



PDF Download
3708319.3733710.pdf
19 December 2025
Total Citations: 0
Total Downloads: 534

 Latest updates: <https://dl.acm.org/doi/10.1145/3708319.3733710>

RESEARCH-ARTICLE

User and Recommender Behavior Over Time: Contextualizing Activity Effectiveness Diversity and Fairness in Book Recommendation

[SAMIRA VAEZ BARENJI](#), Drexel University, Philadelphia, PA, United States

[SUSHOBHAN PARAJULI](#), Drexel University, Philadelphia, PA, United States

[MICHAEL D. EKSTRAND](#), Drexel University, Philadelphia, PA, United States

Open Access Support provided by:

Drexel University

Published: 16 June 2025

[Citation in BibTeX format](#)

UMAP '25: 33rd ACM Conference
on User Modeling, Adaptation and
Personalization

June 16 - 19, 2025
New York City, USA

Conference Sponsors:
SIGCHI

User and Recommender Behavior Over Time

Contextualizing Activity Effectiveness Diversity and Fairness in Book Recommendation

Samira Vaez Barenji

Drexel University

Department of Information Science

Philadelphia, PA, USA

sv849@drexel.edu

Sushobhan Parajuli

Drexel University

Department of Information Science

Philadelphia, PA, USA

sp3886@drexel.edu

Michael D. Ekstrand

Drexel University

Department of Information Science

Philadelphia, PA, USA

mdekstrand@drexel.edu

Abstract

Data is an essential resource for studying recommender systems. While there has been significant work on improving and evaluating state-of-the-art models and measuring various properties of recommender system outputs, less attention has been given to the data itself, particularly how data has changed over time. Such documentation and analysis provide guidance and context for designing and evaluating recommender systems, particularly for evaluation designs making use of time (e.g., temporal splitting). In this paper, we present a temporal explanatory analysis of the UCSD Book Graph dataset scraped from Goodreads, a social reading and recommendation platform active since 2006. We measure the book interaction data using a set of activity, diversity, and fairness metrics; we then train a set of collaborative filtering algorithms on rolling training windows to observe how the same measures evolve over time in the recommendations. Additionally, we explore whether the introduction of algorithmic recommendations in 2011 was followed by observable changes in user or recommender system behavior.

Keywords

recommender systems, gender, authors, fairness, bias, diversity

ACM Reference Format:

Samira Vaez Barenji, Sushobhan Parajuli, and Michael D. Ekstrand. 2025. User and Recommender Behavior Over Time: Contextualizing Activity Effectiveness Diversity and Fairness in Book Recommendation. In *Adjunct Proceedings of the 33rd ACM Conference on User Modeling, Adaptation and Personalization (UMAP Adjunct '25)*, June 16–19, 2025, New York City, NY, USA. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3708319.3733710>

1 Introduction

Datasets are crucial to building, evaluating, and studying recommender systems and other information access systems (IAS). However, the documentation and descriptive analysis accompanying most datasets is slim, focused primarily on high-level descriptive statistics (e.g., the numbers of users, items, and interactions). As recommender systems research has paid increasing attention to time and sequence in training and evaluation data, both through techniques like sequential recommendation [4] and the growing consensus towards global temporal splitting for evaluating even traditional recommender systems [25], it is important to understand

not only the *static* characteristics of a dataset but also how those characteristics *change* over the history of interactions recorded in the dataset. Temporal understanding is also valuable for documenting and contextualizing changes in user or system behavior, both for traditional evaluation and to support research on evolving concerns regarding fairness, isolation, and other social impacts.

In this paper, we present a longitudinal explanatory analysis of the UCSD Book Graph [39, 40], a dataset containing book metadata, reviews, ratings, and interactions from the Goodreads social reading and recommendation platform, spanning more than 10 years from early 2007 to 2017. This analysis describes how the data has evolved over time in its volume as well as genre diversity and gender balance in user interactions, and the effectiveness, diversity, and fairness of collaborative filtering models trained and evaluated at different points in its history.

For the first several years, Goodreads' recommendations were purely social (direct, personal recommendations between users in the social graph). In 2011, Goodreads introduced algorithmic recommendations, generating recommendations for users after they rated at least 20 items on a five-star scale [1]. We further examine whether there are measurable changes in user or recommender behavior on the dimensions we are analyzing associated with the introduction of this recommender system.

Describing data, user, and model changes over time is important for several reasons: i) guiding the design and evaluation of experimental setups or understanding the impact of design decisions (e.g., the effect of selecting different time points for splitting); ii) revealing how different types of bias evolve over time, both to directly understand fairness implications of the data as well as to provide temporal context for further fairness investigations; iii) assessing whether model performance and fairness measures remain consistent across time. While some prior studies have conducted exploratory analyses of recommendation datasets in different domains [26, 37], there remains a gap in integrating longitudinal analyses of real-world recommendation datasets into fairness literature. This paper addresses this gap through the following contributions:

- (1) Conducting a longitudinal exploratory analysis of user activity on Goodreads to understand how activity, fairness, and diversity metrics of user interactions evolve over time.
- (2) Comparing four collaborative filtering algorithms trained on successive time windows to examine how their behavior changes over time in relation to the same set of concerns.
- (3) Investigation of whether the introduction of a recommender system on Goodreads is associated with measurable changes in the system's behavior



