JETIR.ORG ISSN: 2349-5162 | ESTD Year : 2014 | Monthly Issue
JOURNAL OF EMERGING TECHNOLOGIES AND INNOVATIVE RESEARCH (JETIR)
An International Scholarly Open Access, Peer-reviewed, Refereed Journal

HIDDEN BIAS IN BOOK RECOMMENDATION

[1]Mrs. Pragathi., [2]Adarsh Sangam., [3]Adarsha S V, [4]Aditya.,[5]Arjun Garg

[1]Assistant Professor, [2]Student, [3]Student, [4]Student,[5]Student
[1]Computer Science and Engineering,
[1]Sir M Visvesvaraya Institute of Technology, Bangalore, India

Abstract: Book recommendation systems are designed to suggest books to readers based on their interests and reading history. However, these systems can be influenced by hidden biases, which can lead to the exclusion of certain authors or literature and perpetuate harmful stereotypes. These systems use algorithms to analyze users' reading histories and make recommendations based on the books that they have previously read. However, these systems are not perfect, and they can be biased in various ways. This research paper will examine the presence of hidden bias in book recommendation systems and its impact on readers. It will explore the ways in which these biases can be introduced, including through the data used to train the recommendation algorithms and the personal biases of those who design and implement the systems.

***Keywords*:** Book recommendation, Hidden bias,Recommendation algorithms.

I. INTRODUCTION

Book recommendation systems have become an integral part of the reading experience for many individuals, as they offer a convenient way to discover new books and authors. These systems often rely on algorithms that analyze users' reading histories and preferences in order to generate personalized recommendations. However, despite their popularity and convenience, book recommendation systems are not without their flaws, as they are prone to biases that can significantly impact the books and authors that are recommended to readers.

Bias in book recommendation systems can take many forms and can be present at various stages of the recommendation process. For example, the data used to train the algorithms that generate recommendations may be biased towards certain types of books or authors, leading to a lack of diversity in the recommendations. The algorithms themselves may also be biased, either due to their design or the data used to train them. These biases can also come in the form of gender bias, author bias, and popularity bias, and can significantly impact the books and authors that are recommended to readers. Additionally, the recommendations themselves may be biased if they are influenced by external factors such as marketing or the popularity of certain books or authors.

One potential source of bias in book recommendation systems is the data used to train the algorithms. If the data is not representative of the diverse range of books and authors that are available, this can result in biased recommendations. This is particularly concerning given that the publishing industry has a history of excluding marginalized groups and promoting a narrow range of perspectives.

The algorithms themselves may also be biased, either due to their design or the data used to train them. For example, algorithms that are trained on data that is biased towards certain types of books or authors may end up recommending these types of books disproportionately, leading to a lack of diversity in the recommendations. Additionally, algorithms that are designed to prioritize popularity or marketing may end up recommending books that are well-known or heavily promoted, rather than books that may be of interest to the user but may not have the same level of visibility.

Finally, the recommendations themselves may be influenced by external factors such as marketing or the popularity of certain books or authors. This can lead to a self-reinforcing cycle, where popular books or authors are recommended more frequently, leading to even more popularity and visibility. This can make it difficult for lesser-known or marginalized authors to gain exposure and can perpetuate existing inequalities in the publishing industry.

In conclusion, bias in book recommendation systems can have significant consequences for the books and authors that are recommended to readers. It is important to understand the various sources of bias in these systems and to develop strategies for mitigating their effects in order to promote a more diverse and equitable reading experience for all individuals. This research paper aims to examine these issues in greater detail and to explore potential solutions for addressing bias in book recommendation systems.

II. RELATED WORK

Yehuda Koren, et al.,2022, investigates Collaborative filtering which is a technique used in recommender systems to make personalized recommendations to users by predicting their interests based on the preferences of similar users. Advances in the collaborative filtering have led to the development of more sophisticated and effective recommender systems, which can provide high- quality recommendations to users by leveraging large amounts of data and computational power. Some of the key advances in collaborative filtering include the development of matrix factorization techniques, which can learn complex non-linear relationships between users and items.

