items.

Despite these challenges, it is important to highlight that measuring consumption fairness is not dependent on machine-user browsing models. Unlike retrieval fairness, consumption fairness can be measured effectively in long-context models, regardless of how attentions are distributed across items by the generators. Therefore, focusing on consumption fairness ensures that the evaluation remains robust, even in scenarios where exposure patterns vary significantly across different models.

6.4 Measurement of String Utility

In line with the recent call for evaluating various valid output strings [63], we recognize the need for a similar approach to better measure system utility across different rankings given. Recall that our experiments were designed to provide the generator with

different rankings for the same query, leading to varied outputs. This approach is motivated by the idea that items not appearing in the top positions of deterministic rankings may still hold value and should be fairly considered by the system. In this context, the diverse outputs generated from different rankings may still be valid. However, we currently rely on a single target output string for comparison with predictions. Future work could focus on calculating the utility of diffuse predictions, enabling a more nuanced evaluation.

6.5 Limitations

We acknowledge that the evaluation cost of fair RAG systems can be high due to repeated sampling and inference steps. However, in production, only a single ranking is sampled, minimizing the impact on system latency. Also, a limitation in our utility labeling is that it considers single items, while multiple items may yield contrasting utility gains. Despite this, the strong correlation between ranking quality and system effectiveness suggests this approach reasonably approximates item-worthiness for evaluating the impact of fair ranking on RAG systems.

Another limitation lies in the ability to capture the exposure of items beyond those retrieved by the system. Since it is difficult to verify attribution across the entire corpus, our method focuses on measuring the exposure fairness of the attributed sources (e.g., those explicitly cited in the generated output). While this approach does not account for the broader set of potential exposures, it is useful for common RAG applications where the attributed sources are directly displayed to users. In such cases, ensuring fairness among these sources is critical, as they are likely to be the primary content users engage with.

7 Conclusion

This study highlights the impact of fair rankings not only on the ranking and generation quality of RAG systems but also on the equitable attribution of sources in the final output. Through extensive analysis, we demonstrate that fairer RAG models can maintain—and in some cases even surpass—the generation quality of traditional approaches, challenging the assumption of an inherent tradeoff between fairness and effectiveness. Our findings emphasize the importance of fair attribution, showing how improvements at the retrieval stage translate into more equitable exposure of the sources that appear in the generated text. By addressing disparities in both retrieval and attribution, we provide valuable insights for developing responsible and equitable RAG systems.

In future work, we hope to extend this framework to consider graded or missing judgments and exploring the different notions of fairness in RAG systems, ultimately advancing the field of trustworthy RAG systems research.⁷

Acknowledgments

This work was supported by NSF grant 2402874. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect those of the sponsors.

⁷We release our code and dataset at <https://github.com/kimdanny/Fair-RAG>.

