IEEE Access
Multidisciplinary | Rapid Review | Open Access Journal

Received 23 February 2024, accepted 21 May 2024, date of publication 13 June 2024, date of current version 20 June 2024.

Digital Object Identifier 10.1109/ACCESS.2024.3413772

RESEARCH ARTICLE

The Effects of Media Bias on News Recommendations

QIN RUAN[1,3], BRIAN MAC NAMEE[1,2,3], AND RUIHAI DONG[1,2,3], (Member, IEEE)
[1]School of Computer Science, University College Dublin, Dublin, D04 V1W8 Ireland
[2]Insight SFI Research Centre for Data Analytics, Dublin, D04 V1W8 Ireland
[3]Science Foundation Ireland Centre for Research Training in Machine Learning, Dublin, D04 V1W8 Ireland

Corresponding authors: Ruihai Dong (ruihai.dong@ucd.ie) and Qin Ruan (qin.ruan@ucdconnect.ie)

This work was supported by the Science Foundation Ireland under Grant 18/CRT/6183 and 12/RC/2289_P2.

ABSTRACT The negative effects of media bias, such as influencing readers' perceptions and affecting their social decisions, have been widely identified by social scientists. However, the combined impact of media bias and personalised news recommendation systems has remained largely unstudied, especially in real-world news recommendation datasets. Our study bridges this gap by analysing how leading algorithms influence the spread of biased news among news recommendation system users with diverse preferences. In this article, we show that current state-of-the-art news recommendation algorithms amplify the amount of biased media that readers consume and that, while the quality of their recommendations is largely similar, different news recommendation algorithms have differing sensitivities to media bias. We present experimental results that compare the performance of different news recommendation algorithms for users with different subject interests and different levels of prior history of reading biased media. Our analysis reveals that some state-of-the-art news recommendation algorithms that perform well at the recommendation task also lead to large amounts of biased news being recommended to readers. These findings suggest significant potential for negative impacts from increasing volumes of biased media being promoted by news recommendation algorithms. This highlights the importance for organisations to offer more trustworthy personalised news recommendations to mitigate the propagation of bias in news consumption.

INDEX TERMS Algorithmic media bias, filter bubbles, media bias, media bias detection, news recommendation.

I. INTRODUCTION

The primary role of the news media is to present objective and unbiased factual reporting to its readers [1]. Widespread bias in news media, however, means the modern media is failing in its mission of unbiased reporting [2]. Media bias refers to unjustifiable favouritism exhibited by media providers as they cover the news [3]. For example, journalists, and the media organizations that they represent, may only report facts favourable to a particular political view and promote opinions aligned with that view [4], [5]. The following news headlines, compiled from the *allsides* website,[1] that appeared shortly after Donald Trump announced his intention to run in the 2024 USA presidential campaign show examples of media bias:

- **Biased Article 1:** Trump Org. controller said he was ordered to hide benefits on tax forms. – *CNN (Online News), left-oriented*
- **Biased Article 2:** How Trump, infighting and flawed candidates limited Republican gains. – *Washington Post, left-oriented*
- **Biased Article 3:** CPAC Chairman to Newsmax: "Silly" to Blame Trump for GOP Setbacks. – *Newsmax (News), right-oriented*
- **Neutral Article 4:** GOP leaders to Trump after midterms: Delay big announcement. – *NewsNation, center*

Bias can exist in the news in the form of suggestive words, such as "infighting", "flawed" and "silly" in the examples above. It can also be more systematic as part of

The associate editor coordinating the review of this manuscript and approving it for publication was Pasquale De Meo.

[1]https://www.allsides.com/

a political ideology covered by media establishments. The side effects of media bias—for example, distorting readers' perceptions [6] and influencing political elections [7]—have been extensively studied by social scientists. It is also widely agreed that raising public awareness, such as flagging bias in the news, is crucial in a democratic society [8].