Leeandlow Books. et al.,2022, provides data on the diversity of the publishing industry in the United States. The survey was conducted by Lee & Low Books, a publisher of children's and young adult books, and aimed to gather data on the demographics of publishing professionals, including information on race, ethnicity, gender, age, and sexual orientation. The survey found that the publishing industry is largely dominated by white professionals, with white employees making up 78% of the industry. The survey also found that women are underrepresented in certain areas of the publishing industry, particularly in leadership roles, and that people of color are underrepresented across the industry as a whole. The report concludes that the publishing industry needs to increase diversity in order to better reflect the demographics of the country and to ensure that diverse voices and perspectives are represented in the books that are published.

Michael D Ekstrand, et al.,2021, investigates the impact of author gender on book rating and recommendation systems. The paper uses data from a book rating and recommendation website to analyze the relationship between author gender and rating scores and the influence of author gender on book recommendation. The results of the study suggest that there are gender- based differences in book rating and recommendation, with books written by male authors receiving higher rating scores and being more likely to be recommended compared to books written by female authors. The paper discusses the potential implications of these findings for the accuracy and fairness of recommendation systems and suggests future directions for research in this area.

Dominik Kowald, et al.,2020, investigates the phenomenon of popularity bias in music recommendation systems. The paper defines popularity bias as the tendency of recommendation systems to prioritize popular items over less popular ones, and including studies that have found that popular songs are including studies that have found that popular songs are more likely to be recommended to users, and that this can lead to a reinforcement of pre-existing popularity patterns.

Himan Abdollahpouri., et al.,2020, examines the issue of popularity bias in recommendation systems from the perspective of multiple stakeholders, including users, content creators, and platform owners. The paper discusses the negative consequences of popularity bias, including the suppression of diversity and the amplification of existing biases, and suggests potential solutions to address this issue, such as incorporating diverse perspectives in the design of recommendation systems and using diversity- aware evaluation metrics. The paper also highlights the need for a multistakeholder approach to addressing popularity bias in recommendation systems, which takes into account the needs and concerns of all parties involved.

Himan Abdollahpouri., et al.,2019, examines the issue of popularity bias in recommendation systems. The study analyzed the recommendation systems of two music streaming platforms, Spotify and Pandora, and found that both exhibited popularity bias, meaning that they tended to recommend more popular items more frequently than less popular ones. This can have negative consequences, including the amplification of existing biases and the suppression of diverse or alternative content. The paper also discusses the potential causes of popularity bias in recommendation systems and suggests strategies for addressing it. It is an important contribution to the field of artificial intelligence ethics and the design of recommendation systems.

Himan Abdollahpouri., et al.,2019, examines the issue of popularity bias in recommendation systems and its impact on the fairness of these systems. The paper discusses the negative consequences of popularity bias, including the suppression of diversity and the amplification of existing biases, and suggests potential solutions to address this issue, such as incorporating diverse perspectives in the design of recommendation systems and using diversity-aware evaluation metrics. The authors also argue that addressing popularity bias is crucial for ensuring the fairness and ethicality of recommendation systems.

Nettie Finn, et al.,2016, investigates the use of pen names or pseudonyms as a way to mitigate gender bias in the publishing industry. The paper reviews previous research on gender bias in publishing, including studies that have found that male authors are more likely to be published and reviewed than female authors, and that female authors are more likely to use pen names or pseudonyms than male authors. The paper also examines the use of pen names or pseudonyms by female authors in different genres, and the potential impact of these pseudonyms on the perceived quality and success of the author's work.

Erik Brynjolfsson, et al.,2006, explores the concept of the long tail in the context of the digital economy. The long tail refers to the phenomenon where the market is dominated by a small number of popular items, with a large number of less popular items making up the rest of the market. The paper discusses the potential for the long tail to bring about greater diversity and democratization in the digital economy, as it allows for the emergence of niche products and services that may not have been feasible in traditional brick-and-mortar markets.

Greg Linden,. et al.,2003, describes the development and implementation of an item-to-item collaborative filtering system for the online retailer Amazon.com. Collaborative filtering is a technique used in recommender systems to make personalized recommendations to users based on the preferences of similar users. The paper describes how Amazon.com implemented an item-to-item collaborative filtering system, which made recommendations to users based on the items that they had previously purchased or rated, as well as the items that similar users had purchased or rated. The paper discusses the challenges and benefits of implementing this system, including the large scale of the Amazon.com website and the need to handle high levels of data sparsity and skew.

