LIFE: Lifetime-Aware Fair Exposure for Time-Limited Items in Recommendation Systems

Maria Ianne B. Castro [1], Maria de Lourdes M. Silva[1], Iago C. Chaves[1], Javam C. Machado[1]

[1]Laboratório de Sistemas e Bancos de Dados (LSBD)
Departamento de Computação
Universidade Federal do Ceará, Brazil

{ianne.braz, malu.maia, iago.chaves, javam.machado}@lsbd.ufc.br

Abstract. *Recommendation systems play an important role by predicting users' preferences while allocating visibility and economic opportunity to item providers. This becomes especially challenging when items have limited lifetimes (e.g., news articles, political campaigns, perishable goods, and seasonal products), because time-sensitive attention can be disproportionately captured by a small subset of items due to position bias and popularity effects. Building on the notion of fair exposure, we propose LIFE , a multi-objective post-processing approach that adjusts users' top-k lists to balance user relevance, provider fairness, and item urgency, proportionally estimating quality over a time window while prioritizing items closer to the end of their lifetimes.*

1. Introduction

As the amount of online information grows exponentially, recommendation systems become central components in filtering and organizing digital content. Whether recommending products, media, or news, these systems primarily aim to select and rank items according to estimated preferences of the users, facing the challenge of handling this massive volume in order to accurately deliver what users want and need to receive. By organizing elements into rankings, recommendation systems determine which content receives greater visibility, directly influencing attention, interaction, and information consumption patterns. Beyond satisfying user preferences, these systems serve as engines, allocating visibility and opportunity, consequently playing an important role in providing access to the small items providers, who might suffer from underexposure [Morik et al. 2020, Wang et al. 2022, de Lourdes M. Silva et al. 2025]. In the context of items with limited lifetimes, such as news, perishable goods, political campaigns, and seasonal products, the challenge is particularly relevant due to the high volume and rapid turnover of available content.

Given their significant importance in the current scenario, recommendation systems impact not only users but also the items they recommend. Empirical studies, such as those by [Joachims et al. 2017, Collins et al. 2018], have quantified the impact of position bias, demonstrating that higher-ranked items achieve significantly higher click-through rates, regardless of their true relevance, emphasizing the effect of position bias. This dynamism of initial exposure is decisive, since the lack of visibility in early moments prevents the collection of information about the item, leading to its premature neglect and disproportionately benefiting already popular items. The notion of exposure fairness seeks to ensure that the attention allocated to items is proportional to defined criteria, such as relevance or representativeness, thereby avoiding systematic distortions in the distribution of visibility over time.

This exposure problem becomes considerably more complex in online scenarios, where ranking is continuously updated, and items often have a short lifetime of peak relevance, and the absence of exposure during their tight relevance window implies a loss of utility. For instance, recommending a local accident that occurred a week ago in a news article does not seem ideal, but recommending this article in the first hours after the event would be far more effective. In this context, the window of opportunity to ensure fairness of exposure is tight; any delayed correction loses its impact, making the real-time application of fairness of exposure a challenge worth exploring.

This work proposes a framework for analyzing and controlling exposure in recommendation systems with time-sensitive items. The main objective of LIFE is to verify and adjust the exposure distribution, ensuring that the exposure of item groups present in the top-k is proportional to their representativeness or relevance within the overall set. Using post-processing mechanisms and multi-objective reranking, the approach seeks to preserve user relevance. It also ensures that exposure of item groups is proportional to their representativeness in the current window, prioritizing items near the end of their lifetimes.

The work is organized as follows. Section 2 reviews related work, contextualizing the research and providing the theoretical background necessary to position this study's contribution. Section 3 establishes the necessary background, discusses fairness concepts in recommendation systems, and defines the attributes and variables that underpin the framework. Section 4 details the proposed approach, describing its central objective, the mechanism's operation, and the developed algorithm. Section 5 reports the experiments and results, evaluating the effectiveness of the solution across different applications. Finally, Section 6 presents the conclusions and final considerations.

2. Related Work and Background

In this section, we provide the necessary background and related work to support the understanding of our proposed approach, exploring existing works on fairness in recommendation systems and limited lifetime items' scenarios.

2.1. Fairness in recommendation systems

Recommendation platforms involve two principal stakeholder groups: *users*, who consume the available items or content, and *providers*, who make these items available on the platform [Wang et al. 2023]. Although both groups participate in the same ecosystem, their objectives are not necessarily aligned. Users are primarily concerned with receiving accurate and relevant recommendations, whereas providers seek greater consumption of their items, which largely depends on the visibility they receive from the platform.

The study of fairness in recommendation systems has received increasing attention in recent years, giving rise to a variety of fairness notions and algorithmic strategies. In this subsection, we examine representative works that address fairness from the perspectives of both major stakeholders, namely *user-side fairness*, which focuses on fairness toward users, and *item-side fairness*, which concerns fairness toward items or providers.

