FaDeRS: Fairness and Depolarization in Recommender Systems

Rafael Santos
rafaelv@ifes.edu.br
Federal Institute of Espirito Santo
Cachoeiro de Itapemirim, Brazil

Giovanni Comarela
gc@inf.ufes.br
Federal University of Espirito Santo
Vitória, Brazil

ABSTRACT

Recommender systems have become fundamental in modern society as they help navigate the vast amount of available data, enabling users to find information, products, or services more efficiently and personally. They directly impact how people consume data, goods, and resources. Recommender systems often lack fairness and diversity, resulting in unfair services and increased preference polarization. Solution: This work presents FaDeRS (Fairness and Depolarization in Recommender Systems), an approach aimed at increasing fairness and diversity in recommender systems. FaDeRS adjusts predictions through controlled perturbations and optimization to mitigate individual unfairness and polarization without modifying the input data. The research is related to socio-technical theory, addressing one of the socio-algorithmic problems, algorithmic discrimination. We consider a specific set of approaches to encode fair behaviors. The research applied a quantitative method with experimentation using two datasets in distinct contexts, implementing a post-processing algorithm based on the Simulated Annealing meta-heuristic. The proposed method demonstrated significant reductions in polarization (up to 78.64%) and individual unfairness (up to 33.97%), with only a small increase in the Root Mean Square Error (RMSE), indicating an improvement in the socially desirable qualities of the systems without unduly sacrificing accuracy. Notably, FaDeRS consistently outperformed a relevant benchmark methodology across both evaluated datasets. The main contribution is a mechanism that balances personalization and fairness, simultaneously addressing polarization and individual unfairness from the items' perspective, promoting a fairer and more diverse approach to recommendation.

KEYWORDS

recommendation system, fairness, polarization

1 INTRODUCTION

Recommender systems play a vital role in contemporary society, where the explosion of data and the overwhelming amount of available information make it essential to use automated techniques to help users discover relevant content [31]. The ability of these systems to anticipate people's preferences and provide personalized recommendations directly influences the consumption of information, products, and services [12, 30, 34, 45, 46, 55]. However, studies show how recommender systems predicting preferences can offer unfair or unequal service quality to users [51, 54] or lead to social polarization, increasing divergence among user preferences [10].

In: Proceedings of the Brazilian Symposium on Multimedia and the Web (WebMedia'2025). Rio de Janeiro, Brazil. Porto Alegre: Brazilian Computer Society, 2025.

ISSN 2966-2753

This article explores a new approach by examining whether a fairness-focused re-ranking, without the need for input data transformation, can, by itself, improve the social desirability of the resulting recommendations. The issue is addressed by developing a generic method that can be used to enhance two socially relevant properties of the recommender system: individual unfairness and polarization. The proposed method uses controlled perturbations of the original recommendations, combined with an optimization algorithm, to compose a method that generates socially desirable results.

As input, recommendations from a traditional system are used. FaDeRS begins with perturbations of the original recommendations, equalized by a factor (α), aiming to reduce polarization. Then, an optimization algorithm, based on the Simulated Annealing meta-heuristic[1], searches for the best combination of recommendations that minimizes individual unfairness, generally providing reductions in both metrics.

In this article, the method is instantiated with the proposition of metrics that capture polarization and individual unfairness of recommender systems. A series of computational efficiencies that can be explored is presented. In the results, the relationship between improvements in individual unfairness measures and polarization, and changes in the overall system accuracy is considered. Finally, it is shown that a moderate adjustment factor (α) combined with a small number of estimated matrices (h) can generate a substantial reduction in polarization, 78.64% and 77.95%, and individual unfairness, 31.17% and 33.97%, respectively, in the MovieLens-1M[2] and GoodBooks-10k[3] datasets.

The presented method offers a coordinated improvement in both the polarization measure and individual unfairness by adjusting the factor (α). As the value of α is increased, a significant reduction in individual unfairness and polarization is observed, allowing the recommender system to become more balanced and socially fair. However, this improvement is accompanied by an increase in the Root Mean Square Error (*RMSE*), which implies a reduction in recommendation accuracy. Therefore, it is up to the system user to decide on the best configuration, balancing between system accuracy and improvements in the socially desirable qualities of the recommendations.

Previous works have explored a series of approaches incorporating fairness in learning models and recommender systems. These approaches can be broadly categorized into those that depend on (i) pre-processing, i.e., transforming the training data to reduce the potential for unfair outcomes using traditional learning models [6], (ii) in-processing, altering learning objectives and models to ensure fair outcomes, even with unmodified training data [19], and

[1]Used to enhance the searchability of optimization algorithms [40]
[2]https://grouplens.org/datasets/movielens
[3]https://www.kaggle.com/datasets/zygmunt/goodbooks-10k

(iii) post-processing, modifying unfair outcomes from pre-trained learning models [48].

In this work, a post-processing strategy is proposed to mitigate polarization and individual unfairness of items, promoting diversity and ensuring a fair distribution of opportunities among recommended items. Its contributions can be summarized as follows:

- Simultaneous consideration of mitigating polarization and individual unfairness, while most works focus on improving a single metric [9, 35, 49, 62];
- Emphasis on fairness from the perspective of items [1, 33], seeking to ensure fair recommendation opportunity, contrasting with the user perspective [3, 6, 19, 28];
- Elimination of the need for direct modifications in the traditional recommendation algorithm to incorporate fairness principles, in contrast to several other studies [5, 56, 58, 65]. The solution proposed in this article employs a post-processing intervention, based on a re-ranking approach, to ensure fairness.
- Preservation of privacy and integrity of users' sensitive data (such as gender, race, age), while improving social metrics at the individual level, in distinction to group fairness approaches that require this information [15, 22, 25, 64].

The topics will be detailed in subsequent sections, starting with related works in Section 2. Section 3 presents an overview of the proposal. Experimental settings are presented in Section 4, results in Section 5, and Conclusion in Section 6.

2 RELATED WORK

In this section, we address the theoretical foundations of the study, focusing on the definitions of fairness and polarization that underpin our research. We discuss general approaches for incorporating fairness into machine learning algorithms, as well as the specific challenges faced by recommender systems when dealing with equity issues. Additionally, we examine the phenomena of polarization and individual unfairness in these systems and review the strategies proposed to minimize their impact.

2.1 Fairness in Machine Learning

Fairness has become a topic of increased interest in the domain of machine learning. Following this discussion, a recommender system will be considered fair if it ensures uniformity in the quality of service (i.e., accuracy of predictions) for all individuals or user groups [60].

Awareness about the negative social impacts stemming from the application of machine learning algorithms in decision-making contexts has been growing [20, 50]. In response, various metrics and concepts of fairness have been proposed for machine learning tasks, encompassing classification [2, 13, 14, 23, 26, 52], regression [21, 41, 43, 59], ranking [4, 42, 44, 61], and set selection [8, 11, 38, 53].

Compared to learning tasks such as classification and regression, few studies have explored notions of fairness in the context of recommender systems. In the latter case, the proposals are divided into two main categories: fairness assessments at the individual level [16, 57] and at the group level [17, 18, 24].

Most studies address fairness by proposing constraints through regularizers at the group level, identifying sensitive attributes like race or gender to seek parity in statistics such as accuracy and true positive rates [32]. However, these approaches are limited when there is uncertainty in class labeling [63] and may fail due to their aggregative nature [7]. In contrast, individual fairness provides a more detailed assessment by operating at a finer granularity, at the individual level, and is less restrictive as it does not require explicit identification of sensitive attributes [62].

2.2 Individual Fairness

In recommender systems, this refers to the disparity in the treatment of individual users, where the recommendation model favors or disadvantages certain users inconsistently or disproportionately. This occurs when the predictions or recommendations made by the system result in significantly different outcomes for users with similar profiles or preferences, without clear justification based on their interactions with the system [62].