This work is licensed under a Creative Commons Attribution 4.0 International License. *UMAP Adjunct '25*, New York City, NY, USA
© 2025 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-1399-6/25/06
<https://doi.org/10.1145/3708319.3733710>

The temporal changes in data and evaluation metrics we document show that temporal recommender system evaluations are sensitive to the time window at which they are performed. Therefore, evaluating a system at a single point in time does not necessarily reflect its broader behavior. Some periods may reflect peak performance, while others reveal reduced effectiveness or increased disparities. Without a temporal perspective, evaluations risk overlooking how biases evolve as both the user base and content shift. Understanding system behavior over time is therefore essential for capturing its ability to meet users' needs and for identifying when and how bias emerges or is mitigated. It also supports more meaningful metric interpretation and the development of emerging best practices.

In the remainder of this paper, we survey related work on dataset description, bias, and longitudinal analysis (§2), detail the data, algorithms, and experimental setup (§3), present our results and observations (§4), and conclude with implications for recommender systems research on fairness and other topics (§6).

2 Background and Related Work

Describing Datasets. There is a growing call for detailed documentation of AI training and evaluation data, motivated by concerns about transparency, accountability, and the potential for unintended harms [18, 19, 28, 32]. Poorly documented datasets can obscure hidden biases, mislead users about appropriate applications, or lead to flawed evaluations [30, 31]. To address these issues, several documentation frameworks have been proposed. Gebru et al. [18] introduce *datasheets* to describe the dataset's context, content, and recommended uses, similar to documentation practices in engineering. Pushkarna et al. [32] propose *data cards*, which are short summaries that explain how the dataset was created and how it should be used. Nesvijejskaia [28] presents a standardized framework *databook* for dynamically documenting algorithm design throughout data science projects.

In the context of recommender systems, the widely-used MovieLens datasets have detailed historical documentation [21], but this documentation was only published after the platform had been operating for over 15 years. However, such documentation is rare for other datasets. Comprehensive data documentation has been identified as essential to enable reproducible results in AI research [20]. In this paper, we examine the UCSD Book Graph scraped from Goodreads. The dataset is briefly summarized on its website¹ and in the authors' papers [39, 40]. Ekstrand and Kluver [13] augmented the Book Graph with publicly available book and author metadata to study fairness towards authors' gender identities; they provided a brief description of the distributions of author gender data. However, changes to the dataset's size, composition, and distributions over time have not yet been documented.

Fairness and Social Impacts. User-generated data on social media platforms and review sites often reflect societal biases. These biases can enter datasets through uneven participation, visibility, or representation, and may be further amplified during data processing [30]. When such data are used to train recommender systems, it

can reinforce existing patterns of popularity or exclusion and lead to biased system behavior [2, 3, 43].

Research on fairness in recommender systems seeks to document, measure, and mitigate these biases [8, 12, 41]. Ekstrand and Kluver [13] used the Book Graph we examine in this paper to study author gender representation in user profiles and recommender system outputs. Although much of the recommender systems fairness research has focused on static evaluations [7], recent studies highlight the importance of understanding how fairness outcomes evolve as models are retrained on new user feedback [15, 16]. Much of this work, however, is prospective, examining how fairness changes as models are trained and retrained going forward from a static dataset. Our present study complements this direction of research by studying and evaluating how user interactions and model behavior evolve historically over time in the Book Graph.

Longitudinal Analysis. Recommender systems can have long-term effects on user behavior and content exposure [17]. Over time, they may reinforce existing patterns, leading to reduced diversity and unequal visibility across groups [14]. Some studies have shown that short-term engagement signals do not always align with long-term outcomes, like user satisfaction or return visits [42]. Longitudinal studies, which track user interactions over extended periods, are essential to understanding long-term effects. Many longitudinal studies of recommender systems are prospective and examine how systems and user behavior evolve going forward. Some, however, are historical, such as the filter bubble analysis of Nguyen et al. [29] and temporal evaluations of recommender system accuracy [5, 24]. In this study, we conduct a historical longitudinal analysis to understand how the underlying data in a real-world recommendation platform has changed over time.

3 Data and Methods

To document the evolution of book interaction data and its influence on collaborative filtering outputs, we measure the volume and diversity of user interactions across successive time windows, as well as the results of training and evaluating collaborative filtering models at different points in the dataset's temporal span. This section details the specific methods and data preparation used in our analysis.²

3.1 Data

The UCSD Book Graph [39, 40] is a dataset scraped in 2018 from Goodreads, a social reading and recommendation platform. It spans over a decade of book metadata, reviews, ratings, and interactions of users with public profiles, covering the period from January 2007 to November 2017. Ekstrand and Kluver [13] augmented the book graph with public author data from library sources and OpenLibrary through the PIReT/INERTIAL Book Data Tools³, providing further information such as author gender identities. The original book graph includes interactions of approximately 876K users with 1.52M million books. It spans various user activities and is organized into the following components:

- User actions, where a user adds a book to one of their "shelves";

¹<https://cseweb.ucsd.edu/~jmcauley/datasets/goodreads.html>

²The code is available at <https://zenodo.org/records/15333266>

³<https://bookdata.inertial.science>

Table 1: Summary statistics

	Overall	Male	Female
Unique Users	876,145	—	—
Unique Books	1,522,486	348,302	231,345
Unique Authors	612,241	155,888	87,904
Average Rating	3.85	3.80	3.79

- User ratings, provided on a 5-star scale;
- Book metadata, including title, genres, author names, and other attributes.