User-side fairness This dimension of fairness concerns user satisfaction and utility with respect to the recommendations they receive. From a utilitarian perspective, the objective is to maximize social welfare, understood as the aggregate well-being of individuals, while assigning equal moral importance to each person so that no individual is granted unjustified priority or advantage over another [Rawls 2017]. In the context of recommendation systems, this perspective implies that the well-being of no user should be systematically prioritized over others. Along this line, Do et al. [Do et al. 2021] propose an approach that maximizes overall utility while giving particular attention to users who receive the least favorable recommendations. To this end, the authors employ concepts such as Lorenz efficiency to encourage rankings that are Pareto efficient and that improve the utility of the worst-off individuals. From a different perspective on user-side fairness, Li et al. [Li et al. 2021] argue that users belonging to different demographic groups should receive comparable recommendation performance, thereby reducing the risk that certain groups are systematically less satisfied with the recommendations they receive. Moreover, drawing on the notion of envy-freeness [Brams and Taylor 1995, Farhadi et al. 2021], user fairness may also be understood as a condition in which no user prefers another user's recommendation list to their own. However, because recommended items are indivisible and recommendation lists are inherently capacity-constrained, exact envy-freeness is generally unattainable in practice [Amanatidis et al. 2023], which motivates the adoption of relaxed fairness notions [Arnsperger 1994].

Item-side fairness On the other hand, providing fairness addresses disparities between items stemming from popularity bias [Guo et al. 2025], a phenomenon in which items monopolize the top of rankings. Fairness of exposure was formalized to mitigate the marginalization of those items, using the concept of probabilistic allocation of user visual attention in the fairness constraint [Singh and Joachims 2018]. With that concept, the notion of quality-weighted exposure holds that an item or provider's exposure should not necessarily be uniform but rather proportional to its relevance or quality [Heuss et al. 2022]. This approach differs from uniform exposure, as it does not distribute attention equally across all content. Consequently, it avoids disregarding inherent user preferences and the actual relevance distribution within the system's context. A recommendation resulting from quality-weighted exposure is considered fair if each provider receives an overall exposure proportional to the sum of the quality scores of the items they offer across the platform [Wu et al. 2021]. That comes from the fact that position bias inherently prevents any single ranking from achieving full proportionality; systems must treat and amortize fairness over a sequence of rankings [Biega et al. 2018].

2.2. Limited lifetime items

Multiple recommendation scenarios involve items with limited lifetime, whose utility and user interest decay over time. This is particularly evident in domains such as news, fresh content, and newly introduced items, where recommendations must account not only for relevance but also for temporal urgency and visibility [Zhang et al. 2017, Bae et al. 2023].

In these settings, the literature has argued that such items or providers may require a minimum level of exposure within a fixed time window to remain viable, motivating recommendation policies that explicitly incorporate exposure constraints over time [Ben-Porat and Torkan 2023, Bae et al. 2023].

Those lifetime constraints can be seen in news recommendation, for instance, articles rapidly lose relevance as events unfold and public attention shifts [Raza and Ding 2022], so delayed exposure often provides little impact. A similar effect arises in seasonal sales, where promotional items are only attractive during specific commercial periods, such as holidays or end-of-season campaigns [Macedo and Marinho 2014]. In the case of event tickets, the opportunity to benefit from exposure vanishes as the event date approaches or seats become unavailable. Pharmaceutical items may also exhibit lifetime sensitivity, particularly when recommendations involve drugs, vaccines, or medical supplies with expiration dates or time-critical demand. Likewise, perishable items in grocery or food platforms have short commercial viability because freshness degrades over time, making late recommendations less useful for both consumers and suppliers. Finally, job vacancies are also time-constrained items, since positions are eventually filled or closed, and their value diminishes as application deadlines pass.

Under this context, foundational fairness frameworks generally assume an unlimited time horizon for correcting exposure disparities, overlooking the fact that time-sensitive items lose value when their exposure is delayed. Recent advancements have attempted to address fairness in continuously evolving environments. For instance, [Wu et al. 2023] introduced dynamic strategies designed to balance exposure fairness and recommendation quality over sequential interactions. Similarly, [Guo et al. 2025] employed a reinforcement learning framework to enhance the exposure of newly introduced items in dynamic recommender systems. Despite addressing temporal dynamics and accumulated unfairness, these models do not explicitly account for the hard constraints imposed by item expiration. None of the existing studies successfully combine amortized exposure fairness [Biega et al. 2018] with the urgency of item expiration [Bae et al. 2023, Ben-Porat and Torkan 2023] in a dynamic, post-processing reranking environment. LIFE fills this gap by introducing a multi-objective mechanism that adjusts users' top-k lists to simultaneously balance user relevance, provider fairness, and item urgency.

3. Preliminaries

This section formalizes the notation and terminology used throughout this work.