References

- [1] Akari Asai, Zeqiu Wu, Yizhong Wang, Avirup Sil, and Hannaneh Hajishirzi. 2024. Self-RAG: Learning to Retrieve, Generate, and Critique through Self-Reflection. In *The Twelfth International Conference on Learning Representations*.
- [2] Dara Bahri, Yi Tay, Che Zheng, Donald Metzler, and Andrew Tomkins. 2020. Choppy: Cut Transformer for Ranked List Truncation. *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval* (2020).
- [3] Aparna Balagopal, Abigail Z. Jacobs, and Asia J. Biega. 2023. The Role of Relevance in Fair Ranking. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '23)*. Association for Computing Machinery, 2650–2660.
- [4] K. Balan, S. Agarwal, S. Jenni, A. Parsons, A. Gilbert, and J. Collomosse. 2023. EKILA: Synthetic Media Provenance and Attribution for Generative Art. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*. IEEE Computer Society, 913–922.
- [5] Asia J Biega, Krishna P Gummadi, and Gerhard Weikum. 2018. Equity of attention: Amortizing individual fairness in rankings. In *The 41st international acm sigir conference on research & development in information retrieval*. 405–414.
- [6] Bernd Bohnet, Vinh Q Tran, Pat Verga, Roei Aharoni, Daniel Andor, Livio Baldini Soares, Massimiliano Ciaramita, Jacob Eisenstein, Kuzman Ganchev, Jonathan Herzig, et al. 2022. Attributed question answering: Evaluation and modeling for attributed large language models. *arXiv preprint arXiv:2212.08037* (2022).
- [7] Sebastian Bruch, Shuguang Han, Michael Bendersky, and Marc Najork. 2020. A Stochastic Treatment of Learning to Rank Scoring Functions. In *Proceedings of the 13th International Conference on Web Search and Data Mining (WSDM '20)*. Association for Computing Machinery, 61–69.
- [8] Simon Caton and Christian Haas. 2020. Fairness in machine learning: A survey. *Comput. Surveys* (2020).
- [9] Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, Albert Webson, Shixiang Shane Gu, Zhuyun Dai, Mirac Suzgun, Xinyun Chen, Aakanksha Chowdhery, Alex Castro-Ros, Marie Pellat, Kevin Robinson, Dasha Valter, Sharan Narang, Gaurav Mishra, Adams Yu, Vincent Zhao, Yanping Huang, Andrew Dai, Hongkun Yu, Slav Petrov, Ed H. Chi, Jeff Dean, Jacob Devlin, Adam Roberts, Denny Zhou, Quoc V. Le, and Jason Wei. 2024. Scaling Instruction-Finetuned Language Models. *Journal of Machine Learning Research* 25, 70 (2024), 1–53. <http://jmlr.org/papers/v25/23-0870.html>
- [10] Florin Cuconasu, Giovanni Trappolini, Federico Siciliano, Simone Filice, Cesare Campagnano, Yoelle Maarek, Nicola Tonello, and Fabrizio Silvestri. 2024. The Power of Noise: Redefining Retrieval for RAG Systems. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '24)*. Association for Computing Machinery, 719–729.
- [11] Fernando Diaz, Bhaskar Mitra, Michael D. Ekstrand, Asia J. Biega, and Ben Carterette. 2020. Evaluating Stochastic Rankings with Expected Exposure. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management (CIKM '20)*. Association for Computing Machinery, 275–284.
- [12] Michael D. Ekstrand, Lex Beattie, Maria Soledad Pera, and Henriette Cramer. 2024. Not Just Algorithms: Strategically Addressing Consumer Impacts in Information Retrieval. In *Advances in Information Retrieval*. Springer Nature Switzerland, 314–335.
- [13] Michael D Ekstrand, Anubrata Das, Robin Burke, and Fernando Diaz. 2022. Fairness in information access systems. *Foundations and Trends® in Information Retrieval* 16, 1-2 (2022), 1–177.
- [14] Shahul Es, Jithin James, Luis Espinosa Anke, and Steven Schockaert. 2024. RAGAs: Automated Evaluation of Retrieval Augmented Generation. In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics: System Demonstrations*, Nikolaos Aletras and Orphee De Clercq (Eds.). Association for Computational Linguistics, 150–158.
- [15] Thibault Formal, Carlos Lassance, Benjamin Piwowarski, and Stéphane Clinchant. 2022. From Distillation to Hard Negative Sampling: Making Sparse Neural IR Models More Effective. 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, 2353–2359. <https://doi.org/10.1145/3477495.3531857>
- [16] Luyu Gao, Zhuyun Dai, Panupong Pasupat, Anthony Chen, Arun Tejasvi Chaganty, Yicheng Fan, Vincent Zhao, Ni Lao, Hongrae Lee, Da-Cheng Juan, and Kelvin Guu. 2023. RARR: Researching and Revising What Language Models Say, Using Language Models. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, 16477–16508.
- [17] Tianyu Gao, Howard Yen, Jiatong Yu, and Danqi Chen. 2023. Enabling Large Language Models to Generate Text with Citations. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 6465–6488. <https://doi.org/10.18653/v1/2023.emnlp-main.398>