For our analysis, we count Goodreads *works* as items for user interaction and recommendation. A work is the fundamental unit of browsing and cataloging in Goodreads, and represents a particular literary work including its various editions, printings, translations, etc.; individual editions are represented as *books*. Using version 3.0 of the Book Data Tools, we aggregated and deduplicated these individual actions to form an interaction matrix representing each observed user-work pair with a single record. Each record includes the first and last timestamps of any interaction between the user and that work, whether through shelving or rating, along with the last rating the user provided for that work. For this study, we used the last interaction timestamp, as the first timestamps had some erroneous values that made them unreliable for analysis.

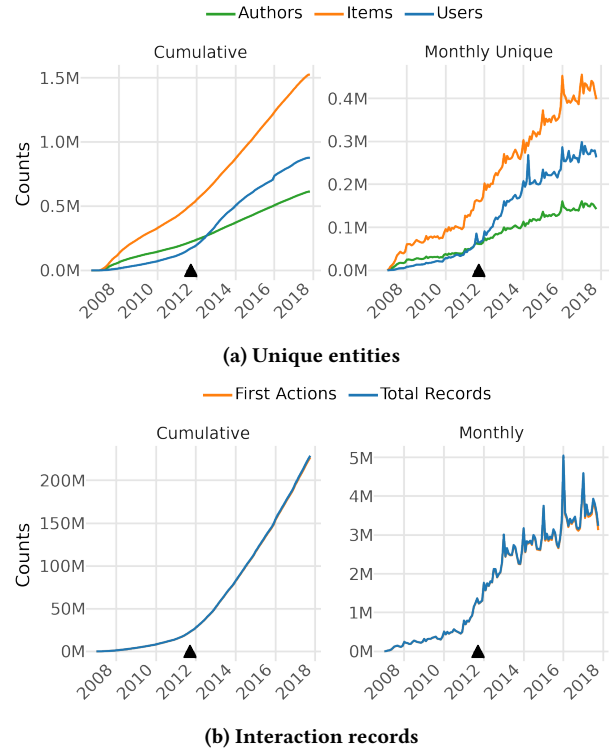
To support the analysis of genre diversity we used the book genre information supplied with the Book Graph (and extracted from users’ genre-related shelf names) to derive a distribution over genres for each work. The book graph genre records report the number of times each genre shelf has been applied to a book by users. We normalized these counts to obtain $P(g | b)$, the probability that book b belongs to genre g . The book graph also contains textual reviews, but we did not consider those in this analysis. Table 1 presents a statistical summary of the dataset, including the number of unique users, books, and authors, as well as the average rating in the ratings data. The gender-specific columns report these same statistics separately for male and female authors based on available gender data. For all experiments and analyses, we excluded authors with ambiguous or missing gender labels (using the same logic as Ekstrand and Kluver [13]).

3.2 Experimental Setup

We used a subset of actions data for our experiments that ranged from January 1, 2007, to October 31, 2017. We begin by analyzing the interaction data within this period on a monthly basis to examine its temporal trends.

To generate and evaluate recommendations, we created successive 2-month windows of test data beginning on Jan. 1, 2009, training the recommender system on the preceding 2 years of interactions and generating top-100 recommendation lists for each user in the test window who also has interactions in the training window. This setup allows us to analyze how model performance evolves over time.

We used LensKit [10] to generate and evaluate recommendations from collaborative filtering models trained on these successive training sets. Following our prior work using this dataset [13, 34],

**Figure 1: Unique entity and interaction counts per month.**

we used three implicit-feedback collaborative filtering algorithms: item-based k -NN [ItemKNN, 9], implicit-feedback ALS Matrix Factorization [ImplicitMF, 36], and Bayesian Personalized Ranking [BPR, 35]; we omit User-based k -NN due to its computational cost. We also considered non-personalized most-popular recommendations (MostPop). This selection helps ensure consistency and alignment with our previous studies, extending that line of inquiry through a longitudinal lens.

4 Results

In this section, we present the results of our analysis on both the interaction data and the recommendation outputs using standard evaluation metrics. We examine how structural and fairness-related properties of the data, such as genre diversity and author gender representation, evolved over time in the data, and how these same properties are reflected in the generated recommendations. This analysis also allows us to assess whether any measurable changes are associated with the introduction of Goodreads’ recommender system in 2011. In each figure, the recommender introduction is marked with a triangle on the x-axis.

4.1 Data

We first explore the temporal patterns in the user interaction data to understand how its structure, diversity, and fairness properties evolved over time, regardless of the recommender outputs. All measurements are on a monthly basis.




Figure 2: Average genre entropy per user.

Data size and growth over time. Figure 1 shows the monthly activity levels in the interaction data. Figure 1a shows the monthly volume of unique users, books, and authors in the data. All three entities show increasing activity over time, with the sharpest increases occurring after 2011. Figure 1b shows the number of records (user-book interactions) and the number of “first” interactions (the first time a user interacts with a particular book). Both lines almost completely overlap and steadily increase as expected, reflecting the platform’s growth over the ten-year period. While there are several fluctuations in monthly user activity, the highest peak occurs around 2016.