Let n and m denote the number of users and items, respectively. The recommendation environment consists of a set of users $U = \{u_1, u_2, \ldots, u_n\}$ and a set of items $I = \{i_1, i_2, \ldots, i_m\}$. Additionally, a set of categories $C = \{c_1, c_2, \ldots, c_o\}$ is considered, where each item is associated with one category.

The system operates over a sequence of time windows $W = \{W_1, W_2, \ldots, W_P\}$, encapsulating the evolution of users' preferences and the continuous entry of new items. Each window $W_t \in W$ represents a discrete interval defined by a start time and a size, which determines its duration. It is also used to parameterize the estimated qualities of items in the W_t time window and the exposures.

For each item i, a lifetime value L_i is defined, representing its expected peak relevance period. Based on empirical notions of temporal decay and importance, it is assumed that items approaching the end of their lifespan require increased exposure to maximize their effective relevance. Consequently, urgency $\mathcal{U}_i(t)$ is defined as a measure of proximity to the end of an item‘s lifetime. The urgency at time t is calculated as the proportion of the consumed lifespan:

$$\mathcal{U}_i(t) = \begin{cases} \frac{t-t_i^{pub}}{L_i}, & \text{if } 0 \leq t - t_i^{pub} \leq L_i \\ 0, & \text{otherwise} \end{cases}$$

where t_i^{pub} denotes the publication time. This ensures that items reach users before losing their peak of relevance. This way, the system functions as a “last-chance“ mechanism, elevating the exposure priority for items about to expire.

The recommendation process is based on a score matrix $S \subseteq \mathbb{R}^{n\times m}$, where every matrix cell $\hat{s}_{ui} \in S$ represents the predicted relevance of the item i associated to the user u. The quantified quality of an item i within a window W_t is denoted as $q_i(W_t)$ and is estimated by the sum of its relevance scores for all users. Hence, the quality of a category c within a window, denoted as $Q_c(W_t)$ for us, is the sum of the qualities of all items belonging to that category.

For each user u, there is a top-k ranking list of the k best system-estimated item suggestions. The weight vector $w = \{w_1, w_2, \ldots, w_k\}$, where w_r represents the exposure received by an item at a rank r, which brings the notion of position bias where, $w_1 \geq w_2 \geq \cdots \geq w_k \geq 0$, representing the diminishing visibility of lower positions [Do et al. 2021, Usunier et al. 2022]. Similarly, the exposure of a category is defined as the sum of the items' exposures that belong to that category.

In recommendation systems, the tendency for items at higher ranks to receive significantly more attention, leading to position bias, often results in an unfair distribution of visibility [Joachims et al. 2017]. To ensure that the exposure is proportional to the quality of each category, the total quality in a window W_t is defined as $Q = \sum_{c\in\mathcal{C}} Q_c(W_t)$, and the total exposure as $E = \sum_{c\in\mathcal{C}} E_c$, which represents the sum of all position bias weights across all rankings. The target exposure for a category c, denoted as T_c is defined as:

$$T_c = E \cdot \frac{Q_c(W_t)}{Q}$$

serving as a fairness benchmark for the recommendation process. While the framework generates rankings, the main objective is to approximate each category's actual exposure E_c to its target T_c, thereby mitigating unfair exposure.

4. Lifetime-aware Item Fair Exposure (LIFE)

The objective of the proposed framework is to address quality-weighted exposure for items with limited lifetimes in recommendation systems, ensuring that the ranking reflects not only users' historical preferences but also the temporal availability of items, thereby achieving weighted visibility over items' categories. The LIFE approach is implemented as a post-processing framework that modifies the output ranks produced by a matrix factorization model [Koren et al. 2009]. The essence of our method is the integration of temporal aspects of items, considering urgency in the exposure decision process, ensuring that items are recommended not only for relevance but also for their remaining window of opportunity. The system operates over a sequence of time windows. Our $\mathcal{U}_I(t)$ acts as a "last chance" mechanism, elevating the exposure priority for items about to expire. In that way, the system can better estimate their true quality over time, helping items with few interactions receive greater exposure before their lifetime ends, and it loses its peak of relevance.

The LIFE algorithm is shown in Algorithm 1. Before initiating the process of changing users' ranks, the framework establishes a fairness benchmark for each category $c \in C$ in the current window, which we call Target Exposure T_c, proportional to the aggregate quality of its active items and users' interactions. This ensures that our categories with higher predicted relevance receive a fair share of visibility while mitigating popularity bias. It is important to note that recommendation models, such as matrix factorization, are inherently susceptible to popularity bias, as they tend to learn more effectively from data with higher density. The algorithm iterates over each user's ranking and performs strategic swaps between items in the user's top-k and candidates outside it. To preserve the recommendations' integrity and avoid a complete overhaul of the personalized experience, we impose a structural constraint by setting a maximum of $\lfloor \frac{k}{2} \rfloor$ swaps per user. This means the algorithm can replace at most 50% of the original recommendations. The swaps decision is oriented by a *priority score* $\mathcal{P}_i$, computed as

$$\mathcal{P}_i = \alpha \cdot \mathcal{G}(c) + \beta \cdot \mathcal{U}_i(t) + \gamma \cdot \hat{s}_{ui}$$

which is composed of three fundamental pillars, each controlled by a specific hyperparameter that defines the relative importance of each objective in the final ranking: *fairness gain*, *temporal urgency*, and *relevance*.