Individual unfairness can be measured from the perspective of users or items. [33] suggested that fair recommendations should consider both user satisfaction and the equitable distribution of opportunities among items.

From the perspective of items, it refers to the phenomenon where some items, despite their potential value, are rarely recommended due to the popularity logic of the algorithms. This not only creates a bias in user exposure but also generates an unequal distribution of opportunities among available items [1].

Most fairness algorithms in recommender systems focus on the user perspective, ensuring fair and balanced recommendations. However, recent studies highlight the importance of the item perspective for a fair allocation of recommendation opportunities. For example, [33] investigated fairness in music recommendation, emphasizing the balance between relevance, fairness, and record label satisfaction. Also in this context, the popularity bias was discussed by [1], who proposed personalized re-ranking techniques.

2.3 Polarization

Polarization is the degree to which opinions, viewpoints, and sentiments diverge within a population. Recommender systems can capture this effect through the ratings that users provide for items [39].

In this context, a filter bubble is a phenomenon where users are selectively exposed to information and content that reinforce their pre-existing beliefs and interests, creating an intellectual isolation environment and limiting informational diversity. This effect can have wide-ranging consequences, including polarization of opinions, formation of echo chambers, and reduced exposure to new ideas and perspectives [36, 47]. Therefore, filter bubbles describe the mechanism, while polarization describes the social effect that may arise as a consequence [37].

Evidence that recommender systems reduce information diversity was found by [36]. Recently, with the implementation of algorithms to promote recommendation diversity, [19] suggested that diversity metrics can help reduce polarization, while [48] proposed calibrated recommendations that balance relevance and diversity.

3 FADERS METHOD

In this section, we present an overview of the proposed method for depolarization and individual fairness.[4] The method is divided into two modules, outlined in Sections 3.1 and 3.2, respectively.

3.1 Module 1: Measures of Polarization and Fairness

In this module, following the specifications and discussions from the previous section, we formally define the metrics for polarization and individual unfairness. All implementations of the fairness measures used in this algorithm were based on the work by [39].

We start by presenting the system configuration, notation, and problem definition. Suppose $X \in \mathbb{R}^{n \times m}$ is a partially observed rating matrix of n users and m items, where element x_{ij} denotes the rating assigned by user i to item j. Let Ω be the set of indices of known ratings in X. Furthermore, let Ω^i be the set of indices of known item ratings for user i, and let Ω_j be the indices of known user ratings for item j.

For a matrix $\mathbf{A}$, $P_\Omega(\mathbf{A})$ is a matrix whose elements in $(i,j) \in \Omega$ are a_{ij} and zero elsewhere. For a vector $\mathbf{a}$, $P_{\Omega_j}(\mathbf{a})$ is a vector whose elements in $i \in \Omega_j$ are the corresponding elements of $\mathbf{a}$ and zero elsewhere. The column j of $\mathbf{A}$ is denoted by vector $\mathbf{a}_j$ and the row i of $\mathbf{A}$ by vector $\mathbf{a}^i$.

In a traditional recommender system, the system input is $P_\Omega(X)$ and the objective is to find an estimated recommendation matrix $\hat{X}$, which is fully filled and provides the unobserved values of X. In this recommendation problem, users are assumed to be in a set $\{u_1, u_2, \ldots, u_n\}$ and items in a set $\{v_1, v_2, \ldots, v_m\}$.

The proposed method considers the recommendations resulting from the traditional recommender system, $\hat{X}$, and based on the calculation of fairness measures, it proposes a solution focused on reducing polarization and individual unfairness.

Polarization. To capture polarization, we measure the degree to which users' ratings disagree. Thus, we consider the estimated ratings $\hat{X}$ and define the polarization metric as the average of the variances of the estimated ratings for each item [39]:

$$\sigma_j^2 = \frac{1}{n} \sum_{i=1}^{n} (\hat{x}_{ij} - \mu)^2 \tag{1}$$

$$R_{pol}(\hat{X}) = \frac{1}{m} \sum_{j=1}^{m} \sigma_j^2 \tag{2}$$

Individual Fairness. It can be considered from the perspective of users or items. In this work, we focus on items, where for each item j, we define ℓ_j, the loss of item j, as the estimate of the mean squared error over the known ratings for item j [39].

$$\ell_j = \frac{||P_{\Omega_j}(\hat{\mathbf{x}}_j - \mathbf{x}_j)||_2^2}{|\Omega_j|} \tag{3}$$

and individual unfairness as the variance of the losses of the items.

$$R_{indv}(X, \hat{X}) = \frac{1}{m^2} \sum_{k=1}^{m} \sum_{l>k} (\ell_k - \ell_l)^2 \tag{4}$$

[4]Source code is available at https://github.com/ravarmes/recsys-faders

3.2 Module 2: Depolarization and Fairness Measures

This module aims to estimate a recommendation matrix that minimizes the polarization and individual unfairness of the items.

Initially, a traditional recommender system is considered, which generates an estimated recommendation matrix $\hat{X}$, where each item v_j receives a set of recommendations $\{u_1, u_2..., u_n | v_j\}$.

The estimated matrix $\hat{X}$ is then used to generate h other estimated matrices $\hat{X}_1, \hat{X}_2..., \hat{X}_h$. Initially, the mean of the elements of $\hat{X}$ is calculated, i.e., $\mu(\hat{X})$. With the mean, the deviation matrix D is determined, representing the difference between $\hat{X}$ and $\mu(\hat{X})$.

An adjustment factor α is then randomly selected from a uniform distribution in the interval $[0, 0.8]$. The idea is to bring the recommendations closer to the mean to reduce polarization, without equalizing them to the mean, which would increase the *RMSE*. In experimental tests, this range of values provided effective polarization reduction, keeping increases in *RMSE* below 10%.

The estimated matrices $\hat{X}_1, \hat{X}_2..., \hat{X}_h$ are obtained by subtracting the product of the deviation matrix D and the adjustment factor α from the matrix $\hat{X}$. Representing matrices with a focus on depolarization.

$$\hat{X}_p = \hat{X} - \alpha D \tag{5}$$

Each estimated matrix $\hat{X}_p$, where $\{1 \leq p \leq h\}$, has a list of m individual losses $\{\ell_1, \ell_2..., \ell_m | \hat{X}_p\}$, each loss associated with one of the m items.

We define the matrix formed by the values of all the individual losses as $Z = [Z_{ij}]_{h \times m}$ which represents the m individual losses calculated for each of the h estimated matrices $\hat{X}_p$, where $Z_{ij} \in \{\mathbb{R}_+\}$, $\{1 \leq i \leq h\}$, and $\{1 \leq j \leq m\}$, index estimated matrices and items, respectively. Each matrix/item pair is associated with a score Z_{ij} representing the individual loss of item j in relation to the estimated matrix i.

The binary matrix $W = [W_{ij}]_{h \times m}$ is defined to indicate which individual losses will be considered in forming the final estimated matrix $\hat{X}_\pi$, where $W_{ij} \in \{0, 1\}$, $\{1 \leq i \leq h\}$, and $\{1 \leq j \leq m\}$ also indexing estimated matrices and items, respectively. Specifically, if the individual loss of estimated matrix i is considered for item j, then $W_{ij} = 1$; otherwise, $W_{ij} = 0$. The determination of this matrix W is performed by an optimization algorithm, as detailed in the subsequent Section 3.3.

For a better understanding of the strategy, refer to Figure 1 and the Pseudocode in Algorithm 1, whose five main steps can be detailed as follows.

1. Prediction: the matrix X, partially filled, is considered by a traditional recommender system to make recommendation predictions in $\hat{X}$;
2. Estimated matrices: h matrices generated by controlled perturbations in $\hat{X}$ aiming to minimize R_{pol};
3. Individual losses: individual errors of items in each of the h matrices are calculated to construct the matrix Z;
4. Binary solution matrix: via optimization algorithm, the binary matrix W is generated, which represents the best combination of recommendations that minimizes R_{indv};
5. Final matrix: from W the solution matrix $\hat{X}_\pi$ is structured.