III. METHODOLOGY

Data Collection

To investigate the hidden bias in book recommendation, we chose the widely used Book-Crossing dataset [34], which was created in 2004 by crawling the website of Book-Crossing for four weeks. Book-Crossing is a social network and book exchange website that allows people to share their books with others around the world. The concept behind Book-Crossing is to encourage book sharing and the spread of literacy while also promoting recycling and reducing waste. Users update their profile with the information about the books they have read along with a rating from 1 to 10. The dataset consists of three sub-datasets: book ratings, user profiles with their location information, and book information such as the name of the author, publisher, and year of publication. Since the dataset lacked complete author information, we enhance it with additional data from external sources. We took the following steps, illustrated in Figure 1, to gather more data:

1. Linked the Book-Crossing dataset with Google Books using ISBN.
2. Validated and corrected the author name provided in Book-Crossing using the author name included in Google Books.
3. Linked the dataset to Virtual International Authority File (VIAF) using author name.
4. Validated and enriched the data with author country of citizenship from WikiData using VIAF ID.

In order to investigate bias in book recommendations. We incorporated the following approach, which involved:

• Removing implicit ratings

• Removing users with more than 200 ratings

• Removing users with less than 5 ratings

• Removing items with less than 5 ratings

Such cut-offs are commonly applied in the recommendation literature due to the sparsity of data on online platforms. Therefore, we also analysed the effect of these cut-offs on the bias of the dataset.

Moreover, we processed the dataset to remove duplicate items with different ISBN numbers and ISBN numbers included in the ratings dataset but not in the items dataset, which we considered as mistakes. We have processed the entire dataset and present its general characteristics in the subsequent section.

Analysing Bias

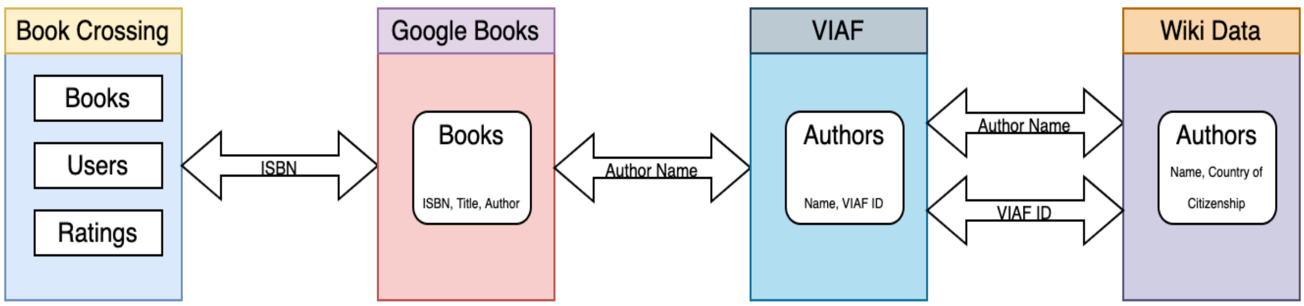

Fig: analysing bias

To assess the presence of bias in book recommendations, we developed a method to measure the potential bias based on country of citizenship and compared it to the impact of popularity bias. We used the Book Crossing dataset to analyse the relationship between the popularity of books and the country of citizenship of their authors. Popularity was measured by the number of ratings an item received. Next, we trained 11 different recommender algorithms, including 9 collaborative filtering algorithms and 2 dummy approaches, using an 80-20% split. We then generated recommended lists for each user based on the predicted rankings of each algorithm and examined the distribution of author country of citizenship in both the users' profiles and recommended lists for each algorithm. By comparing the distribution before and after recommendations, we were able to observe the potential impact of each algorithm's recommendations on the distribution of books from different countries. This allowed us to compare the propagation of country of citizenship bias to that of popularity bias and evaluate the effectiveness of each algorithm in mitigating bias.

III.RESULTS

Bias in the dataset

Our analysis of the dataset revealed a clear bias in favour of books written by US citizen authors. 36% of the unique books in the dataset were written by American authors. However, when the cut-offs were introduced, the percentage of American-authored books increased to 68.5%, as shown in Figure. The same trend was observed in the ratings datasets, as seen in Figures. In the entire dataset, 55.3% of the ratings were given to books authored by US citizens, and in the Bookcrossing ratings, the percentage increased to 74.4%.