- [18] John Guiver and Edward Snelson. 2009. Bayesian inference for Plackett-Luce ranking models. In *Proceedings of the 26th Annual International Conference on Machine Learning*. ACM, 377–384.
- [19] Emil Julius Gumbel. 1954. *Statistical theory of extreme values and some practical applications: a series of lectures*. Vol. 33. US Government Printing Office.
- [20] Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Ming-Wei Chang. 2020. REALM: Retrieval-Augmented Language Model Pre-Training. In *Proceedings of the 37th International Conference on Machine Learning (ICML '20)*. JMLR.org, Article 368, 10 pages.
- [21] Junqing He, Kunhao Pan, Xiaoqun Dong, Zhuoyang Song, LiuYiBo LiuYiBo, Qiangsuun Qiangsuun, Yuxin Liang, Hao Wang, Enming Zhang, and Jiaxing Zhang. 2024. Never Lost in the Middle: Mastering Long-Context Question Answering with Position-Agnostic Decompositional Training. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Lun-Wei Ku, Andre Martins, and Vivek Srikumar (Eds.). Association for Computational Linguistics, Bangkok, Thailand, 13628–13642. <https://doi.org/10.18653/v1/2024.acl-long.736>
- [22] Peter Henderson, Xuechen Li, Dan Jurafsky, Tatsunori Hashimoto, Mark A. Lemley, and Percy Liang. 2023. Foundation Models and Fair Use. *Journal of Machine Learning Research* 24, 400 (2023), 1–79.
- [23] Sebastian Hofstätter, Jiecao Chen, Karthik Raman, and Hamed Zamani. 2023. Fidelity: Efficient and effective retrieval-augmented text generation. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1437–1447.
- [24] Or Honovich, Roei Aharoni, Jonathan Herzig, Hagai Taitelbaum, Doron Kukliansky, Vered Cohen, Thomas Scialom, Idan Szpektor, Avinatan Hassidim, and Yossi Matias. 2022. TRUE: Re-evaluating Factual Consistency Evaluation. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Marine Carpuat, Marie-Catherine de Marneffe, and Ivan Vladimir Meza Ruiz (Eds.). Association for Computational Linguistics, Seattle, United States, 3905–3920. <https://doi.org/10.18653/v1/2022.naacl-main.287>
- [25] Cheng-Yu Hsieh, Yung-Sung Chuang, Chun-Liang Li, Zifeng Wang, Long Le, Abhishek Kumar, James Glass, Alexander Ratner, Chen-Yu Lee, Ranjay Krishna, and Tomas Pfister. 2024. Found in the middle: Calibrating Positional Attention Bias Improves Long Context Utilization. In *Findings of the Association for Computational Linguistics: ACL 2024*, Lun-Wei Ku, Andre Martins, and Vivek Srikumar (Eds.). Association for Computational Linguistics, Bangkok, Thailand, 14982–14995. <https://doi.org/10.18653/v1/2024.findings-acl.890>
- [26] Gautier Izacard, Mathilde Caron, Lucas Hosseini, Sebastian Riedel, Piotr Bojanowski, Armand Joulin, and Edouard Grave. 2022. Unsupervised Dense Information Retrieval with Contrastive Learning. *Transactions on Machine Learning Research* (2022).
- [27] Gautier Izacard, Patrick Lewis, Maria Lomeli, Lucas Hosseini, Fabio Petroni, Timo Schick, Jane Dwivedi-Yu, Armand Joulin, Sebastian Riedel, and Edouard Grave. 2023. Atlas: Few-shot learning with retrieval augmented language models. *Journal of Machine Learning Research* 24, 251 (2023), 1–43.
- [28] Thomas Jaenicke, Graham McDonald, and Iadh Ounis. 2024. Fairness-Aware Exposure Allocation via Adaptive Reranking. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '24)*. Association for Computing Machinery, 1504–1513.
- [29] Meng Jiang, Keqin Bao, Jizhi Zhang, Wenjie Wang, Zhengyi Yang, Fuli Feng, and Xiangnan He. 2024. Item-side Fairness of Large Language Model-based Recommendation System. In *Proceedings of the ACM on Web Conference 2024 (WWW '24)*. Association for Computing Machinery, 4717–4726.
- [30] Diane Kelly, Jaime Arguello, and Robert Capra. 2013. NSF workshop on task-based information search systems. *SIGIR Forum* 47, 2 (2013), 116–127.
- [31] Urvashi Khandelwal, Omer Levy, Dan Jurafsky, Luke Zettlemoyer, and Mike Lewis. 2020. Generalization through Memorization: Nearest Neighbor Language Models. In *International Conference on Learning Representations*.
- [32] To Eun Kim, Alireza Salemi, Andrew Drozdov, Fernando Diaz, and Hamed Zamani. 2024. Retrieval-Enhanced Machine Learning: Synthesis and Opportunities. *arXiv preprint arXiv:2407.12982* (2024).
- [33] Patrick S. H. Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6–12, 2020, virtual*, Hugo Larochelle, Marc'Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin (Eds.).
- [34] Nelson F Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2024. Lost in the middle: How language models use long contexts. *Transactions of the Association for Computational Linguistics* 12 (2024), 157–173.
- [35] Yinhan Liu. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692* 364 (2019).
- [36] Lingjuan Lyu, C Chen, and J Fu. 2023. A Pathway Towards Responsible AI Generated Content. In *IJCAI*. 7033–7038.