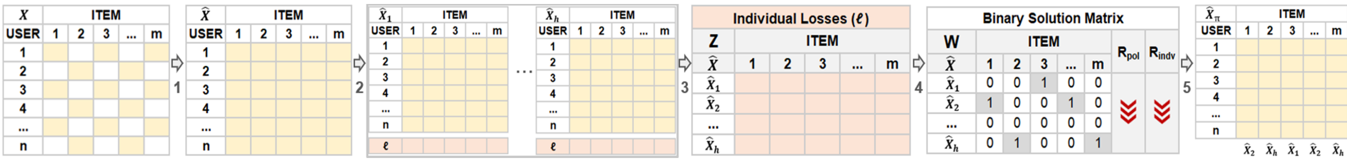

Figure 1: Method FaDeRS operational architecture

Algorithm 1 Operation of the FaDeRS Method

Input: Partially observed matrix $X \in \mathbb{R}^{n \times m}$ of n users and m items, number h of estimated matrices to be generated
Output: Estimated matrix $\hat{X}_\pi$ with the lowest indices of R_{pol} and R_{indv}

Calculate $\hat{X}$
Calculate R_{pol}, R_{indv}, *RMSE* of $\hat{X}$
Calculate h estimated matrices $\{\hat{X}_1, \hat{X}_2, ..., \hat{X}_h\}$
for $p \leftarrow 1$ to h **do**
  Calculate $\{\ell_1, \ell_2..., \ell_m\}$, R_{pol}, R_{indv}, *RMSE* of $\hat{X}_p$
end for
$Z_{\ h \times m} \leftarrow$ Build the individual loss matrix (ℓ)
$W_{\ h \times m} \leftarrow$ Build the binary solution matrix using **Simulated Annealing meta-heuristic** (Algorithm 2)
$\hat{X}_\pi \leftarrow$ Build the final estimated matrix via matrix W

3.3 Optimization with Simulated Annealing

The final stage of FaDeRS involves optimizing the selection of individual item predictions from the h generated estimated matrices ($\hat{X}_p$) to form the final recommendation matrix $\hat{X}_\pi$. This optimization aims to minimize both individual unfairness (R_{indv}) and, to a lesser extent, polarization (R_{pol}). To achieve this, we employ the meta-heuristic **Simulated Annealing (SA)**. SA is particularly suitable for this combinatorial optimization problem due to its robust ability to explore vast search spaces and escape local optima.

The Simulated Annealing algorithm operates on the previously calculated individual loss matrix (ZIL) and polarization loss matrix (ZIP). Both matrices are normalized before being used in the objective function to ensure a balanced contribution from each metric. The objective function that SA seeks to minimize is defined as:

$$\begin{aligned} \text{Objective} =& w_{pol} \cdot \text{Mean}(\text{ZIP}_{\text{norm}} \text{ where } W = 1) \\ &+ w_{indv} \cdot \text{Variance}(\text{ZIL}_{\text{norm}} \text{ where } W = 1) \end{aligned} \tag{6}$$

Based on empirical tests, the weights $w_{pol} = 0.01$ for polarization and $w_{indv} = 1.0$ for individual unfairness were used. This specific weighting was adopted to address the differing scales of the two normalized metrics (R_{pol} being a mean and R_{indv} a variance, typically resulting in R_{pol} having a larger magnitude). By assigning a significantly smaller weight to R_{pol}, the objective function prioritizes the minimization of individual unfairness, which is the primary focus of this optimization stage, while still considering the polarization effect as a secondary, balancing factor. This approach ensures that the optimization process is effectively and consistently driven by the desired reduction in R_{indv}, without being disproportionately influenced by the inherent scale, variability, or contextual sensitivity of R_{pol} across different scenarios.

The Simulated Annealing process unfolds as follows:

(1) **Initialization:** The algorithm starts by randomly generating an initial candidate solution. This solution is represented by the binary matrix $W \in \{0, 1\}^{h \times m}$, where each column (corresponding to an item) contains exactly one '1'. This '1' indicates which of the h estimated matrices ($\hat{X}_p$) is selected to provide the predictions for that specific item in the final $\hat{X}_\pi$ matrix. An initial high temperature (`initial_temp`) is set, allowing for extensive exploration of the solution space.
(2) **Iteration and Neighborhood Exploration:** In each iteration, a neighboring solution ($\mathtt{W}_{\mathtt{new}}$) is generated by applying a small perturbation to the current solution ($\mathtt{W}_{\mathtt{current}}$). This perturbation involves randomly selecting an item's column in W and moving the '1' from its current row (representing the currently selected $\hat{X}_p$) to a different, randomly chosen row (another $\hat{X}_p$).
(3) **Acceptance Criterion:** The new solution's quality (*objective$_{new}$*) is evaluated using the defined objective function. If *objective$_{new}$* is better (lower) than *objective$_{current}$*, the new solution is always accepted. If *objective$_{new}$* is worse, it can still be accepted with a probability governed by the Boltzmann function: $e^{(-\Delta\text{Objective}/\text{Temperature})}$. This probabilistic acceptance allows the algorithm to escape local optima by occasionally accepting less optimal solutions, particularly at higher temperatures.
(4) **Cooling Schedule:** The system's temperature is gradually decreased in each iteration by multiplying it with a `cooling_rate` (e.g., 0.995). As the temperature drops, the probability of accepting worse solutions diminishes, guiding the algorithm towards converging on a near-optimal solution.

The entire process of building the binary solution matrix W using Simulated Annealing is detailed in Algorithm 2.

4 MATERIALS AND METHODS

This section describes the scientific methodology employed in the computational experiments. The research adopts a quantitative approach and is epistemologically positioned in the critical field, with the main objective of promoting social fairness by mitigating polarization and individual unfairness in recommender systems. The experimentation method is used to manipulate and control variables that influence recommendation outcomes. The aim of the study is to explain the phenomenon of mitigating polarization and individual unfairness, aiming for a more equitable and diverse distribution of recommendations. Data was collected from two datasets from different contexts and prepared for experimentation. The analysis

Algorithm 2 Simulated Annealing for Binary Solution Matrix W

Input: Individual loss matrix ZIL, Polarization loss matrix ZIP, h, m, *num_iterations*, *initial_temp*, *cooling_rate*
Output: Optimal binary solution matrix $W \in \{0,1\}^{h \times m}$

Normalize ZIL and ZIP to ZIL_{norm} and ZIP_{norm}
Define objective function $f(W)$ as:

$$\text{Objective} = 0.01 \cdot \text{Mean}(\text{ZIP}_{\text{norm}} \text{ where } W = 1) + \text{Variance}(\text{ZIL}_{\text{norm}} \text{ where } W = 1)$$

Initialize W_{current} randomly (one '1' per column)
$W_{\text{best}} \leftarrow W_{\text{current}}$
$\text{objective}_{\text{best}} \leftarrow f(W_{\text{best}})$
temperature $\leftarrow$ initial_temp
for $k \leftarrow 1$ to *num_iterations* **do**
  Generate W_{new} by moving a '1' in a random column of W_{current}
  $\text{objective}_{\text{new}} \leftarrow f(W_{\text{new}})$
  $\Delta\text{objective} \leftarrow \text{objective}_{\text{new}} - \text{objective}_{\text{current}}$
  if $\Delta\text{objective} < 0$ **then**
    Accept W_{new} as W_{current}
    if $\text{objective}_{\text{current}} < \text{objective}_{\text{best}}$ **then**
      Update W_{best} and $\text{objective}_{\text{best}}$
    end if
  else ▷ Accept worse solution with probability
    if $\text{Random}(0,1) < e^{(-\Delta\text{objective}/\text{temperature})}$ **then**
      Accept W_{new} as W_{current}
    end if
  end if
  temperature $\leftarrow$ temperature $\cdot$ cooling_rate
end for
return W_{best}

of the results was conducted through statistical analysis, highlighting the method's performance, comparing it to the ALS method in terms of polarization (R_{pol}), individual unfairness (R_{indv}), and Root Mean Square Error ($RMSE$). The distributions of known (X) and estimated ratings ($\hat{X}$ and $\hat{X}_\pi$) were examined, along with the method's impact on the variance of recommendations per item and variability of individual errors.