Fairness Gain ($\mathcal{G}$): Focused on balancing visibility across different item categories, it measures how much the inclusion of a specific item helps bring the category's actual exposure E_c closer to the target exposure T_c calculated for that category. The algorithm calculates the absolute distance between the current state and the target exposure, where the gain is positive, if inserting an item reduces this distance. Fairness weight (α) parameter controls the strength of the categorical exposure adjustment. A higher α pushes the system to prioritize items from under-exposed categories, even if they have slightly lower relevance.

Temporal Urgency ($\mathcal{U}$): It is based on the lifetime L_i. As time passes and the item approaches the end of its peak of relevance, the value of $\mathcal{U}_i(t)$ scales from 0 to 1, elevating items that are about to become less relevant. Urgency weight (β) is crucial in dynamic catalogs to ensure that short-lived content receives sufficient exposure to be evaluated by users before it expires, giving preference to items whose $\mathcal{U}_i(t)$ is close to 1.0

Algorithm 1: LIFE: Lifetime-Aware Fair Exposure Reranking

foreach *user* $u \in U$ **do**
  $R_u \leftarrow$ current Top-k list for user u;
  $swaps \leftarrow 0$;
  while $swaps < \left\lfloor \frac{k}{2} \right\rfloor$ **do**
    Identify eligible item $i_{out} \in R_u$ for removal such that:
    1. $i_{out} \in c$ where $E_c > T_c$
    2. Relevance score $\hat{s}_{ui_{out}} <$ safety_threshold;
    if *no eligible* i_{out} *found* **then**
      break
    $max\_P \leftarrow -\infty$;
    $i_{in} \leftarrow$ null;
    foreach *candidate item* $j \notin R_u$ **do**
      Calculate Urgency: $\mathcal{U}_i(t) = \begin{cases} \frac{t - t_i^{pub}}{L_i}, & \text{if } 0 \leq t - t_i^{pub} \leq L_i \\ 0, & \text{otherwise} \end{cases}$;
      Calculate Fairness Gain: $G(c_j) \leftarrow$ reduction in $|E_{c_j} - T_{c_j}|$ if j is added;
      Calculate Priority: $P_j \leftarrow \alpha \mathcal{G}(c_j) + \beta \mathcal{U}_j(t) + \gamma \hat{s}_{uj}$;
      if $P_j > max\_P$ **then**
        $max\_P \leftarrow P_j$;
        $i_{in} \leftarrow j$;
    Replace i_{out} with i_{in} in R_u;
    Update global exposure counters $E_{c_{out}}$ and $E_{c_{in}}$;
    $swaps \leftarrow swaps + 1$;
return *Reranked top-k lists for all users in* W_t;

Relevance ($\hat{s}_{ui}$): Ensures that, despite fairness and adjustments, the recommendation remains useful for the user. It utilizes the predicted scores generated by the Matrix Factorization model and represents the original intent of the pure recommendation algorithm. Relevance weight (γ) ensures the preservation of the user experience. A high γ prevents the framework from making aggressive swaps that would significantly decrease the perceived quality of the recommendation list.

Although the priority score guides the selection of new items, the framework aims not to pursue fairness at the expense of the user experience. To preserve some quality, the algorithm applies constraints during the swap process: An item is only eligible for removal if it belongs to an over-exposed category ($E_c > T_c$) and its relevance score is below a safety threshold. Also every time a swap is performed, the global category exposure counters are updated in real time, this ensures that fairness objectives are distributed across the entire user base within the window.

Figure 1 illustrates by example the pipeline of the algorithm developed in our studies. In the first module, the system starts with a set of items categorized into groups. Each item $i \in I$ carries a quality score q_i and a lifetime L_i, representing how much time