Genre diversity. To quantify genre diversity in the interaction data, we derive user-specific genre distributions by aggregating the genre distributions of the books each user has interacted with during each month. For a given user, the probability assigned to each genre reflects the cumulative probability mass of that genre across all books in their interaction history, based on the book-level distributions $P(g | b)$ defined in §3. We then compute the Shannon entropy of this user-level distribution to measure the genre diversity in their interactions, with higher entropy indicating a more diverse selection. The final metric is reported as the average entropy across all users.

Figure 2 shows the average genre entropy across user interactions. We can see that genre diversity stabilizes after an initial rise, holding relatively stable with some fluctuations until about 2012, after which it decreases. The patterns suggest that user interaction with genres remained relatively consistent over time.

Individual fairness and popularity bias. We use the *Gini index* to measure the distribution of interactions across individual items and authors in the data. Higher Gini values indicate greater inequality, i.e., a limited number of books or authors receiving most of the interactions.

Figure 3 shows that book-level inequality is consistently high and relatively stable, especially after 2012, with author-level inequality even higher. These results suggest persistent bias in user engagement toward popular books and authors.

Author gender representation. Finally, we examine author gender representation in user interactions by computing the monthly proportion of books interacted with that are written by female authors




Figure 3: Gini of monthly interactions.




Figure 4: Author gender representation.

(only considering books for which the first author’s gender identity is known).

As shown in Figure 4, the proportion of female authors rises sharply from around 30% in the early years to just over 50% by 2011, after which it remains relatively stable through 2017 with a small jump around 2014. This is consistent with the analysis of Thelwall and Kousha [37] suggesting that Goodreads exhibits a female-author bias, which in turn aligns with the hypothesis that female authors were more commercially successful than male authors due to being more frequently rated online [38], though further evidence would be needed to confirm this.

4.2 Recommendations

We now turn to the evaluation of recommendations generated by our selected collaborative filters for each time window. All measures are computed on top-100 recommendation lists for all users in a 2-month test window. We evaluate these recommendation lists using the same set of metrics applied in Section 4.1, alongside the ranking performance measures and rank-biased versions of entropy and Gini.

To account for user attention bias toward higher-ranked items in a recommendation list, we apply a position-based exposure weighting model to entropy and accuracy metrics in recommendations. Specifically, we use the geometric cascade browsing model [27, 33], defining the exposure weight w_i for position i in the ranked list as:

$$w_i = \gamma^{(i-1)} \quad (1)$$




Figure 5: Monthly unique entity counts in recommendations.

where $0 < \gamma < 1$ is a tunable *patience* parameter, which we set to LensKit’s default of 0.85 in our experiments.

Coverage and accuracy. Figure 5 shows the number of unique items and authors that appear in the top-100 recommendation lists for each algorithm over time. ItemKNN and BPR recommend substantially more unique items and authors compared to MostPop and ImplicitMF, which is a consistent pattern over time. This suggests that the former algorithms provide broader coverage across both items and authors.

We also report RBP [27], NDCG [22], and Mean Reciprocal Rank as ranking performance metrics over time (Figures 6a, 6b, 6c). ImplicitMF achieves the highest performance scores in all three metrics, followed by ItemKNN, BPR, and MostPop. The ordering of systems is mostly consistent over time, but there are noticeable fluctuations in performance. The performance gap between the best-performing models (ImplicitMF and ItemKNN) and the worst-performing models (MostPop and BPR) also appears to widen over time. The low scores of MostPop suggest that personalized models have performed better in capturing meaningful variations in user taste.

Genre diversity. For each recommendation list, we compute a weighted probability distribution over genres. The probability of genre g in a user’s list L is defined as:

$$P(g | L) \propto \sum_{i=1}^{|L|} w_i \cdot P(g | L(i)) \quad (2)$$

where w_i represents the exposure weight of position i and $P(g | L(i))$ is the probability that the book in position i belongs to genre g (see Section 3.1). We then compute entropy over the resulting genre distribution to obtain the genre entropy of a list, and average across lists to obtain an overall diversity measure of *mean rank-biased genre entropy*. As with profile entropy, larger entropy values indicate more diverse recommendation lists (i.e., more uniform distribution across genres).

Figure 7 shows the mean genre entropy across recommendation lists. Perhaps surprisingly, MostPop yields the highest entropy. We see at least two possible reasons for this. First, since MostPop recommends a fixed set of widely popular books, these items are more likely to have user-generated genre tags, possibly resulting

in higher measured entropy. In contrast, the other algorithms recommend a broader set of less popular items (as seen in Figure 5), which may include books with sparse or missing genre data that are excluded from the entropy calculation. Second, the most popular books themselves may be more diverse in genre, for example, including one highly popular book from each genre category.

BPR consistently exhibits the lowest genre diversity, likely due to its narrower focus on user-specific patterns in recommending items. Overall, all algorithms show relatively stable or slightly declining entropy over time, with more variation in MostPop. This suggests that personalized recommendation models are less influenced by changes in the genre distribution of the most popular books in any particular time period.