4.1 Datasets

To evaluate our proposed method, we conducted experiments on two real-world datasets, whose statistics can be found in Table 1.

Dataset 1 used the MovieLens-1M dataset, which contains approximately 1 million ratings of around 4000 movies made by roughly 6000 users, with ratings on a 5-point scale [27]. We filtered the top 1000 users with the most ratings, along with the top 1000 most-rated movies.

In Dataset 2, we used the GoodBooks-10k dataset, which contains approximately 1 million ratings of about 10000 books made by 53424 users. To ensure a similar sparsity (ρ), we filtered the top 1000 users with the most ratings, along with the top 1000 most-rated books.

We randomly split the filtered datasets into two main parts: one for training (80%) and another for testing (20%).

Table 1: Description of datasets

Dataset	# Users	# Items	# Ratings	ρ(%)
MovieLens-1M	1000	1000	360784	63.92
GoodBooks-10k	1000	1000	68943	93.11

4.2 Traditional Recommendation Algorithm

For estimating unknown ratings we used the **ALS** (Alternating Least Squares) optimization method. This model-based method is employed in matrix factorization, minimizing the quadratic error, as demonstrated by [29]. In our experiments, we optimized two hyperparameters to ensure a balance between accuracy and computational efficiency: the *rank*, which was set to 20 to define the number of latent factors, and the regularization parameter *lambda*, also set to 20, to reduce the risk of overfitting.

4.3 Simulated Annealing Hyperparameters

For the Simulated Annealing meta-heuristic employed within the FaDeRS method, the following hyperparameters were determined and set after exploratory testing to optimize performance and computational efficiency:

- **Number of Iterations (`num_iterations`):** Set to **10,000** iterations. This value was chosen to ensure adequate exploration of the search space, balancing convergence speed with solution quality.
- **Initial Temperature (`initial_temp`):** Adjusted to **1,000**. A high initial temperature allows the algorithm greater freedom to accept worse solutions in the early stages, thereby facilitating escape from local optima and promoting a broader exploration.
- **Cooling Rate (`cooling_rate`):** Defined as **0.995**. This gradual cooling rate ensures that the exploration of the solution space is sufficiently slow, allowing the algorithm enough time to converge towards a high-quality solution without premature stagnation.

These specific values represent a practical balance between the quality of the solution found and the computational cost observed during our experiments.

4.4 Number of Estimated Matrices (h)

To ensure the accuracy and reliability of the results and determine the optimal number of estimated matrices to be calculated by the fairness algorithm, we tested ten possible values of h: 4, 8, 12, 16, 20, 24, 28, 32, 36, and 40. We performed 5 repetitions for each value of h. This approach allows for the measurement of the mean and standard deviation, providing a reliable estimate of data dispersion and facilitating the detection of outliers.

5 RESULTS AND DISCUSSION

In this section, we present the performance of the proposed method, emphasizing the quality of recommendations and the effectiveness of fairness compared to a traditional recommendation method (ALS). In the results presented in tables, the cells with the largest reductions in item polarization and individual unfairness are highlighted in red. Results with the largest increases in *RMSE* are highlighted in blue.

The summary data from the experiments are presented in Tables 2 and 3, providing the following information:

- α: Adjustment factor for the fraction of data deviation towards the mean;
- μ: Average execution of the algorithm over 5 repetitions for all values of h;
- %: Percentage of reduction or increase comparing the original mean (ALS) and the mean after running FaDeRS.

Table 2: ALS *vs* FaDeRS on MovieLens-1M

ALS | Polarization (R_{pol}) = 0.127 ALS | Individual Unfairness (R_{indv}) = 0.034 ALS | Root Mean Square Error ($RMSE$) = 0.875						
	R_{pol}		R_{indv}		$RMSE$	
α	(μ)	(%)	(μ)	(%)	(μ)	(%)
0.2	0.092	-27.20	0.033	-3.09	0.879	+0.48
0.4	0.065	-48.72	0.030	-13.06	0.892	+1.89
0.6	0.045	-64.73	0.026	-24.73	0.910	+3.95
0.8	0.027	-78.64	0.023	-31.17	0.935	+6.89

Table 3: ALS *vs* FaDeRS on GoodBooks-10K

ALS | Polarization (R_{pol}) = 0.254 ALS | Individual Unfairness (R_{indv}) = 0.058 ALS | Root Mean Square Error ($RMSE$) = 0.855						
	R_{pol}		R_{indv}		$RMSE$	
α	(μ)	(%)	(μ)	(%)	(μ)	(%)
0.2	0.177	-30.11	0.056	-4.26	0.860	+0.59
0.4	0.128	-49.54	0.051	-11.73	0.873	+2.03
0.6	0.083	-67.21	0.044	-23.57	0.895	+4.60
0.8	0.056	-77.95	0.038	-33.97	0.920	+7.53

The tables provide a detailed analysis of the fairness and efficiency metrics of the method compared to the ALS method on the MovieLens-1M and GoodBooks-10K datasets.

In the MovieLens-1M dataset (Table 2), polarization (R_{pol}) gradually reduces with the increase of α, reaching a maximum reduction of 78.64% with $\alpha = 0.8$. Individual unfairness (R_{indv}) also decreases, with the largest reduction at 31.17%. The Root Mean Square Error (*RMSE*) increases with α, reaching a maximum of 6.89%.

In the GoodBooks-10K dataset (Table 3), a similar trend is observed. Polarization (R_{pol}) and individual unfairness (R_{indv}) significantly decrease with $\alpha = 0.8$, with maximum reductions of 77.95% and 33.97%, respectively. The *RMSE* increases up to 7.53%.

The higher the adjustment factor (α), the lower the polarization (R_{pol}) and individual unfairness (R_{indv}). The reductions in polarization are proportionally greater since the adjustment factor (α) directly implies bringing the data towards the mean. The reduction in individual unfairness is subsequent and dependent on the optimization algorithm. Considering that the root mean square error (*RMSE*) increases proportionally with the adjustment factor (α), users of the method will need to decide between fairness (polarization and individual unfairness) and accuracy (root mean square error).

In Figure 2, a significant reduction in item polarization (R_{pol}) and individual unfairness (R_{indv}) is observed, with values presented in terms of percentage variation to more clearly illustrate the decrease in R_{pol} and R_{indv} relative to the increase in *RMSE*. The comparison was made between the polarization and individual unfairness of the estimated matrix $\hat{X}_\pi$, resulting from the application of the method (FaDeRS), and the estimated matrix $\hat{X}$, obtained through a traditional recommendation algorithm (ALS). Both analyses were conducted with the datasets MovieLens-1M and GoodBooks-10k.

Figure 3 illustrates the impact of applying the fairness algorithm on the estimated ratings for the MovieLens-1M and GoodBooks-10k datasets. The distribution of known ratings represents the baseline scenario of classifications (X), which frequently exhibit polarization. In these experiments, two distinct scenarios were investigated: the first with $\alpha = 0.2$, aiming for minimal perturbations in the initial estimated matrix ($\hat{X}$) to generate the final recommendation matrix ($\hat{X}_\pi$), and the second with $\alpha = 0.8$, where perturbations are greater. In the scenario with $\alpha = 0.8$, the greater perturbations result in estimates closer to the mean, indicating a reduction in the distribution range of estimates, and hence a reduction in polarization.