Our findings suggest that introducing the cut-offs in the dataset exacerbated the bias in favour of books authored by US citizens. This observation implies a direct relationship between item popularity and the country of citizenship of the author, as Book Crossing excludes items with less than 5 ratings. According to the Book-Crossing website, only 20% of their users are based in the US, which suggests that American authors are overrepresented in the dataset compared to the number of American users.

Lastly, we conducted a t-test to compare the mean popularity of books written by American authors and those written by non-American authors. The result showed a significant difference between the two sets, indicating that American-authored books are generally more popular than non-American-authored books.

Recommendation Bias Analysis

Following the process described earlier, we analysed the ratio of American-authored books in users' profiles versus the recommended list of each algorithm. Figure illustrates that, except for Random and Most Popular, all collaborative filtering recommenders increase the ratio, with NMF, MF, and PMF being exceptions. To compare these tendencies to popularity bias, we used the %ΔGAP metric. This metric expresses the relative increase of average item popularity in a user's recommended list compared to their profile. Figure shows that all collaborative filtering algorithms, except for PMF, MF, and NMF, produce recommended lists with an increased average item popularity, which means that the same algorithms that recommend excessively American-authored books also recommend excessively popular books. We then compared the ratio of US citizen authors in a user's profile to the recommended list of each recommender algorithm for each user individually.

The results of our study reveal interesting insights into the behavior of recommendation algorithms in terms of popularity bias. Below Fig. displays the average ΔGAP values for the different algorithms across three user groups - Niches, Diverse, and Bestseller-focused. Our findings indicate that Niches users receive significantly higher ΔGAP values, followed by Diverse and Bestseller-focused users, respectively. These results confirm the findings of a previous study by Abdollahpouri et al., highlighting the fact that popularity bias impacts users who are interested in unpopular items the most. Additionally, our study shows that although Bestseller-focused users receive the most favorable recommendations, the average ΔGAP value is 126.55, indicating that algorithms can propagate the popularity bias even further than the user groups' interest in popular items. Furthermore, our analysis reveals that the investigated algorithms exhibit similar behavior in terms of popularity bias across all user groups.

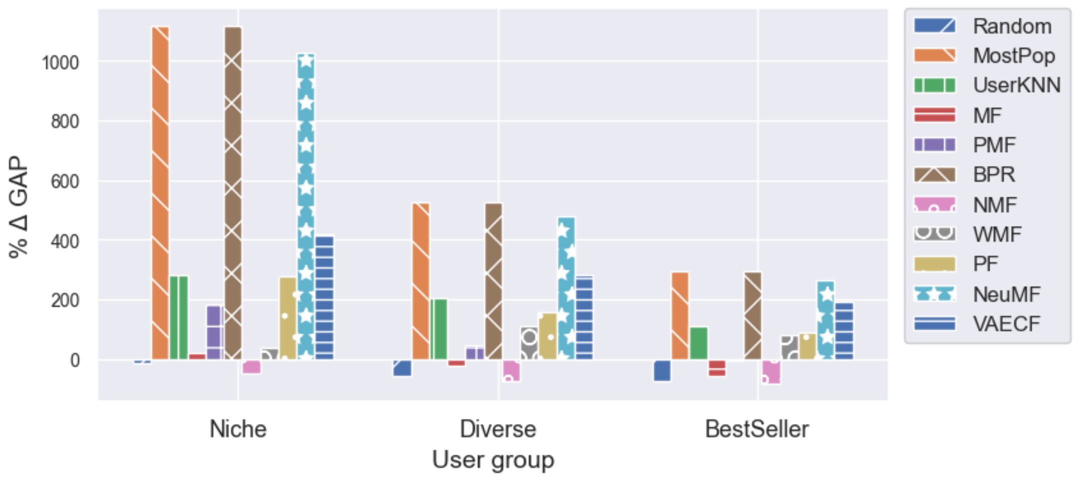

Fig: the behavior of recommendation algorithms in terms of popularity bias

Analysis

Firstly, analysing the gender distribution of authors in the Book Crossing dataset helped us to identify any potential gender biases that may be present in book recommendation systems. If the recommendation system is biased towards recommending books by male authors over female authors, we see equal distribution for male and female users.

Similarly, analysing the country of citizenship distribution of authors in the dataset helped us to identify any potential biases towards certain countries or cultures. If the recommendation system is biased towards recommending books by authors from certain countries or cultures over others, this may indicate a cultural or national bias in the system. In the analysis, we saw that out of the entire dataset, 55.3% of the ratings were given to books authored by US citizens, and in the Book Crossing ratings, the percentage increased to 74.4%.