- [37] Chris J. Maddison, Andriy Mnih, and Yee Whye Teh. 2017. The Concrete Distribution: A Continuous Relaxation of Discrete Random Variables. In *International Conference on Learning Representations*.
- [38] Rishabh Mehrotra, Ashton Anderson, Fernando Diaz, Amit Sharma, Hanna Wallach, and Emine Yilmaz. 2017. Auditing Search Engines for Differential Satisfaction Across Demographics. In *Proceedings of the 26th International Conference on World Wide Web Companion (WWW '17 Companion)*. International World Wide Web Conferences Steering Committee, 626–633.
- [39] Alistair Moffat and Justin Zobel. 2008. Rank-biased precision for measurement of retrieval effectiveness. *ACM Transactions on Information Systems (TOIS)* 27, 1 (2008), 1–27.
- [40] Abhiman Neelakanteswara, Shreyas Chaudhari, and Hamed Zamani. 2024. RAGs to Style: Personalizing LLMs with Style Embeddings. In *Proceedings of the 1st Workshop on Personalization of Generative AI Systems (PERSONALIZE 2024)*. 119–123.
- [41] Harrie Oosterhuis. 2021. Computationally efficient optimization of plackett-luce ranking models for relevance and fairness. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1023–1032.
- [42] Harrie Oosterhuis. 2022. Learning-to-rank at the speed of sampling: Plackett-luce gradient estimation with minimal computational complexity. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 2266–2271.
- [43] Robin L. Plackett. 1975. The analysis of permutations. *Journal of the Royal Statistical Society Series C: Applied Statistics* 24, 2 (1975), 193–202.
- [44] Amifa Raj and Michael D Ekstrand. 2020. Comparing fair ranking metrics. *arXiv preprint arXiv:2009.01311* (2020).
- [45] Hannah Rashkin, Vitaly Nikolaev, Matthew Lamm, Lora Aroyo, Michael Collins, Dipanjan Das, Slav Petrov, Gaurav Singh Tomar, Iulia Turc, and David Reitter. 2023. Measuring Attribution in Natural Language Generation Models. *Computational Linguistics* 49, 4 (Dec. 2023), 777–840. https://doi.org/10.1162/coli_a_00486
- [46] Stephen Robertson, S. Walker, S. Jones, M. M. Hancock-Beaulieu, and M. Gatford. 1995. Okapi at TREC-3. In *Proceedings of the Third Text REtrieval Conference (TREC-3)*. Gaithersburg, MD: NIST, 109–126.
- [47] Dongyu Ru, Lin Qiu, Xiangkun Hu, Tianhang Zhang, Peng Shi, Shuaichen Chang, Cheng Jiayang, Cunxiang Wang, Shichao Sun, Huan Yu Li, Zizhao Zhang, Binjie Wang, Jiarong Jiang, Tong He, Zhiguo Wang, Pengfei Liu, Yue Zhang, and Zheng Zhang. 2024. RAGChecker: A Fine-grained Framework for Diagnosing Retrieval-Augmented Generation. In *The Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track*. <https://openreview.net/forum?id=J9oefdGUuM>
- [48] Jon Saad-Falcon, Omar Khattab, Christopher Potts, and Matei Zaharia. 2024. ARES: An Automated Evaluation Framework for Retrieval-Augmented Generation Systems. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, Kevin Duh, Helena Gomez, and Steven Bethard (Eds.). Association for Computational Linguistics, 338–354.
- [49] Alireza Salemi, Surya Kallumadi, and Hamed Zamani. 2024. Optimization Methods for Personalizing Large Language Models through Retrieval Augmentation. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '24)*. Association for Computing Machinery, 752–762. <https://doi.org/10.1145/3626772.3657783>