Figure 4 demonstrates more directly the effect of the fairness algorithm on polarization. The distribution presented refers to the variances of the estimated ratings of the items. R_{pol} is calculated as the mean of the variances of the item estimates. R_{pol} (before) considers the matrix $\hat{X}$, while R_{pol} (after) considers the matrix $\hat{X}_\pi$. After applying the algorithm, a reduction in the distribution range of variances is observed, resulting in an overall decrease in polarization.

Figure 5 presents the effects of the method focusing on individual unfairness from the item perspective. This metric evaluates the variability of the individual errors of the estimates of each item compared to the original ratings. In other words, R_{indv} measures the variance between the error estimates of an item. On the MovieLens-1M dataset, the application of the algorithm resulted in a reduction of R_{indv} from 0.03 to 0.02. In the GoodBooks-10k dataset, the reduction was from 0.06 to 0.04. A decrease in the value of R_{indv} after applying the algorithm indicates a more balanced distribution of estimation errors for each item, highlighting the algorithm's effectiveness in promoting lower variability of errors and, thus, greater fairness in the distribution of recommendations for the items.

These figures highlight that the method reduces polarization and promotes a fairer distribution of individual losses. This dual effect is essential for building recommender systems with broader fairness.

Our study references the work "Fighting Fire with Fire: Using Antidote Data to Improve Polarization and Fairness of Recommender Systems" by [39]. For comparative purposes, we evaluated FaDeRS against this benchmark on the **MovieLens-1M dataset**, where FaDeRS demonstrated substantial improvements, reducing R_{pol} by 78.74%, compared to the benchmark's 47.27% [39]. Regarding R_{indv},

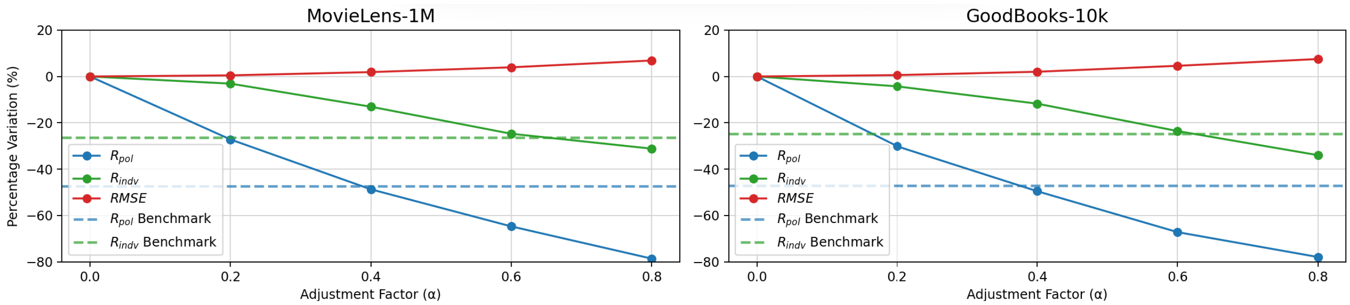

Figure 2: Polarization R_{pol}, Individual Unfairness R_{indv}, and $RMSE$

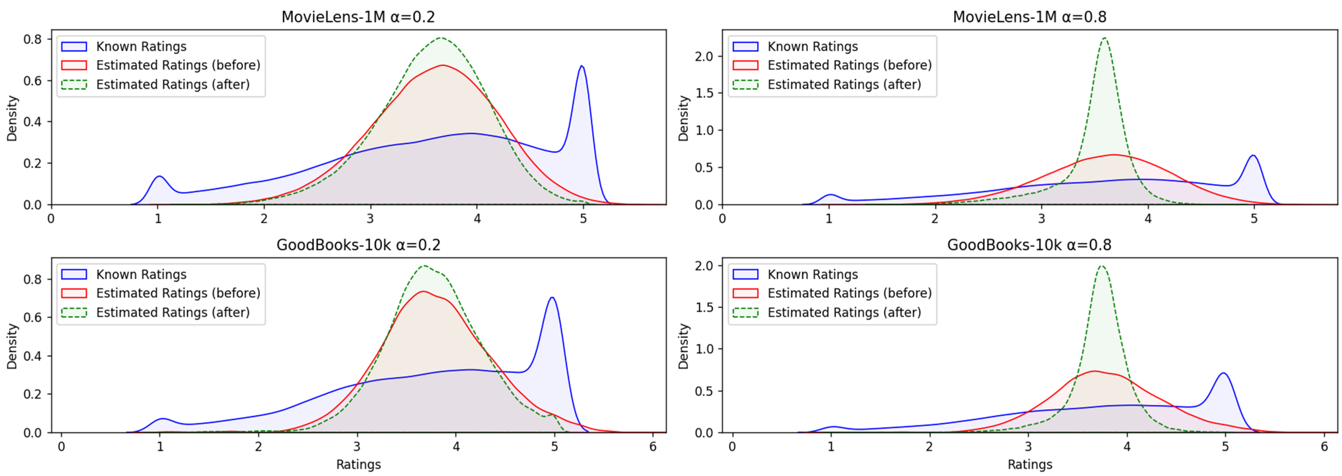

Figure 3: Known Ratings (X), Estimated Ratings Before ($\hat{X}$) and Estimated Ratings After ($\hat{X}_\pi$)

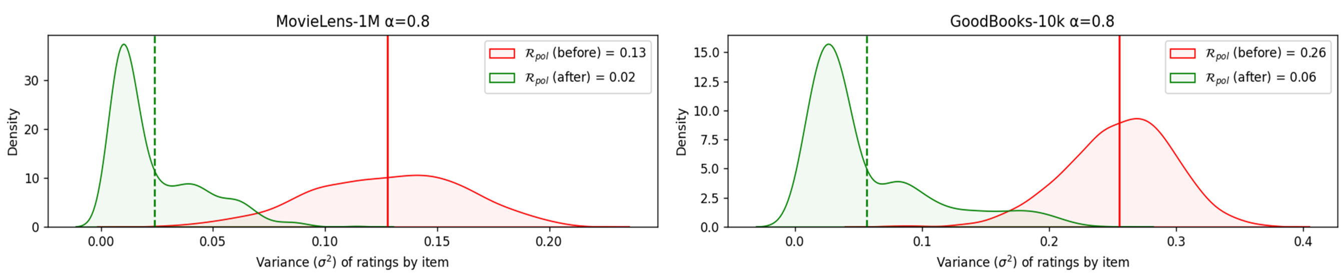

Figure 4: FaDeRS Effect on Polarization (R_{pol})

FaDeRS reduced it by 31.17%, surpassing the benchmark's 26.38%. These results underscore the effectiveness of FaDeRS in optimizing reduction metrics in the MovieLens-1M dataset.

To provide a broader comparative context, the methodology described by [39] was also implemented and executed using the **GoodBooks-10k dataset**. This execution yielded reductions of 47.16% in polarization and 24.70% in individual unfairness. In comparison, FaDeRS achieved more significant reductions on GoodBooks-10k, with 77.95% for polarization and 33.97% for individual unfairness. These comparative results across distinct datasets highlight the enhanced and robust effectiveness of FaDeRS.