News recommendation systems use personalised recommendation algorithms to help readers navigate collections of news articles quickly and efficiently [9]. Modern personalised news recommendation algorithms [10] typically employ machine learning models to learn from user data (including users' interests and reading history) allowing them to learn representations of news articles and user behaviours. These representations are then employed to estimate the probability that a user will click on a particular article, and then to promote the articles most likely to lead to clicks. The main differences among leading news recommendation algorithms lie in the data sources used to generate news representations, the network architectures used to learn news and user representations, and the specific methods used to generate click probability predictions.

In the digital age, the intersection of media bias and personalised news recommendation systems has garnered some attention. Studies [11], [12], [13] have suggested that the political polarisation seen on social networks might be driven by personalised recommendation algorithms. These social networks can create a feedback loop where users are increasingly exposed to content that aligns with their existing beliefs, potentially intensifying media bias and reinforcing political polarisation. Furthermore, simulation experiments [14], [15], [16], [17] have shown that the latest news recommendation algorithms tend to suggest more biased news to simulated users who prefer such content, and more neutral news to those who like that instead.

Our work offers a new perspective on the role of news recommendation algorithms in disseminating biased news articles. By employing a real-world news recommendation dataset, we investigate how leading news recommendation algorithms affect the volume of biased news articles recommended to users with varying preferences for such content. We trained news recommendation algorithms using this news recommendation dataset, integrating a media bias detector to identify biased news articles from real user reading histories. We designed two user grouping strategies: one that groups users based on the proportion of biased news in their reading histories, and another that groups users based on the interests (e.g. sports, news, politics) represented in their reading histories as well as the proportion of biased news. We applied these two grouping strategies to segment real users within our news recommendation dataset. Presenting the same dataset of news articles to all users and analysing the biased content within the top recommendations allows us to uncover the relationship between different user groups and the dissemination of biased news articles by news recommendation algorithms.

We address two critical research questions that guide our exploration of the tendencies of news recommendation algorithms to disseminate biased content:

- **RQ1**: Are news recommendation systems, influenced by users' historical reading biases, leading to recommendations of more biased articles?
- **RQ2**: If news recommendation systems are indeed influenced by users' historical reading biases, is the level of influence different for different recommendation algorithms?

To answer these research questions, we conduct experiments that observe changes in the proportion of biased news articles recommended by well-trained recommendation algorithms to user groups with differing degrees of biased news in their reading histories. In the experimental setting, we focus on predicting whether news articles are biased or unbiased without distinguishing between left-, centre- or right-leaning bias. We do this because previous work has found that left- and right-leaning news articles are similar to each other but have a significant difference from centre-leaning news [4], [18]. Moreover, we are more concerned with whether, and to what extent, recommendation algorithms are affected by media bias overall than with different kinds of bias. The results of our experiments show that news recommendation systems are indeed influenced by media bias, and that the sensitivity of different recommendation algorithms to media bias varies depending on how user behaviour sequences are modelled, and how recommendation algorithms generate recommendations. We believe that our exploration of the side effects of media bias on news recommendation algorithms is the first of its kind in the literature, that it lays a solid foundation for identifying and researching media bias in news recommendation systems, and that it will be a springboard for further work.

The remainder of this article is organized as follows: Section II describes relevant existing work on news recommendation and media bias; Section III outlines the methodologies used in this study, including the generation of news recommendations, the detection of media bias, the creation of user groups, and the analysis of bias proportion within the top-k recommendations; Section IV describes the experimental method; results are discussed in Section V; and, finally, Section VI provides conclusions and directions for future work. The source code for our framework and experiments is available to the public.[2]

II. RELATED WORK

In this section, we review the relevant literature on news recommendation algorithms, media bias, media bias detection, and the amplification of media bias through algorithmic personalisation.