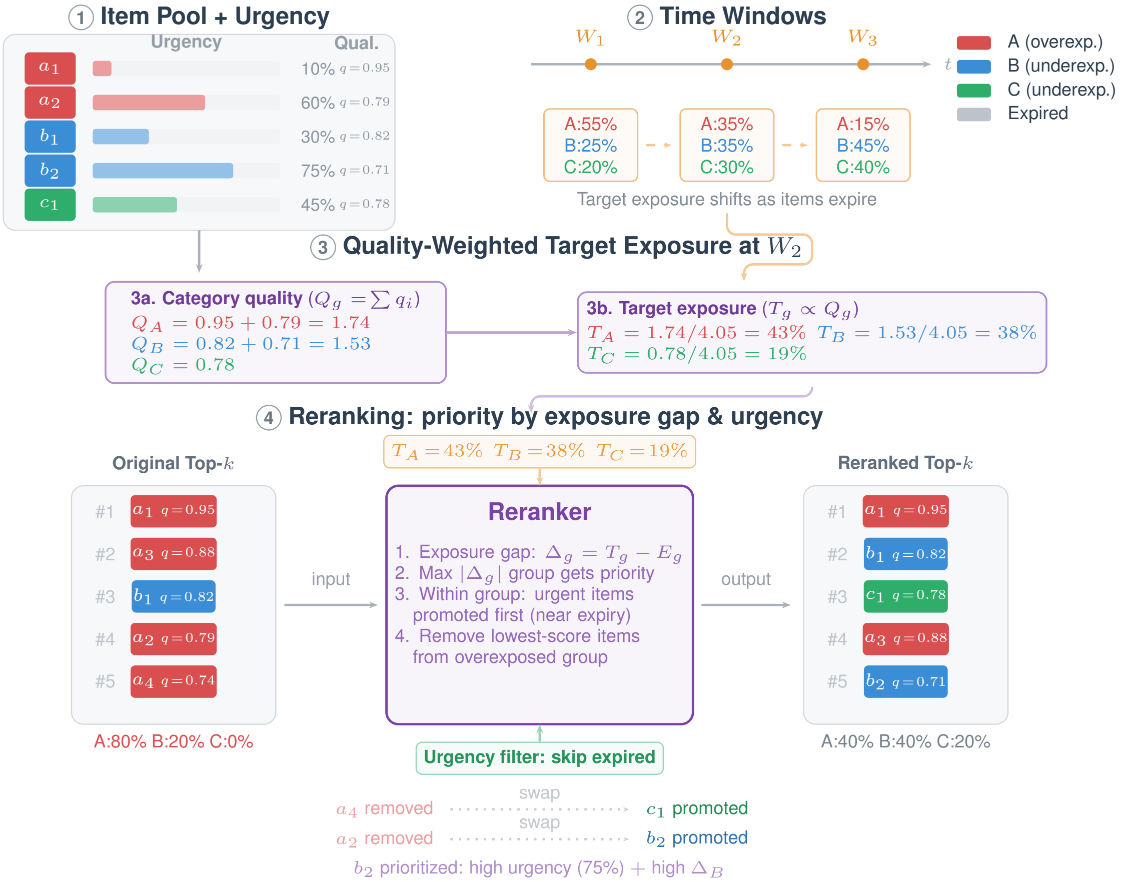

Figure 1. Quality-weighted post-processing reranking with time-aware group fairness.

remains before the item expires, and is a critical variable for prioritizing content about to disappear. The framework operates across time windows, where, as time progresses, the target exposure for each group shifts in order to account for the items present in the window. At stage 3, we see the ideal exposure target for each group in the current window being calculated: 3a is the sum of the individual quality scores of active items for each group, and 3b is the exposure proportional to each group's total quality relative to the entire pool. The 4 step is the decision maker; the *reranker* takes the original top-k and transforms it. It identifies the difference between what a group should have and what it currently has. Underexposed groups are prioritized, and items from those groups with low time-to-live are promoted first to ensure they get a chance to be seen by users. The result is a reranked top-k distribution that is more similar to the expected target in that time window.

Let A_t denote the set of active users within a given time window W_t, and let $\mathcal{I} \subseteq I$ represent the total set of candidate items available for swapping. The computational complexity of LIFE depends primarily on the number of active users per time window, the total number of candidate items, and the maximum allowed swaps. During the reranking phase, the algorithm evaluates potential replacements for each active user, which, in the worst-case scenario, this evaluation reaches the maximum swap limit, resulting in a time complexity of $\mathcal{O}(|A_t||\mathcal{I}|\left\lfloor \frac{k}{2} \right\rfloor)$. Since $\frac{k}{2}$ is a constant value derived from the top-k

recommendation size, the predominant complexity is $\mathcal{O}(|A_t||\mathcal{I}|)$ per time window.

5. Experiments and Results

This section presents the experimental evaluation of the proposed approach. We describe the datasets and their experimental setups, the hyperparameter settings, the temporal evaluation protocol, and discuss the results obtained with LIFE [1].

Datasets and experimental setup

We conduct the experimental evaluation on three benchmark datasets: MovieLens 100K [Harper and Konstan 2015] and two subsets of Amazon Product Data, namely Grocery and Gourmet Food and Health and Personal Care[2]. Each dataset was preprocessed so that each record corresponds to a tuple containing a user, an item, a rating, and a timestamp.

To ensure that each evaluation step contains a sufficient number of observations, we impose a minimum density requirement of 100 active users per window. Whenever a window does not satisfy this threshold, it is extended by incorporating the subsequent temporal block until the minimum number of active users is reached. This adaptive procedure ensures that each evaluation window remains statistically meaningful.