6 CONCLUSIONS AND FUTURE WORK

The results highlight the importance of considering fairness measures, such as polarization (R_{pol}) and individual unfairness (R_{indv}), in recommendation algorithms. The proposed method reduced these metrics, promoting a more diverse and fair distribution of recommendations. As shown in the results section, polarization and individual unfairness were significantly reduced in the two tested datasets, with maximum reductions of 78.64% and 33.97%, respectively. This positive impact on the fairness of recommendations demonstrates the effectiveness of FaDeRS in decreasing

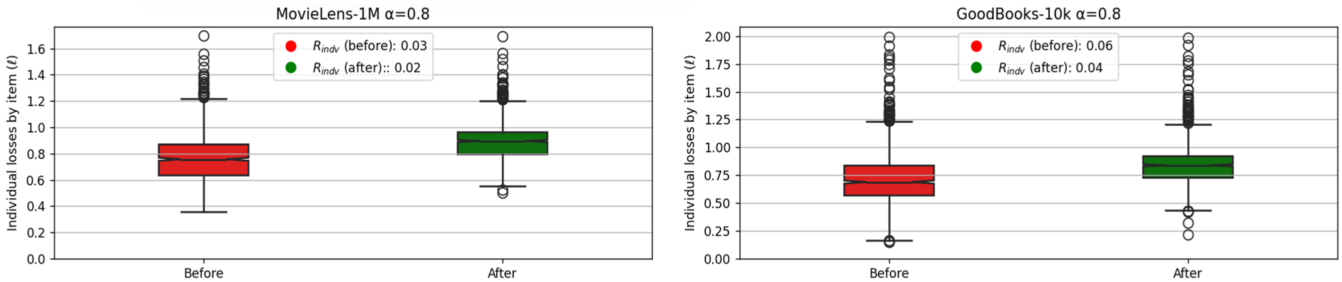

Figure 5: FaDeRS Effect on Individual Unfairness (R_{indv})

variances, both among the individual losses of items and in the estimated ratings, addressing the growing demand for fairer practices in recommendation systems.

The increase in the adjustment factor α improves fairness but degrades precision. With $\alpha = 0.8$, we achieved the greatest reductions in R_{pol} and R_{indv} in the MovieLens-1M and GoodBooks-10K datasets, but with a greater degradation in *RMSE*. The choice of the adjustment factor should be weighed according to the priorities of the application. It is important to emphasize that as α increases, the increment in *RMSE* becomes evident, suggesting the critical need to balance the trade-offs between precision and fairness to meet the specific needs of users.

The limitations of this work include the dependence on the factor α and the evaluation on only two datasets (MovieLens-1M and GoodBooks-10K), which may not generalize to other contexts. Furthermore, our approach relies on manual configurations for optimizing the adjustment factor, which can be a bottleneck in more dynamic or scalable scenarios.

Future research may focus on hybrid optimizations that reduce the loss of precision while promoting fairness. Exploring other fairness measures in different contexts, testing the method on various datasets, and applying it to other recommendation algorithms are important steps. It is also essential to develop techniques to dynamically adjust the factor α according to the characteristics of the dataset and the needs of the application. Additionally, gaining a better understanding of the longitudinal impact of these fairness measures in real-world environments could provide valuable insights for the continuous evolution and adaptation of more inclusive and fair recommendation systems.

Based on the positive performance in relation to the benchmarks, as detailed in the results section, the FaDeRS method has considerable potential for continuous innovation and practical implementation in recommendation systems that seek to balance effectiveness and ethics.

ACKNOWLEDGMENTS

We would like to thank the Federal Institute of Espírito Santo (IFES) for the institutional support provided during the development of this research.

REFERENCES

[1] Himan Abdollahpouri, Robin Burke, and Bamshad Mobasher. 2019. Managing Popularity Bias in Recommender Systems with Personalized Re-ranking. *CoRR* abs/1901.07555 (2019), 1–6. arXiv:1901.07555 http://arxiv.org/abs/1901.07555

[2] André Amaral, André Araújo, Bruno Tomazela, Tatiana Escovedo, and Marcos Kalinowski. 2023. Machine Learning Applied to the Classification of Technical Inspection Recommendations Regarding the Trend to Increase Criticality. In *Proceedings of the XIX Brazilian Symposium on Information Systems* (Maceió, Brazil) *(SBSI '23)*. Association for Computing Machinery, New York, NY, USA, 17–27. doi:10.1145/3592813.3592884

[3] Alexander Beutel, Jilin Chen, Zhe Zhao, and Ed H. Chi. 2019. Fairness in Recommendation Ranking through Pairwise Comparisons. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD '19)*. ACM, Anchorage, Alaska, USA, 2212–2220.

[4] A. Bhargava, A. Bansal, and V. Goyal. 2022. Machine Learning–Based Detection and Sorting of Multiple Vegetables and Fruits. *Food Analytical Methods* 15 (2022), 228–242. doi:10.1007/s12161-021-02086-1

[5] Robin Burke, Nasim Sonboli, and Aldo Ordonez-Gauger. 2018. Balanced Neighborhoods for Multi-sided Fairness in Recommendation. In *Proceedings of the 1st Conference on Fairness, Accountability and Transparency (Proceedings of Machine Learning Research, Vol. 81)*, Sorelle A. Friedler and Christo Wilson (Eds.). PMLR, New York University, NYC, USA, 202–214. https://proceedings.mlr.press/v81/burke18a.html

[6] Robin Burke, Nasim Sonboli, and Aldo Ordonez-Gauger. 2018. Fairness and Transparency in Recommendation Systems. In *Proceedings of the 2018 Conference on Fairness, Accountability, and Transparency (FAT'18)*. ACM, New York, NY, USA, 1–48.

[7] Simon Caton and Christian Haas. 2024. Fairness in Machine Learning: A Survey. *ACM Comput. Surv.* 56, 7, Article 166 (apr 2024), 38 pages. doi:10.1145/3616865

[8] L. Elisa Celis, Amit Deshpande, Tarun Kathuria, and Nisheeth K. Vishnoi. 2016. How to be Fair and Diverse? *CoRR* abs/1610.07183 (2016), 5. arXiv:1610.07183 http://arxiv.org/abs/1610.07183

[9] Sanjay Chandrasekaran, Giulia De Pasquale, Giuseppe Belgioioso, and Florian Dörfler. 2024. Mitigating Polarization in Recommender Systems via Network-aware Feedback Optimization. arXiv:2408.16899 [eess.SY] https://arxiv.org/abs/2408.16899

[10] Federico Cinus, Marco Minici, Corrado Monti, and Francesco Bonchi. 2022. The Effect of People Recommenders on Echo Chambers and Polarization. *Proceedings of the International AAAI Conference on Web and Social Media* 16, 1 (May 2022), 90–101. doi:10.1609/icwsm.v16i1.19275

[11] Miriam Colpo, Bruno Alves, Kevin Pereira, Anna Flávia Brandão, Marilton Aguiar, and Tiago Primo. 2021. Attribute selection based on genetic and classification algorithms in the prediction of hospitalization need of COVID-19 patients. In *Anais do XVII Simpósio Brasileiro de Sistemas de Informação* (Uberlândia). SBC, Porto Alegre, RS, Brasil, 8 pages. https://sol.sbc.org.br/index.php/sbsi/article/view/17691

[12] Jadna Cruz, Amanda Oliveira, Diego Silva, and Frederico Durão. 2022. GRSPOI: A Point-of-Interest Recommender Systems for Groups Using Diversification. In *Anais do XVIII Simpósio Brasileiro de Sistemas de Informação* (Curitiba). SBC, Porto Alegre, RS, Brasil, 8 pages. https://sol.sbc.org.br/index.php/sbsi/article/view/21349

[13] S. Dargan, M. Kumar, and A. Mittal et al. 2024. Handwriting-based gender classification using machine learning techniques. *Multimedia Tools and Applications* 83 (2024), 19871–19895. doi:10.1007/s11042-023-16354-1

[14] Joselito de Sousa Júnior, Vinícius Machado, Rodrigo Souza Veras, Roney de Sales Santos, Bruno Alves de Lima, Aline Leal Silva, Francisco da Silva Sousa, and Francisco Imperes Filho. 2022. Public Administration Suppliers Classification Model Based on Supervised Machine Learning. In *Anais do XVIII Simpósio Brasileiro de Sistemas de Informação* (Curitiba). SBC, Porto Alegre, RS, Brasil, 8 pages. https://sol.sbc.org.br/index.php/sbsi/article/view/21382

[15] Yashar Deldjoo, Vito Walter Anelli, Hamed Zamani, Alejandro Bellogín, and T. D. Noia. 2021. A flexible framework for evaluating user and item fairness in recommender systems. *User Modeling and User-Adapted Interaction* 31 (2021), 457 – 511. https://api.semanticscholar.org/CorpusID:231591295

[16] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Rich Zemel. 2011. Fairness Through Awareness. arXiv:1104.3913 [cs.CC] https://arxiv.org/abs/1104.3913

[17] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard S. Zemel. 2011. Fairness Through Awareness. *CoRR* abs/1104.3913 (2011), 214–226. arXiv:1104.3913 http://arxiv.org/abs/1104.3913

[18] 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. doi:10.1561/1500000079

[19] Michael D. Ekstrand, Maria Tian, Juyeon Maeng Azpiazu, Robin Burke, John D. Ekstrand, and Augustine Anuyah. 2018. All the Cool Kids, How Do They Fit In?: Popularity and Demographic Biases in Recommender Evaluation and Effectiveness. In *Proceedings of the 2018 Conference on Fairness, Accountability, and Transparency*. PMLR, New York, NY, USA, 172–186.