Individual fairness and popularity bias. We compute the Gini index ($Gini@k$) over the exposure weights w_i , which estimate the exposure assigned to an item or author at position i in the ranked recommendation lists. This effectively measures the allocation of “user attention” as a resource, accounting for the fact that users are likely to give more of their attention to books at the top of a recommendation list.

Figure 8 shows the Gini index for author-level (left) and item-level (right) exposure. As expected, MostPop displays the highest inequality in exposure, with both item- and author-level Gini values consistently close to 1.0. ItemKNN and BPR produce more balanced exposure, especially in earlier years. The ordering of ItemKNN and ImplicitMF we observe here differs from that found in the movie recommendation by Ekstrand et al. [11], indicating that models’ relative fairness properties may differ between domains or datasets.

All models show a gradual increase in Gini values over time and maintain relatively high levels of exposure inequality. This indicates a growing exposure bias toward a smaller set of items and authors, even in more personalized algorithms. The high Gini values are also influenced by the domain of items: for consistent comparison across both models and time windows, we computed Gini over all items in the dataset, with the unrecommended items receiving zero exposure.

Author gender representation. Figure 9 shows the proportion of books in the recommendations that have female authors. ItemKNN shows the highest and most steady increase in female author representation over time. ImplicitMF also increases gradually, while BPR stays more stable with a smaller rise. MostPop starts with the lowest proportion and increases slowly until a sharp jump near the end. All models show upward trends, but the speed and pattern of growth are different. The overall trend reflects the underlying patterns seen in the interaction data discussed in Section 4.1.

4.3 Effects of Goodreads Recommendation

To evaluate whether the introduction of Goodreads’ recommender system in 2011 is associated with measurable changes in system behavior, we analyzed temporal trends in both the interaction data and the outputs of our trained models. Overall, we did not observe clear or immediate changes in diversity, fairness, or effectiveness metrics around 2011 that can be directly associated with the launch of the recommender system.




Figure 6: Recommender system performance by offline metrics.

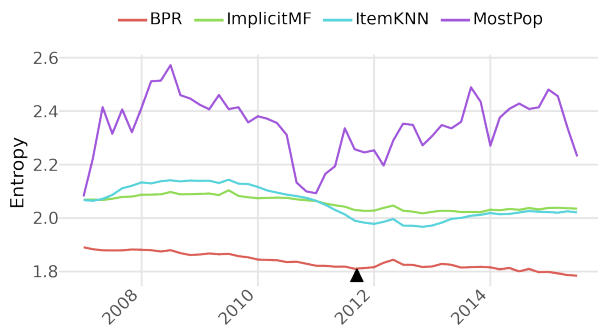


Figure 7: Genre entropy in recommendations.

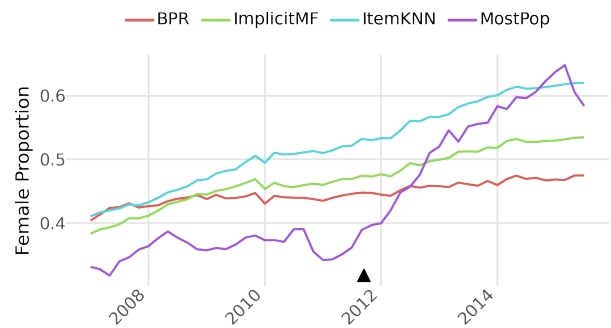


Figure 9: Author gender representation in recommendations.




Figure 8: Item and author popularity concentration (Gini @ 100).

However, there are some patterns that are worth further investigation. One such pattern is a gradual decrease in genre entropy in the interaction data after 2012, which may suggest an increasing influence of recommendation or narrowing user exploration. While the models differ in how they reflect underlying data trends, ImplicitMF and ItemKNN show signs of stabilization in both genre entropy and exposure Gini metrics following the system’s introduction. We also did not see identifiable changes in recommender effectiveness associated with the introduction of algorithmic recommendation; increased effectiveness on offline evaluation metrics may be a sign

of user interaction behavior becoming more predictable (due to the influence of algorithmic recommendations), but we do not see clear evidence of such an effect.

Overall, the observed trends appear to reflect broader, ongoing developments in the system, such as the accumulation of user interaction history, growth of the item catalog, and evolving user behavior, rather than the effects of a single intervention. Existing-system bias (bias in item exposure due to the recommender system being active while data is collected) is a significant problem in offline recommender system evaluation [6, 7, 23]. Goodreads’ lack of algorithmic recommendations for its first 4–5 years of operation provides an opportunity to detect potential existing-system biases. Our analyses were not able to detect specific problems attributable to the addition of a recommender system in this particular dataset. Algorithmic recommendations have significantly less prominence in the Goodreads user experience than they do in many other platforms, so they may have less of an impact on user behavior.

5 Discussion

We observed different trends in diversity and fairness measures across both the interaction data and the generated recommendations. We were particularly interested in understanding how aspects of the data that affect bias, fairness, diversity, and recommender performance changed over the history of the platform.

While user activity and the item corpus have increased over time, the popularity concentration in the interaction data (as measured by the Gini index) remains relatively stable, especially after 2012. This suggests that increased user engagement is not solely concentrated on the most popular books, but that users have had a relatively steady spread of interest in both more- and less-popular books over time. Moreover, although the interaction Gini stabilized after 2012, the exposure Gini for models like BPR and ImplicitMF continues to rise. This raises the question of whether there are other patterns in user interactions that lead to increased concentration in model outputs despite stable patterns in the underlying data. Future research should investigate this possibility to better understand the drivers of popularity bias in collaborative filtering outputs.