Because the datasets differ in their temporal scope and interaction density, the sliding-window configuration is adjusted for each scenario in order to better reflect a real-world setting.

The MovieLens 100K dataset spans a short, dense time period from September 1997 to April 1998. Although it spans less than one year, it contains multiple interactions on the same day. To capture user preferences over this compressed temporal range, we use a 30-day warm-up period and a 7-day sliding window. This configuration enables evaluation across a larger number of windows despite the dataset's limited overall duration.

In contrast, the Amazon Product Data subsets cover a much longer period, with reviews spanning from May 1996 to July 2014. To accommodate this longitudinal scale and ensure that each evaluation step remains statistically representative, we adopt a 180-day warm-up period and a 30-day sliding window. Whereas MovieLens exhibits a denser, more continuous stream of ratings, the Amazon subsets exhibit substantially sparser, more irregular interaction patterns.

In the experimental design and categories assignment, we impose a high imbalance of category distribution within each dataset. Specifically, during the random assignment of categories to items in the dataset, we ensured that Category 1 received a higher probability, resulting in significantly more information about this first category. By forcing this imbalance, we can evaluate the framework´s ability to override the model's base popularity bias and achieve exposure fairness in line with our quality-weighted targets.

[1] For reproducibility purposes, the experiments are available on `https://github.com/MiBC03/life_algorithm.git`

[2] `https://jmcauley.ucsd.edu/data/amazon/`

Also, for online experimental purposes, we define L_i as a random value within a specified range.

Hyperparameter settings

To address the multi-objective aspect of our approach and facilitate the interpretation of each component, we normalize the hyperparameters such that the weights for fairness (α), urgency (β), and relevance (γ) sum to 1. To identify the optimal configuration and evaluate hyperparameter sensitivity, we conducted a grid search across all evaluated datasets, testing weight combinations. The empirical results demonstrated that the configuration $\alpha = 0.4, \beta = 0.2$ and $\gamma = 0.4$ provides an ideal, balanced baseline, assigning a substantial and equal weight to categorical justice and user relevance, while allocating sufficient priority to the last chance mechanism. For our experiments, the recommendation list size k was set to 10, which inherently defined our maximum swap limit to 5 items per user.

Temporal evaluation protocol

As the proposed approach simulates a real-world production environment in which new interactions become available over time, the experimental protocol follows an iterative training procedure. The model is first trained on an initial set of historical interactions before the reranking strategy is applied. Then, for each time window W_t, the model is retrained using the initial interaction set together with all interactions observed from W_1 to W_{t-1}, and subsequently evaluated on the interactions associated with the current window W_t. At each iteration, the training and test sets remain disjoint, thereby preventing data leakage and supporting the validity of the evaluation.

To assess the exposure fairness achieved by LIFE , we compute, for each time window, the percentage of exposure allocated to each category and then report the average across all windows. We compare three distributions: the target exposure distribution, the exposure induced by the baseline ranking (i.e., the raw top-k), and the exposure obtained after reranking. Furthermore, we also calculate standard quality metrics, such as precision, recall, NDCG and MAP, to evaluate the trade-off between fairness and the ranking utility. This allows us to verify whether the adjustments made to ensure fair exposure compromise the relevance of the recommendations provided to the users and negatively impact the overall user experience.

Discussion

The results of LIFE applied in the MovieLens dataset are summarized in Table 1. It has an expected concentration of exposure in the dominant category across all ranking strategies. Yet, it presents a distinct dynamic in its baseline ranking, where category 2 experiences significant overexposure in the traditional, receiving 26.59% of visibility, which exceeds its quality-based target of 19.46%. This disproportional concentration prevents the dominant group from reaching its true potential, leaving it with only 57.63% despite a target of 61.23%. Furthermore, Category 3 is also underexposed in the baseline model, receiving only 6.55% against a target of 10.10%, while Category 4 already aligns perfectly with its target at 9.23%. Ultimately, the re-ranking strategy corrects these imbalances, adjusting every category to closely match its respective target by correcting the over-exposed category 2 and redistributing its excess visibility precisely to categories 1 and 3, all while maintaining the existing stability of category 4.

Table 1. Comparison between Target, top-k, and reranked results in MovieLens

Category	Orig.	Rerank	Target	Metric	Base	Drop (%)
1	57.63%	**61.23%**	**61.23%**	Precision@10	0.1724	-6.07%
2	26.59%	**19.78%**	**19.46%**	Recall@10	0.0915	-6.07%
3	6.55%	**9.79%**	**10.10%**	NDCG@10	0.1932	-4.18%
4	9.23%	**9.22%**	**9.22%**	MAP@10	0.1042	-5.41%

The Health dataset results, in Table 2, exhibit a tendency to favor the top groups in the baseline ranking. Despite that, category 2 is slightly overexposed at 24.57% against its 21.59% target. Consequently, the smaller groups are suppressed. Category 3 has visibility of only 7.53%, falling short of its 9.98% target, and category 4 is similarly at 7.01%, below its 9.08% expected exposure. The reranking algorithm adapts to these irregularities, curbing the overrepresentation by reducing category 1 to 58.81%, which is slightly below its target, and pulling category 2 down to a much more balanced 22.48%. Meanwhile, the suppressed categories experience a crucial recovery.