Analysing the native language distribution of authors in the dataset can also help to identify potential biases towards certain languages. If the recommendation system is biased towards recommending books by authors who write in certain languages over others, this may indicate a language bias in the system. In the analysis, we saw that out of the entire dataset, 23% of the books were written and published by English authors, and in the Book Crossing ratings, the percentage increased to 29.8%.

Finally, analysing the birth year distribution of authors can help to identify potential biases towards certain time periods or generations. If the recommendation system is biased towards recommending books by authors from certain time periods or generations over others, this may indicate a generational bias in the system. We see that most birthyears are as concentrated in the 20th century.

Overall, by analysing the various distributions in the Book Crossing dataset, we can gain a better understanding of potential biases that may exist in book recommendation systems. This can help to identify and mitigate these biases, resulting in a more equitable and inclusive recommendation system.

Based on our findings, we can confidently state that certain recommender algorithms exhibit bias against authors' country of citizenship, even when this is not a feature included in the training process. When we compared the propagation of American author bias and popularity bias, we found that they shared similar characteristics. Therefore, we believe that the observed author bias was directly influenced by the known phenomenon of popularity bias.

Our results suggest that American book platforms accessible to users worldwide may have a US-centric bias, and this bias can be perpetuated when certain recommenders are used without proper scrutiny. The lack of geodiversity in open data is a well-known issue that needs to be addressed, whether the data is used for research or commercial purposes. Our findings also indicate that popularity bias in recommendation should not be viewed merely as an issue of weak performance that mainly affects e-commerce.

When we reviewed the birth year distribution of authors in the Book Crossing dataset, we observed that the majority of authors were born in the 20th century. This finding suggests that there may be a bias towards recommending books by authors from a particular time period, which could potentially exclude newer or more diverse voices in literature.

The fact that the majority of authors were born in the 20th century is not necessarily problematic on its own. However, if a recommendation system is biased towards recommending books by authors from a particular time period, it may result in a limited and potentially outdated selection of books being recommended to users. This could lead to a lack of diversity in the books being recommended and potentially exclude newer or more diverse voices in literature.

When analysing the native language distribution of authors in the Book Crossing dataset, we found that 23% of the books in the dataset were written and published by English authors. However, when looking at the ratings given by users in the Book Crossing ratings dataset, the percentage of books written by English authors increased to 29.8%.

This finding suggests that there may be a bias towards recommending books by English authors in the Book Crossing ratings dataset. It could also indicate that books by English authors are more highly rated by users in the dataset, leading to a higher percentage of English-authored books in the ratings data.

However, it is important to note that this finding may not necessarily indicate a bias in the recommendation system itself. Instead, it may reflect the preferences of users in the Book Crossing ratings dataset, who may be more likely to rate books by English authors highly.

We must be careful not to demote certain books based on their authors' sensitive features, which may be correlated with past unpopularity. For this reason, our work highlights the need for researchers to conduct a thorough examination of potential biases in recommendation systems and make an active effort to address them, even when popularity appears to be the cause.