For genre diversity, the models largely mirror the gradual changes seen in the interaction data. However, there is a slight difference around 2012: while genre entropy in user interactions peaks around that time, most algorithms show a decline. One possible explanation is that changes in user behavior took time to influence model training data, leading to a delayed rise in entropy in the recommendations. Most of the models begin showing upward trends after this point, which may reflect a lagged response to shifts in genre diversity in the data.

Our findings, in general, and particularly the changing patterns in recommender system effectiveness metrics, indicate that evaluation outcomes are dependent on the choice of test and train time windows. This means the selection of time windows is important for the comparative evaluation of recommender systems, and suggests that results may be time-specific and potentially difficult to generalize. While we did see a few inversions of recommendation model performance rankings, we do observe a shift in the performance gap between models as time progresses in the dataset. This will influence the effect sizes reported in experiments on different temporal test windows. Temporal, data-centered audits of recommender systems can uncover evolving patterns in performance, fairness, and diversity that may not be visible in static, snapshot evaluations.

Our analysis does have several limitations. Some of our results are affected by sparse or incomplete metadata, such as missing genre tags or author information, which limit the accuracy of certain measures. The UCSD book graph also only contains public profiles that were still active in 2017 or 2018, so data from private or deleted accounts is not available. We also do not know what specific recommendation techniques Goodreads deployed in 2011 (or later), so we cannot model the performance changes in their recommendation models or directly test for their influence on the data. Finally, our results are observational and correlational, so we cannot establish causality between any of our metrics or between the introduction of the Goodreads recommender system and any user behavior changes. Nevertheless, we believe such descriptive and exploratory analyses are valuable for understanding the contours of the data used by the recommender systems and user modeling research communities to train and evaluate models and to study the behavior and social impacts of humans and algorithms on social data. Our findings also point to the need for further exploration of underlying drivers such as user engagement, metadata quality, and patterns our metrics could not detect but which nonetheless influence recommendation models over time.

6 Conclusion

Documenting and understanding datasets is important for designing experiments and contextualizing analyses and findings. The evolution of recommender systems datasets over time is often overlooked, but it is crucial for informing time-based experimental design decisions and provides context both for interpreting experimental results and for understanding historical changes when performing simulation studies of potential future dynamics [6, 15, 16].

We provided such an analysis for the UCSD Book Graph, documenting both the evolution of the dataset itself and the behavior of collaborative filtering models retrained at different time points. We found increasing representation of female authors, decreasing genre diversity, and slightly increasing concentration of interactions on popular users and items over the history of the dataset. We did not observe any clear impacts of the introduction of algorithmic recommendations on the behaviors of either users or collaborative filtering recommender models.

Our analysis lays the retrospective groundwork for future work on the temporal dynamics of recommender system behavior, including both traditional accuracy concerns and social impact considerations such as diversity and fairness. They also highlight the potential impact of the choice of splitting time on recommender experimental results. We encourage groups preparing and releasing new datasets for recommender systems research to include temporal analyses along with the descriptive statistics they publish with their datasets.

Acknowledgments

Compute time was provided by the Dragon Deep Learning Cluster, funded by NSF grant 23-20600.

References

- [1] 2011. Announcing Goodreads Personalized Recommendations. <https://www.goodreads.com/blog/show/303-announcing-goodreads-personalized-recommendations>.
- [2] Himan Abdollahpouri, Robin Burke, and Bamshad Mobasher. 2017. Controlling Popularity Bias in Learning-to-Rank Recommendation. In *Proceedings of the Eleventh ACM Conference on Recommender Systems (RecSys '17)*. 42–46. doi:10.1145/3109859.3109912
- [3] Alex Beutel, Jilin Chen, Tulsee Doshi, Hai Qian, Li Wei, Yi Wu, Lukasz Heldt, Zhe Zhao, Lichan Hong, Ed H Chi, et al. 2019. Fairness in Recommendation Ranking Through Pairwise Comparisons. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 2212–2220. doi:10.1145/3292500.3330745
- [4] Tesfaye Fenta Boka, Zhendong Niu, and Rama Bastola Neupane. 2024. A Survey of Sequential Recommendation Systems: Techniques, Evaluation, and Future Direction. *Information Systems* (2024), 102427. doi:10.1016/j.is.2024.102427
- [5] Robin Burke. 2010. Evaluating the Dynamic Properties of Recommendation Algorithms. In *Proceedings of the Fourth ACM Conference on Recommender Systems*. 225–228. doi:10.1145/1864708.1864753
- [6] Allison J. B. Chaney, Brandon M. Stewart, and Barbara E. Engelhardt. 2018. How Algorithmic Confounding in Recommendation Systems Increases Homogeneity and Decreases Utility. In *Proceedings of the 12th ACM Conference on Recommender Systems (RecSys '18)*. 224–232. doi:10.1145/3240323.3240370
- [7] Jiawei Chen, Hande Dong, Xiang Wang, Fuli Feng, Meng Wang, and Xiangnan He. 2023. Bias and Debias in Recommender System: A Survey and Future Directions. *ACM Transactions on Information Systems* 41, 3 (Feb. 2023), 1–39. doi:10.1145/3564284
- [8] Yashar Deldjoo, Dietmar Jannach, Alejandro Bellogin, Alessandro Difonzo, and Dario Zanzonelli. 2024. Fairness in Recommender Systems: Research Landscape and Future Directions. *User Modeling and User-Adapted Interaction* 34, 1 (2024), 59–108. doi:10.1007/s11257-023-09364-z
- [9] Mukund Deshpande and George Karypis. 2004. Item-Based Top-N Recommendation Algorithms. *ACM Transactions on Information Systems* 22, 1 (2004), 143–177. doi:10.1145/963770.963776