Table 2. Comparison between Target, top-k, and reranked results in Amazon Health and Personal Care

Category	Orig.	Rerank	Target	Metric	Base	Drop (%)
1	60.88%	**58.81%**	**59.36%**	Precision@10	0.0102	-8.00%
2	24.57%	**22.48%**	**21.59%**	Recall@10	0.0618	-8.79%
3	7.53%	**9.61%**	**9.98%**	NDCG@10	0.0358	-6.08%
4	7.01%	**9.09%**	**9.08%**	MAP@10	0.0240	-4.59%

Table 3 contains the results of LIFE for the Grocery dataset. It reveals an imbalance in the baseline ranking, primarily by the dominant group. In this execution, category 1 is significantly overexposed, capturing 62.88% of the visibility compared to its 59.82% merit. This over-concentration leads to an under-exposure for all the remaining groups. Category 4 has visibility of 9.37%, below its 10.36% target, and category 2 is suppressed to 17.13%, falling short of its 18.89% expected exposure. Our method addresses this disparity by redistributing the excess of exposure from category 1 to where it is needed. Consequently, category 4 is elevated to 10.66%, slightly surpassing its target. Similarly, category 2 rises from its initial 17.13% to a stable 18.54%, significantly closer to its expected quality, and category 3 receives a proportional boost, reaching 11.12%. With that, the algorithm demonstrates its ability to identify multi-layered exposure biases and reallocate visibility from multiple over-represented classes directly to the most starved groups.

Analyzing the impact of the re-ranking algorithm shows the expected trade-off between fairness and accuracy in recommender systems. By redistributing exposure toward

Table 3. Comparison between Target, top-k, and reranked results in Amazon Grocery and Gourmet Food

Category	Orig.	Rerank	Target	Metric	Base	Drop (%)
1	62.88%	**59.68%**	**59.82%**	Precision@10	0.0089	-7.41%
2	17.13%	**18.54%**	**18.89%**	Recall@10	0.0633	-7.41%
3	10.63%	**11.12%**	**10.93%**	NDCG@10	0.0358	-5.13%
4	9.37%	**10.66%**	**10.36%**	MAP@10	0.0245	-3.52%

underrepresented categories, the method inevitably introduces some loss in traditional recommendation metrics. The magnitude of the trade-off varies across datasets, with some domains showing greater sensitivity to ranking modifications than others. Nevertheless, the proposed approach reduces exposure imbalance and increases the visibility of disadvantaged categories. Overall, the experiments demonstrate that the re-ranking strategy balances recommendation utility and exposure fairness, achieving meaningful fairness gains.

6. Conclusion

We propose LIFE , a post-processing approach that addresses the fairness of exposure in recommendation systems for items with limited lifetimes. We propose that ideally, the attention allocated to items should be proportional to relevance within the overall set. However, for time-sensitive items, failing to expose them during their peak relevance window results in a loss of utility. We apply a "last chance" mechanism that elevates the exposure priority of items approaching the end of their lifespan, ensuring they reach users before they lose their peak relevance. As a result, the system balances user relevance, provider fairness, and item urgency by proportionally estimating quality over a sliding time window. We compare our approach against a Matrix Factorization baseline in scenarios that simulate a chronological real-world data flow. We empirically demonstrate that the LIFE reranking strategy corrects exposure imbalances with high precision, achieving alignment with quality-based targets without compromising the integrity of the personalized user experience.

Moving forward, we aim to replace randomized lifetime variables (L_i) with real-world temporal data, such as product expiry dates or the shelf-life of news articles, to enhance model accuracy. Furthermore, replacing the synthetic category assignments used to create imbalance in this study with authentic item taxonomies or provider-level groupings would enable a more robust evaluation of the framework's effectiveness in real-world scenarios.

Acknowledgments

We acknowledge the support of the Conselho Nacional de Desenvolvimento Científico e Tecnológico (CNPq) grant 308635/2025-6, of the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior (CAPES) grant 88887.694847/2022-00, of the Fundação de Amparo à Pesquisa do Estado de São Paulo (FAPESP) grant 2025/04319-9, and of Lenovo, as part of its R&D investment under Brazil's Informatics Law.

References

Amanatidis, G., Aziz, H., Birmpas, G., Filos-Ratsikas, A., Li, B., Moulin, H., Voudouris, A. A., and Wu, X. (2023). Fair division of indivisible goods: Recent progress and open questions. *Artificial Intelligence*, 322:103965.

Arnsperger, C. (1994). Envy-freeness and distributive justice. *Journal of Economic Surveys*, 8(2):155–186.