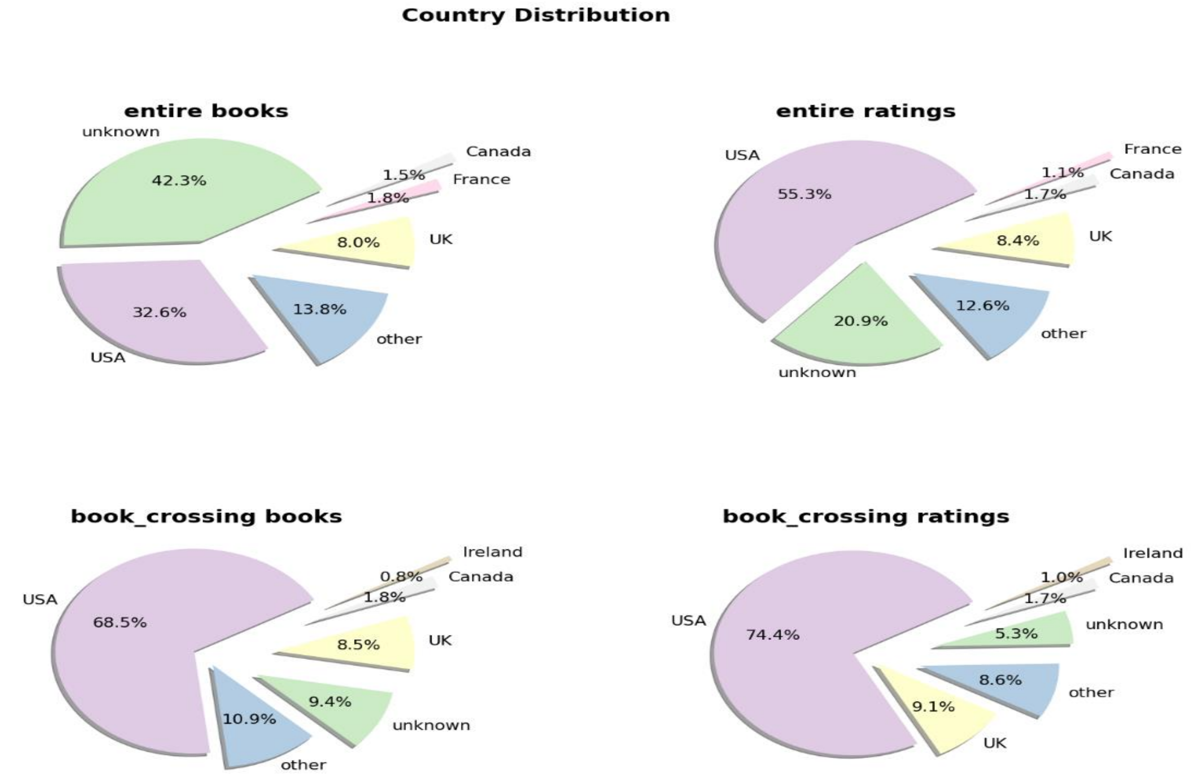

Fig: country distribution

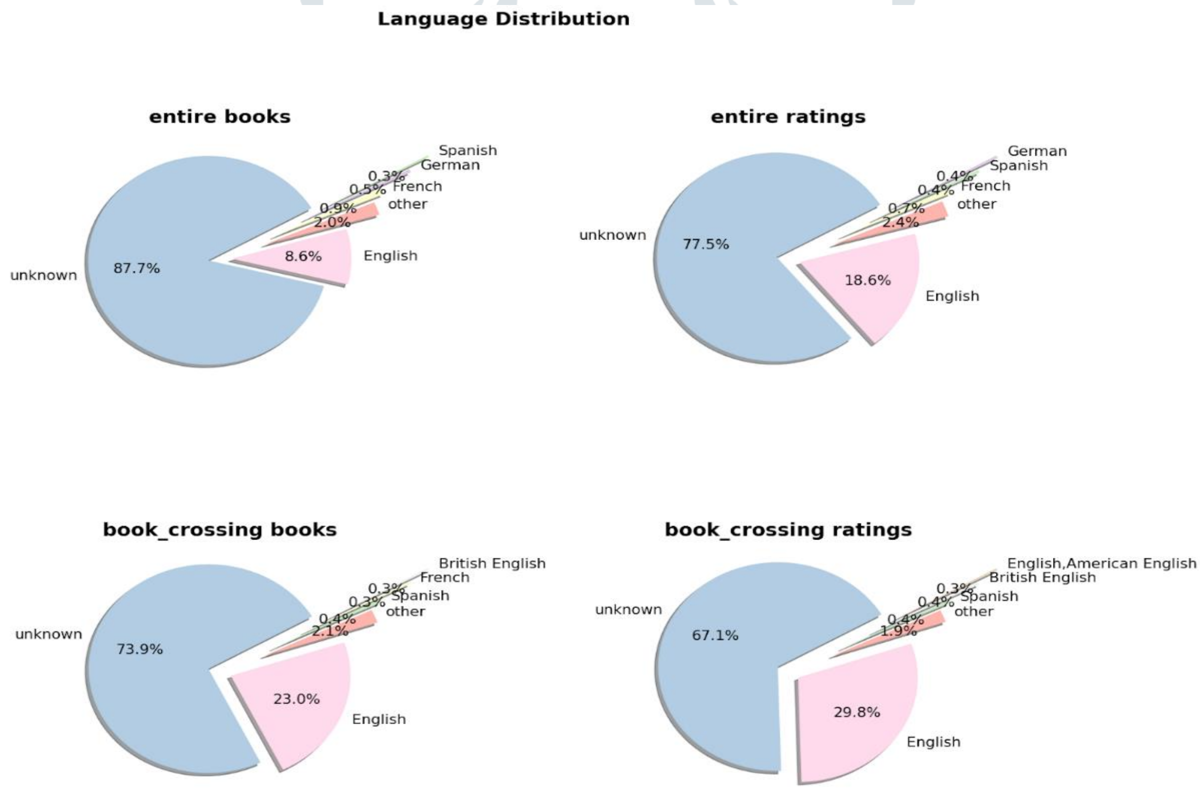

Fig: language distribution

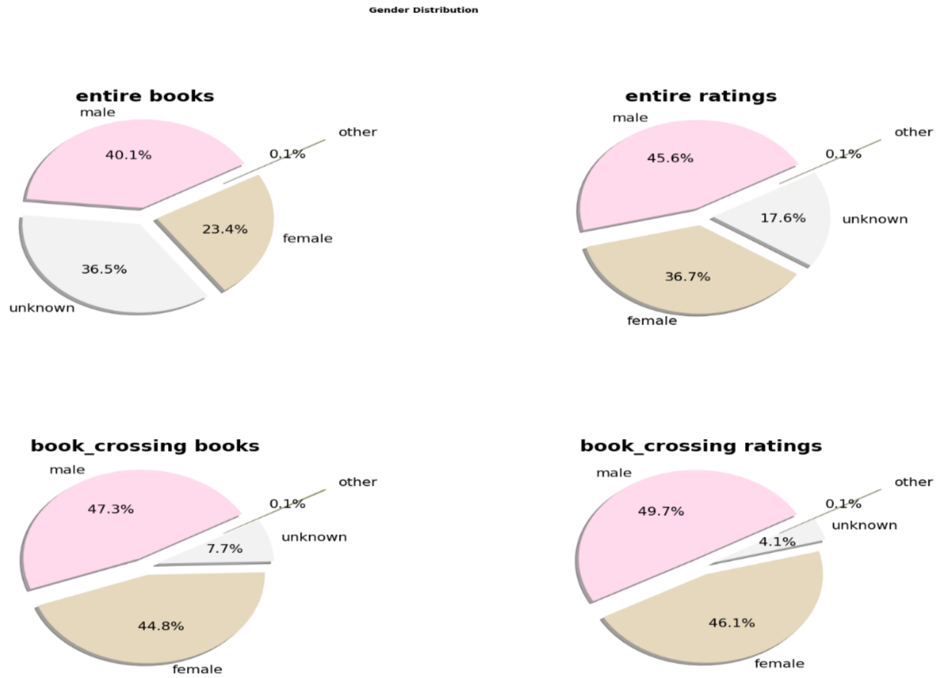

Fig: gender distribution

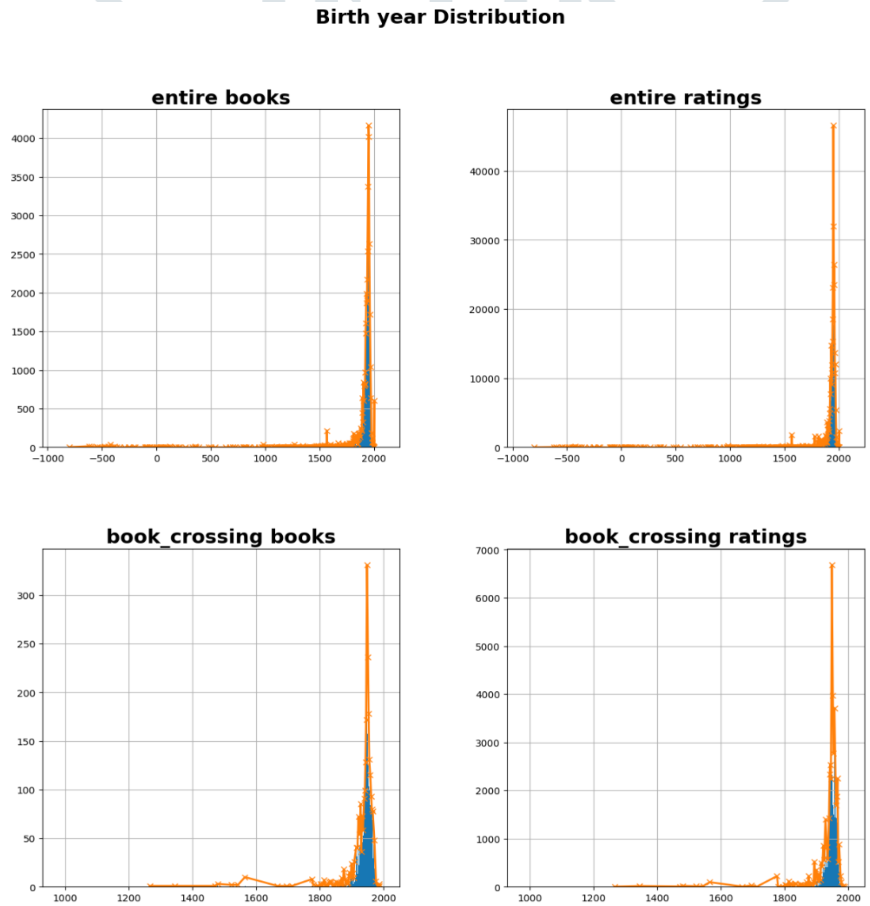

Fig: birth year distribution

IV. CONCLUSION

The conclusion of this research paper emphasizes the importance of addressing hidden bias in book recommendation systems. The study focused on the phenomenon of hidden bias introduced by commonly used recommender algorithms that do not take any item features as input. The researchers theorized that feature bias comes as a direct result of popularity bias and investigated this hypothesis in the context of book recommendations using a well-known book ratings dataset.

The findings showed that books written by American authors within the dataset were significantly more popular compared to the rest, and that many commonly used collaborative filtering algorithms on average recommend more American-authored books than in the users' profile. This indicates a potential bias towards recommending books by American authors over authors from other countries.