- [10] Michael D. Ekstrand. 2020. LensKit for Python: Next-Generation Software for Recommender Systems Experiments. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management (2020-10-19)*. 2999–3006. doi:10.1145/3340531.3412778
- [11] Michael D. Ekstrand, Ben Carterette, and Fernando Diaz. 2024. Distributionally-Informed Recommender System Evaluation. *ACM Transactions on Recommender Systems* 2, 1, Article 6 (2024). doi:10.1145/3613455
- [12] Michael D. Ekstrand, Anubrata Das, Robin Burke, Fernando Diaz, et al. 2022. Fairness in Information Access Systems. *Foundations and Trends® in Information Retrieval* 16, 1-2 (2022), 1–177. doi:10.1561/1500000079
- [13] Michael D. Ekstrand and Daniel Kluver. 2021. Exploring Author Gender in Book Rating and Recommendation. *User Modeling and User-Adapted Interaction* 31 (2021), 377–420. doi:10.1007/s11257-020-09284-2
- [14] Francesco Fabbri, Maria Luisa Croci, Francesco Bonchi, and Carlos Castillo. 2022. Exposure Inequality in People Recommender Systems: The Long-Term Effects. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 16. 194–204. doi:10.1609/icwsm.v16i1.19284
- [15] Andres Ferraro, Michael D. Ekstrand, and Christine Bauer. 2024. It's Not You, It's Me: The Impact of Choice Models and Ranking Strategies on Gender Imbalance in Music Recommendation. In *Proceedings of the 18th ACM Conference on Recommender Systems (RecSys '24)*. 884–889. doi:10.1145/3640457.3688163
- [16] Andres Ferraro, Xavier Serra, and Christine Bauer. 2021. Break the Loop: Gender Imbalance in Music Recommenders. In *Proceedings of the 2021 Conference on Human Information Interaction and Retrieval (CHIIR '21)*. 249–254. doi:10.1145/3406522.3446033
- [17] Yingqiang Ge, Shuchang Liu, Ruoyuan Gao, Yikun Xian, Yunqi Li, Xiangyu Zhao, Changhua Pei, Fei Sun, Junfeng Ge, Wenwu Ou, et al. 2021. Towards Long-Term Fairness in Recommendation. In *Proceedings of the 14th ACM International Conference on Web Search and Data Mining*. 445–453. doi:10.1145/3437963.344182
- [18] Tinnit Geburu, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé III, and Kate Crawford. 2021. Datasheets for Datasets. *Commun. ACM* 64, 12 (2021), 86–92. doi:10.1145/345872
- [19] Joan Giner-Miguel, Abel Gómez, and Jordi Cabot. 2022. DescribeML: A Tool for Describing Machine Learning Dataset. In *Proceedings of the 25th International Conference on Model Driven Engineering Languages and Systems: Companion Proceedings*. 22–26. doi:10.1145/3550356.3559087
- [20] Odd Erik Gundersen and Sigbjørn Kjensmo. 2018. State of the Art: Reproducibility in Artificial Intelligence. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 32. doi:10.1609/aaai.v32i1.11503
- [21] F. Maxwell Harper and Joseph A. Konstan. 2015. The MovieLens Datasets: History and Context. *ACM Transactions on Interactive Intelligent Systems* 5, 4, Article 19 (Dec. 2015). doi:10.1145/2827872
- [22] Kalervo Järvelin and Jaana Kekäläinen. 2002. Cumulated Gain-Based Evaluation of IR Techniques. *ACM Transactions on Information Systems* 20, 4 (Oct. 2002), 422–446. doi:10.1145/582415.582418
- [23] Olivier Jeunen. 2019. Revisiting Offline Evaluation for Implicit-Feedback Recommender Systems. In *Proceedings of the 13th ACM Conference on Recommender Systems (RecSys '19)*. 596–600. doi:10.1145/3298689.3347069
- [24] Neal Lathia, Stephen Hailes, and Licia Capra. 2009. Evaluating Collaborative Filtering Over Time. In *SIGIR '09 Workshop on the Future of IR Evaluation*. http://www.cs.ucl.ac.uk/staff/n.lathia/papers/lathia_ireval09.pdf
- [25] Zaiqiao Meng, Richard McCreadie, Craig Macdonald, and Iadh Ounis. 2020. Exploring Data Splitting Strategies for the Evaluation of Recommendation Models. In *Proceedings of the 14th ACM Conference on Recommender Systems*. 681–686. doi:10.1145/3383313.3418479
- [26] Monika Mishra, Jaydeep Choppe, Maitri Shah, Pankti Parikh, Rakshith Chandan Babu, and Jongwook Woo. 2019. Big Data Predictive Analysis of Amazon Product Review. In *KSIIT The 14th Asia Pacific International Conference on Information Science and Technology, APIC-IST, KSIIT, Beijing, China*. <https://www.calstatela.edu/sites/default/files/amazonprodreviewapic-ist2019.pdf>