Bae, H.-K., Ahn, J., Lee, D., and Kim, S.-W. (2023). Lancer: A lifetime-aware news recommender system. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 4141–4148.

Ben-Porat, O. and Torkan, R. (2023). Learning with exposure constraints in recommendation systems. In *Proceedings of the ACM Web Conference 2023*, pages 3456–3466.

Biega, A. J., Gummadi, K. P., and Weikum, G. (2018). Equity of attention: Amortizing individual fairness in rankings. *CoRR*, abs/1805.01788.

Brams, S. J. and Taylor, A. D. (1995). An envy-free cake division protocol. *The American Mathematical Monthly*, 102(1):9–18.

Collins, T., Beel, J., and Carey, S. (2018). A study of position bias in digital library recommender systems. In *Proceedings of the 18th ACM/IEEE Joint Conference on Digital Libraries (JCDL '18)*, pages 335–344, New York, NY, USA. Association for Computing Machinery.

de Lourdes M. Silva, M., Chaves, I., Mendonça, A. L., Neto, E. D., and Machado, J. (2025). Twix: Balancing fairness and utility in item exposure for recommendation systems. In *Anais do XL Simpósio Brasileiro de Bancos de Dados*, pages 427–440, Porto Alegre, RS, Brasil. SBC.

Do, V., Corbett-Davies, S., Atif, J., and Usunier, N. (2021). Two-sided fairness in rankings via lorenz dominance. *Advances in Neural Information Processing Systems*, 34:8596–8608.

Farhadi, A., Hajiaghayi, M., Latifian, M., Seddighin, M., and Yami, H. (2021). Almost envy-freeness, envy-rank, and nash social welfare matchings. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 5355–5362.

Guo, H., Sun, Z., Wang, D., Wei, T., Li, J., and Zhang, J. (2025). Enhancing new-item fairness in dynamic recommender systems. In *Proceedings of the 48th International ACM SIGIR conference on research and development in Information Retrieval*, pages 1707–1716.

Harper, F. M. and Konstan, J. A. (2015). The movielens datasets: History and context. *Acm transactions on interactive intelligent systems (tiis)*, 5(4):1–19.

Heuss, M., Sarvi, F., and de Rijke, M. (2022). Fairness of exposure in light of incomplete exposure estimation. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 759–769.

Joachims, T., Granka, L., Pan, B., Hembrooke, H., and Gay, G. (2017). Accurately interpreting clickthrough data as implicit feedback. In *Acm Sigir Forum*, volume 51, pages 4–11. Acm New York, NY, USA.

Koren, Y., Bell, R., and Volinsky, C. (2009). Matrix factorization techniques for recommender systems. *Computer*, 42(8):30–37.

Li, Y., Chen, H., Fu, Z., Ge, Y., and Zhang, Y. (2021). User-oriented fairness in recommendation. In *Proceedings of the web conference 2021*, pages 624–632.

Macedo, A. and Marinho, L. (2014). Event recommendation in event-based social networks. *Proc. of Int. Work. on Social Personalization*.

Morik, M., Singh, A., Hong, J., and Joachims, T. (2020). Controlling fairness and bias in dynamic learning-to-rank. In *Proceedings of the 43rd international ACM SIGIR conference on research and development in information retrieval*, pages 429–438.

Rawls, J. (2017). A theory of justice. In *Applied ethics*, pages 21–29. Routledge.

Raza, S. and Ding, C. (2022). News recommender system: a review of recent progress, challenges, and opportunities. *Artificial Intelligence Review*, 55(1):749–800.

Singh, A. and Joachims, T. (2018). Fairness of exposure in rankings. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*, pages 2219–2228.

Usunier, N., Do, V., and Dohmatob, E. (2022). Fast online ranking with fairness of exposure. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, pages 2157–2167.

Wang, J., Ma, W., Li, J., Lu, H., Zhang, M., Li, B., Liu, Y., Jiang, P., and Ma, S. (2022). Make fairness more fair: Fair item utility estimation and exposure re-distribution. In *Proceedings of the 28th ACM SIGKDD conference on knowledge discovery and data mining*, pages 1868–1877.

Wang, Y., Ma, W., Zhang, M., Liu, Y., and Ma, S. (2023). A survey on the fairness of recommender systems. *ACM Transactions on Information Systems*, 41(3):1–43.

Wu, Y., Cao, J., and Xu, G. (2023). Faster: A dynamic fairness-assurance strategy for session-based recommender systems. *ACM Transactions on Information Systems (TOIS)*, 42(1):1–26.

Wu, Y., Cao, J., Xu, G., and Tan, Y. (2021). Tfrom: A two-sided fairness-aware recommendation model for both customers and providers. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '21, page 1013–1022, New York, NY, USA. Association for Computing Machinery.

Zhang, F., Liu, Q., and Zeng, A. (2017). Timeliness in recommender systems. *Expert Systems with Applications*, 85:270–278.