[20] G. Emirhüseyinoğlu, M. Shahhosseini, G. Hu, et al. 2023. Validation of scenario generation for decision-making using machine learning prediction models. *Optim Lett* 17 (2023), 520–528. doi:10.1007/s11590-023-02023-7

[21] Felipe Faria, Álvaro Jr., and Luiz Merschmann. 2015. Predição de Rankings de Artistas por Meio de Regressão. In *Anais do XI Simpósio Brasileiro de Sistemas de Informação* (Goiânia). SBC, Porto Alegre, RS, Brasil, 95–102. doi:10.5753/sbsi.2015.5805

[22] Golnoosh Farnadi, Pigi Kouki, Spencer K. Thompson, Sriram Srinivasan, and Lise Getoor. 2018. A Fairness-aware Hybrid Recommender System. arXiv:1809.09030 [cs.IR] https://arxiv.org/abs/1809.09030

[23] X. Feng, Y. Cai, and R. Xin. 2023. Optimizing diabetes classification with a machine learning-based framework. *BMC Bioinformatics* 24 (2023), 428. doi:10.1186/s12859-023-05467-x

[24] Sorelle A. Friedler, Carlos Scheidegger, and Suresh Venkatasubramanian. 2016. On the (im)possibility of fairness. arXiv:1609.07236 [cs.CY] https://arxiv.org/abs/1609.07236

[25] Avijit Ghosh, Ritam Dutt, and Christo Wilson. 2021. When Fair Ranking Meets Uncertain Inference. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '21)*. ACM, Virtual Event, Canada, 1033–1043. doi:10.1145/3404835.3462850

[26] Wesckley Faria Gomes and Methanias Colaço. 2022. Applications of Artificial Intelligence for Auditing and Classification of Incongruent Descriptions in Public Procurement. In *Proceedings of the XVIII Brazilian Symposium on Information Systems* (Curitiba, Brazil) *(SBSI '22)*. Association for Computing Machinery, New York, NY, USA, Article 40, 8 pages. doi:10.1145/3535511.3535551

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

[28] Toshihiro Kamishima, Shotaro Akaho, Hideki Asoh, and Jun Sakuma. 2018. Recommendation Independence. In *Proceedings of the 1st Conference on Fairness, Accountability and Transparency (Proceedings of Machine Learning Research, Vol. 81)*, Sorelle A. Friedler and Christo Wilson (Eds.). PMLR, New York, NY, USA, 187–201. https://proceedings.mlr.press/v81/kamishima18a.html

[29] Ni Li and Yinshui Xia. 2024. Movie recommendation based on ALS collaborative filtering recommendation algorithm with deep learning model. *Entertainment Computing* 51 (2024), 100715. doi:10.1016/j.entcom.2024.100715

[30] Rodrigo Lima, Giovanni Comarella, Fabiano Belém, and Julio Reis. 2022. Justiça em Sistemas de Recomendação: Uma Análise de Técnicas de Regularização. In *Anais do XXXVII Simpósio Brasileiro de Bancos de Dados* (Búzios). SBC, Porto Alegre, RS, Brasil, 177–189. doi:10.5753/sbbd.2022.224352

[31] Nícolas Terra Maia and Tiago Cruz de França. 2023. Recomendação de Conteúdo Baseada em Localização e Interesse: uma aplicação com o CidadeSocial. In *Anais Estendidos do XXXVIII Simpósio Brasileiro de Bancos de Dados* (Belo Horizonte/MG). SBC, Porto Alegre, RS, Brasil, 71–77. doi:10.5753/sbbd_estendido.2023.233535

[32] Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. 2022. A Survey on Bias and Fairness in Machine Learning. arXiv:1908.09635 [cs.LG] https://arxiv.org/abs/1908.09635

[33] Rishabh Mehrotra, James McInerney, Hugues Bouchard, Mounia Lalmas, and Fernando Diaz. 2018. Towards a Fair Marketplace: Counterfactual Evaluation of the Trade-off between Relevance, Fairness & Satisfaction in Recommendation Systems. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*. ACM, Singapore, 2243–2252. doi:10.1145/3269206.3272027

[34] Juliano Menzen and Kleinner Oliveira. 2023. 4Experts: A Task Recommendation Approach Using Machine Learning and Biometric Data. In *Anais do XIX Simpósio Brasileiro de Sistemas de Informação* (Maceió/AL). SBC, Porto Alegre, RS, Brasil, 1–8. https://sol.sbc.org.br/index.php/sbsi/article/view/25108

[35] Debarghya Mukherjee, Mikhail Yurochkin, Moulinath Banerjee, and Yuekai Sun. 2020. Two simple ways to learn individual fairness metrics from data. In *Proceedings of the 37th International Conference on Machine Learning (ICML'20)*. JMLR.org, Virtual, Article 658, 11 pages.

[36] Tien T. Nguyen, Philip M. Hui, F. Maxwell Harper, Loren Terveen, and Joseph A. Konstan. 2014. Exploring the Filter Bubble: The Effect of Using Recommender Systems on Content Diversity. In *Proceedings of the 23rd International Conference on World Wide Web*. Association for Computing Machinery, Seoul, South Korea, 677–686. doi:10.1145/2566486.2568012

[37] Eli Pariser. 2011. *The Filter Bubble: How the New Personalized Web Is Changing What We Read and How We Think.* Penguin Press, New York, NY.

[38] Bruno Prece, Edson Pacheco, Rodolfo Barros, and Sylvio Barbon. 2019. Improvements on diagnostic assessment questionnaires of Maturity Level Management with feature selection. In *Anais do XV Simpósio Brasileiro de Sistemas de Informação* (Aracajú). SBC, Porto Alegre, RS, Brasil, 48–55. https://sol.sbc.org.br/index.php/sbsi/article/view/13888

[39] Bashir Rastegarpanah, Krishna P. Gummadi, and Mark Crovella. 2019. Fighting Fire with Fire: Using Antidote Data to Improve Polarization and Fairness of Recommender Systems. In *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining* (Melbourne VIC, Australia) *(WSDM '19)*. Association for Computing Machinery, New York, NY, USA, 231–239. doi:10.1145/3289600.3291002

[40] Faria Sajjad, Muhammad Rashid, Afia Zafar, Kainat Zafar, Benish Fida, Ali Arshad, Saman Riaz, Ashit Kumar Dutta, and Joel J.P.C. Rodrigues. 2023. An efficient hybrid approach for optimization using simulated annealing and grasshopper algorithm for IoT applications. *Discover Internet of Things* 3, 1 (Dec. 2023), 19 pages. doi:10.1007/s43926-023-00036-3 Publisher Copyright: © The Author(s) 2023..