- [27] Alistair Moffat and Justin Zobel. 2008. Rank-Biased Precision for Measurement of Retrieval Effectiveness. *ACM Transactions on Information Systems* 27, 1, Article 2 (Dec. 2008). doi:10.1145/1416950.1416952
- [28] Anna Nesvijevskaia. 2021. DATABOOK: A Standardised Framework for Dynamic Documentation of Algorithm Design During Data Science Projects. *IASSIST Quarterly* 45, 2 (2021). doi:10.29173/iq989
- [29] Tien T Nguyen, Pik-Mai 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 (WWW '14)*. 677–686. doi:10.1145/2566486.2568012
- [30] Alexandra Olteanu, Carlos Castillo, Fernando Diaz, and Emre Kiciman. 2019. Social Data: Biases, Methodological Pitfalls, and Ethical Boundaries. *Frontiers in Big Data* 2 (2019), 13. doi:10.3389/fdata.2019.00013
- [31] Tiago P Pagano, Rafael B Loureiro, Fernanda VN Lisboa, Rodrigo M Peixoto, Guilherme AS Guimarães, Gustavo OR Cruz, Maira M Araujo, Lucas L Santos, Marco AS Cruz, Ewerton LS Oliveira, et al. 2023. Bias and Unfairness in Machine Learning Models: A Systematic Review on Datasets, Tools, Fairness Metrics, and Identification and Mitigation Methods. *Big Data and Cognitive Computing* 7, 1 (2023), 15. doi:10.3390/bdcc7010015
- [32] Mahima Pushkarna, Andrew Zaldivar, and Oddur Kjartansson. 2022. Data Cards: Purposeful and Transparent Dataset Documentation for Responsible AI. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency (FAccT '22)*. 1776–1826. doi:10.1145/3531146.3533231
- [33] Amifa Raj and Michael Ekstrand. 2023. Unified Browsing Models for Linear and Grid Layouts. *arXiv preprint arXiv:2310.12524* (2023). doi:10.48550/arXiv.2310.12524
- [34] Amifa Raj and Michael D. Ekstrand. 2024. Towards Optimizing Ranking in Grid-Layout for Provider-Side Fairness. In *Advances in Information Retrieval (2024)*, Nazli Goharian, Nicola Tonello, Yulan He, Aldo Lipani, Graham McDonald, Craig Macdonald, and Iadh Ounis (Eds.). 90–105. doi:10.1007/978-3-031-56069-9_7
- [35] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2009. BPR: Bayesian Personalized Ranking from Implicit Feedback. In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence (UAI '09)*. 452–461.
- [36] Gábor Takács, István Pilászy, and Domonkos Tikk. 2011. Applications of the Conjugate Gradient Method for Implicit Feedback Collaborative Filtering. In *Proceedings of the Fifth ACM Conference on Recommender Systems (2011-10-23)*. 297–300. doi:10.1145/2043932.2043987
- [37] Mike Thelwall and Kayvan Kousha. 2017. Goodreads: A Social Network Site for Book Readers. *Journal of the Association for Information Science and Technology* 68, 4 (2017), 972–983. doi:10.1002/asi.23733
- [38] Marc Verboord. 2011. Cultural Products Go Online: Comparing the Internet and Print Media on Distributions of Gender, Genre and Commercial Success. *Communications* 36, 4 (2011), 441–462. doi:10.1515/comm.2011.022
- [39] Mengting Wan and Julian J. McAuley. 2018. Item Recommendation on Monotonic Behavior Chains. In *Proceedings of the 12th ACM Conference on Recommender Systems*. 86–94. doi:10.1145/3240323.3240369
- [40] Mengting Wan, Rishabh Misra, Ndapa Nakashole, and Julian J. McAuley. 2019. Fine-Grained Spoiler Detection from Large-Scale Review Corpora. In *Proceedings of the 57th Conference of the Association for Computational Linguistics*. 2605–2610. doi:10.18653/v1/P19-1248
- [41] Yifan Wang, Weizhi Ma, Min Zhang, Yiqun Liu, and Shaoping Ma. 2023. A Survey on the Fairness of Recommender Systems. *ACM Transactions on Information Systems* 41, 3, Article 52 (Feb. 2023). doi:10.1145/3547333
- [42] Yuyan Wang, Mohit Sharma, Can Xu, Sriraj Badam, Qian Sun, Lee Richardson, Lisa Chung, Ed H. Chi, and Minmin Chen. 2022. Surrogate for Long-Term User Experience in Recommender Systems. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 4100–4109. doi:10.1145/3534678.3539073
- [43] Emre Yalcin and Alper Bilge. 2022. Evaluating Unfairness of Popularity Bias in Recommender Systems: A Comprehensive User-Centric Analysis. *Information Processing & Management* 59, 6 (2022). doi:10.1016/j.ipm.2022.103100