The researchers highlight the need for fair treatment of social groups and the importance of ensuring that the recommendation system is inclusive of authors from different countries and cultures. They also acknowledge the limitations of using existing data sources such as WikiData and the potential bias that may be introduced through these sources.

The conclusion emphasizes the need for a well-rounded understanding of hidden bias in book recommendation and the importance of considering various dimensions of the problem, including the bias towards authors, books, and users. By properly accounting for hidden bias, libraries can benefit from the automation recommender systems offer in attracting users, while ensuring that their value of inclusivity is being properly adhered to. Future work will explore these dimensions further and consider methods for addressing hidden bias in book recommendation systems.

V. REFERENCES

[1] Himan Abdollahpouri. 2020. Popularity bias in recommendation: a multistakeholder perspective. Ph.D. Dissertation. University of Colorado at Boulder.

[2] Himan Abdollahpouri, Masoud Mansoury, Robin Burke, and Bamshad Mobasher. 2019. The Unfairness of Popularity Bias in Recommendation. RecSys Workshop on Recommendation in Multistakeholder Environments (RMSE) ; Conference date: 20-09-2019

[3] Erik Brynjolfsson, Yu Jeffrey Hu, and Michael D Smith. 2006. From niches to riches: Anatomy of the long tail. Sloan management review 47, 4 (2006), 67–71.

[4] Michael D Ekstrand and Daniel Kluver. 2021. Exploring author gender in book rating and recommendation. User modeling and user adapted interaction 31, 3 (2021), 377–420.

[5] Nettie Finn. 2016. Pseudonymous Disguises: Are Pen Names An Escape From the Gender Bias in Publishing (2016).

[6] Yehuda Koren, Steffen Rendle, and Robert Bell. 2022. Advances in collaborative filtering. Recommender systems handbook (2022), 91–142.

[7] Greg Linden, Brent Smith, and Jeremy York. 2003. Amazon. com recommendations: Item-to-item collaborative filtering. IEEE Internet computing 7, 1 (2003), 76–80.

[8] Leeandlow Books. 2022. Where is the diversity in publishing? The 2019 Diversity Baseline Survey Results.

[9] Harald Steck. 2011. Item popularity and recommendation accuracy. In Proceedings of the fifth ACM conference on Recommender systems. 125–132.

[10] Shrikant Saxena and Shweta Jain. 2021. Exploring and Mitigating Gender Bias in Recommender Systems with Explicit Feedback. arXiv preprint arXiv:2112.02530 (2021).

.