[41] Everton Santana, Saulo Mastelini, and Sylvio Jr. 2017. Deep Regressor Stacking for Air Ticket Prices Prediction. In *Anais do XIII Simpósio Brasileiro de Sistemas de Informação* (Lavras). SBC, Porto Alegre, RS, Brasil, 25–31. doi:10.5753/sbsi.2017.6022

[42] Paulo Santos, Karina Delgado, Marcelo Lauretto, and Marcio Ribeiro. 2017. Revisão Sistemática da Literatura sobre ranking de Relacionamentos na Web Semântica. In *Anais do XIII Simpósio Brasileiro de Sistemas de Informação* (Lavras). SBC, Porto Alegre, RS, Brasil, 32–39. doi:10.5753/sbsi.2017.6023

[43] A. Sikora, T. Zhang, and D.J. Murphy et al. 2023. Machine learning vs. traditional regression analysis for fluid overload prediction in the ICU. *Scientific Reports* 13 (2023), 19654. doi:10.1038/s41598-023-46735-3

[44] Adriano Souza, Reinaldo Fortes, and Joubert Lima. 2017. OLAP Textual com Múltiplas Hierarquias de Tópicos e Rankings Segmentados. In *Anais do XIII Simpósio Brasileiro de Sistemas de Informação* (Lavras). SBC, Porto Alegre, RS, Brasil, 480–487. doi:10.5753/sbsi.2017.6078

[45] Ediana Souza, Daniel Lichtnow, and Isabela Gasparini. 2022. Estratégia de Pós-processamento Aplicada a um Sistema de Recomendação de Artigos para a Melhora da Diversidade. In *BraSNAM 2022 - XI Brazilian Workshop on Social Network Analysis and Mining* (Niterói). Sociedade Brasileira de Computação, Sociedade Brasileira de Computação, Porto Alegre, RS, Brasil, 216–221. doi:10.5753/brasnam.2022.222805

[46] Luan De Souza and Marcelo Manzato. 2023. A User Study on Explanations with Different Levels of Detail in Recommender Systems. In *Anais do XXIX Simpósio Brasileiro de Sistemas Multimídia e Web* (Ribeirão Preto/SP). SBC, Porto Alegre, RS, Brasil, 208–212. doi:10.1145/3617023.3617051

[47] Dominik Spohr. 2017. Fake News and Ideological Polarization: Filter Bubbles and Selective Exposure on Social Media. *Business Information Review* 34, 3 (2017), 150–160. doi:10.1177/0266382117722446

[48] Harald Steck. 2018. Calibrated recommendations. In *Proceedings of the 12th ACM Conference on Recommender Systems* (Vancouver, British Columbia, Canada) *(RecSys '18)*. Association for Computing Machinery, New York, NY, USA, 154–162. doi:10.1145/3240323.3240372

[49] Jonathan Stray. 2021. Designing Recommender Systems to Depolarize. arXiv:2107.04953 [cs.IR] https://arxiv.org/abs/2107.04953

[50] V. A. Sudakov and Y. P. Titov. 2023. Pandemic Forecasting by Machine Learning in a Decision Support Problem. *Math Models Comput Simul* 15 (2023), 520–528. doi:10.1134/S2070048223030171

[51] Jiakai Tang, Shiqi Shen, Zhipeng Wang, Zhi Gong, Jingsen Zhang, and Xu Chen. 2023. When Fairness meets Bias: a Debiased Framework for Fairness aware Top-N Recommendation. In *Proceedings of the 17th ACM Conference on Recommender Systems* (Singapore, Singapore) *(RecSys '23)*. Association for Computing Machinery, New York, NY, USA, 200–210. doi:10.1145/3604915.3608770

[52] Antonio Vale, Pedro Santos, Hélcio Soares, and Raimundo Moura. 2023. Automatic Classification of Public Expenses in the Fight against COVID-19: A Case Study of TCE/PI. In *Anais do XIX Simpósio Brasileiro de Sistemas de Informação* (Maceió/AL). SBC, Porto Alegre, RS, Brasil, 221–228. https://sol.sbc.org.br/index.php/sbsi/article/view/25135

[53] R. Vashistha, Z. Noor, and S. Dasgupta et al. 2023. Application of statistical machine learning in biomarker selection. *Scientific Reports* 13 (2023), 18331. doi:10.1038/s41598-023-45323-9

[54] Yifan Wang, Weizhi Ma, Min Zhang, Yiqun Liu, and Shaoping Ma. 2023. A Survey on the Fairness of Recommender Systems. *ACM Trans. Inf. Syst.* 41, 3, Article 52 (feb 2023), 43 pages. doi:10.1145/3547333

[55] Heitor Werneck, Nicollas Silva, Carlos Mito, Adriano Pereira, Diego Dias, Elisa Albergaria, and Leonardo Rocha. 2022. Introducing Contextual Information in an Ensemble Recommendation System for Fashion Domains. In *Anais do XXVIII Simpósio Brasileiro de Sistemas Multimídia e Web* (Curitiba). SBC, Porto Alegre,

RS, Brasil, 237–244. doi:10.1145/3539637.3557058

[56] Chuhan Wu, Fangzhao Wu, Xiting Wang, Yongfeng Huang, and Xing Xie. 2021. FairRec: Fairness-aware News Recommendation with Decomposed Adversarial Learning. arXiv:2006.16742 [cs.IR] https://arxiv.org/abs/2006.16742

[57] Yao Wu, Jian Cao, Guandong Xu, and Yudong Tan. 2021. TFROM: A Two-sided Fairness-Aware Recommendation Model for Both Customers and Providers. arXiv:2104.09024 [cs.AI] https://arxiv.org/abs/2104.09024

[58] Bin Xia, Junjie Yin, Jian Xu, and Yun Li. 2019. WE-Rec: A fairness-aware reciprocal recommendation based on Walrasian equilibrium. *Knowledge-Based Systems* 182 (2019), 104857. doi:10.1016/j.knosys.2019.07.028

[59] J. Xue, W. Yao, and S. Xiang. 2024. Machine learning embedded EM algorithms for semiparametric mixture regression models. *Computational Statistics* 22 (2024), 159–171. doi:10.1007/s00180-024-01482-5

[60] Muhammad Bilal Zafar, Isabel Valera, Manuel Gomez Rodriguez, and Krishna P. Gummadi. 2017. Fairness Beyond Disparate Treatment & Disparate Impact: Learning Classification without Disparate Mistreatment. In *Proceedings of the 26th International Conference on World Wide Web* (Perth, Australia) *(WWW '17)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 1171–1180. doi:10.1145/3038912.3052660

[61] Meike Zehlike, Tom Sühr, Ricardo Baeza-Yates, Francesco Bonchi, Carlos Castillo, and Sara Hajian. 2022. Fair Top-k Ranking with multiple protected groups. *Processing & Management* 59, 1 (2022), 102707. doi:10.1016/j.ipm.2021.102707

[62] Wenbin Zhang, Zichong Wang, Juyong Kim, Cheng Cheng, Thomas Oommen, Pradeep Ravikumar, and Jeremy Weiss. 2023. Individual Fairness under Uncertainty. arXiv:2302.08015 [cs.LG] https://arxiv.org/abs/2302.08015

[63] Wenbin Zhang and Jeremy C. Weiss. 2022. Longitudinal Fairness with Censorship. arXiv:2203.16024 [cs.LG] https://arxiv.org/abs/2203.16024

[64] Yong Zheng, Tanaya Dave, Neha Mishra, and Harshit Kumar. 2018. Fairness In Reciprocal Recommendations: A Speed-Dating Study. In *Adjunct Publication of the 26th Conference on User Modeling, Adaptation and Personalization* (Singapore, Singapore) *(UMAP '18)*. Association for Computing Machinery, New York, NY, USA, 29–34. doi:10.1145/3213586.3226207

[65] Ziwei Zhu, Xia Hu, and James Caverlee. 2018. Fairness-Aware Tensor-Based Recommendation. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management* (Torino, Italy) *(CIKM '18)*. Association for Computing Machinery, New York, NY, USA, 1153–1162. doi:10.1145/3269206.3271795