- [50] Alireza Salemi, Sheshera Mysore, Michael Bendersky, and Hamed Zamani. 2024. LaMP: When Large Language Models Meet Personalization. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, 7370–7392.
- [51] Alireza Salemi and Hamed Zamani. 2024. Evaluating retrieval quality in retrieval-augmented generation. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 2395–2400.
- [52] Piotr Sapiezynski, Wesley Zeng, Ronald E Robertson, Alan Mislove, and Christo Wilson. 2019. Quantifying the Impact of User Attention on Fair Group Representation in Ranked Lists. In *Companion Proceedings of The 2019 World Wide Web Conference (WWW '19)*. Association for Computing Machinery, 553–562.
- [53] Tefko Saracevic. 2016. *The Notion of Relevance in Information Science: Everybody knows what relevance is. But, what is it really?* Morgan & Claypool Publishers.
- [54] Robik Shrestha, Yang Zou, Qiuyu Chen, Zhiheng Li, Yusheng Xie, and Siqi Deng. 2024. FairRAG: Fair Human Generation via Fair Retrieval Augmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 11996–12005.
- [55] Ashudeep Singh and Thorsten Joachims. 2018. Fairness of exposure in rankings. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*. 2219–2228.
- [56] Ashudeep Singh and Thorsten Joachims. 2019. Policy learning for fairness in ranking. In *Proceedings of the 33rd International Conference on Neural Information Processing Systems*. Number 487. Curran Associates Inc., 5426–5436.
- [57] Yi Tay, Mostafa Dehghani, Vinh Q. Tran, Xavier Garcia, Jason Wei, Xuezhi Wang, Hyung Won Chung, Dara Bahri, Tal Schuster, Steven Zheng, Denny Zhou, Neil Houlsby, and Donald Metzler. 2023. UL2: Unifying Language Learning Paradigms. In *The Eleventh International Conference on Learning Representations*. <https://openreview.net/forum?id=6ruVLB727MC>
- [58] Ke Yang and Julia Stoyanovich. 2017. Measuring fairness in ranked outputs. In *Proceedings of the 29th international conference on scientific and statistical database management*. 1–6.
- [59] Hamed Zamani and Michael Bendersky. 2024. Stochastic RAG: End-to-End Retrieval-Augmented Generation through Expected Utility Maximization. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '24)*. Association for Computing Machinery, 2641–2646. <https://doi.org/10.1145/3626772.3657923>
- [60] Hamed Zamani, Fernando Diaz, Mostafa Dehghani, Donald Metzler, and Michael Bendersky. 2022. Retrieval-Enhanced Machine Learning. In *Proceedings of the 45th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*.
- [61] Meike Zehlike, Francesco Bonchi, Carlos Castillo, Sara Hajian, Mohamed Megahed, and Ricardo Baeza-Yates. 2017. Fa* ir: A fair top-k ranking algorithm. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*. 1569–1578.
- [62] Hengran Zhang, Ruqing Zhang, Jiafeng Guo, Maarten de Rijke, Yixing Fan, and Xueqi Cheng. 2024. Are Large Language Models Good at Utility Judgments?. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '24)*. Association for Computing Machinery, 1941–1951.
- [63] Yiming Zhang, Avi Schwarzschild, Nicholas Carlini, J Zico Kolter, and Daphne Ippolito. 2024. Forcing Diffuse Distributions out of Language Models. In *First Conference on Language Modeling*. <https://openreview.net/forum?id=9JY1QLVFPZ>