[2]https://github.com/ruanqin0706/NewsRec.git

A. NEWS RECOMMENDATION ALGORITHMS

It is common to consider modelling news recommendations as a sequential task—which is also a widespread practice in other recommendation domains, such as e-commerce recommendations [19] and movie recommendations [20]. Sequential recommendation algorithms assume that, in the short-term, historical user behaviours influence users' future decisions [21]. Based on this assumption, researchers have made many attempts to leverage information about historical news reading behaviours to construct user representations and news representations upon which to base recommendations—with modern approaches typically based on deep learning [22].

For example, the Long Short-Term User Representation (LSTUR) model [23] models user interest by passing a historical sequence of news articles clicked upon by a reader into a long short-term memory neural network. The Neural Attentive Multi-View Learning (NAML) model [24], the Neural Personalized Attention (NPA) approach [25], and the Neural Recommendation Model with Self-Attention (NRMS) [26] learn user representations by modelling news clicked upon by users using an attention mechanism. Dual Attention Networks (DAN) [27] leverage an attention-based convolutional neural network to aggregate a user's interests and an attention-based recurrent neural network to capture the user's click behaviour, and then combines both to make recommendations. The Fine-grained Interest Matching (FIM) model [28] leverages a three-dimensional convolution neural network to mine the user intentions hidden in reading records to reflect the user's fine-grained interests. User-News Matching BERT (UNBERT) [29] and PLM-empowered news recommendations (PLM-empowered) [30] use pre-trained language models to introduce out-of-domain knowledge to enhance the algorithm's ability to capture user interests.

B. MEDIA BIAS

Media bias exists in a variety of forms in news articles. For example, obfuscating by over-reporting, censoring events, cherry-picking facts [31], [32], presenting political ideology in a way that is biased to one side, or ignoring or attacking alternative points of view [33]. Although bias in the news media is discussed and analyzed in academic research [4], [6], [31], [34], [35], many people still consider news articles to be reliable factual reports about events [36]. This trust that readers have in the media may lead to the adoption of biased views, thus allowing the media to have a significant impact on society and public opinion [34], [35], [37], [38]. While complete elimination of bias may be an unrealistic goal, detecting and highlighting instances of media bias can warn readers that content is not balanced, and allow journalists and publishers to evaluate their work objectively [39].

Researchers from the social sciences have a long history of analysing media bias through manual analysis methods [40], [41]. Content analysis [40], [42], [43] is one of the primary methods used to identify and quantify media bias in news texts. In content analysis, researchers collect relevant news data and send them to coders to systematically read and label paragraphs related to media bias in articles, and articles are analysed with these labels. Researchers also use frame analysis [44] to analyse media bias by investigating readers' perceptions of the message conveyed by news articles, and their understanding of how the message is conveyed. In addition, researchers use meta-analysis to infer the factors that cause media bias based on reviewing existing work [33]. All of these methods require significant manual effort and expert domain knowledge, which makes them almost impossible to apply to large-scale news article corpora. However, recent advances in automated media bias detection methods using natural language processing have helped to address this challenge [45].

Computational approaches for detecting media bias in articles have been studied since the work of Lin et al. [46]. Up to that point, media bias had been investigated under different names, including opinion, ideology, authenticity and hyper-partisanship [5]. The most common current approach frames media bias detection as a text classification problem addressed using supervised machine learning approaches applied to an annotated dataset [4], [46], [47]. For example, Recasens et al. [48] considered linguistic bias and developed a word-based logistic regression model that treats bias-inducing words as indicators of a biased article, and all other words as indicators of an unbiased article. They perform a linguistic analysis to find bias inducing words. Jiang et al. [49] built an ELMo-based [50] sentence encoder to predict the biased ideology of an article. Baly et al. [31] employed a pre-trained BERT model [51] to encode news content by averaging the word representations extracted from BERT's last two layers, and build a bias detector based on these encodings.

There are also approaches to automatically detect news bias that are not simply based on text classification. For example, Ogawa et al. [52] leveraged ideas from text mining to propose a stakeholder mining mechanism that identifies news bias by comparing participants described in a news event. Chen et al. [5] argued that feature-based and neural network text classification models only capture low-level lexical information. They designed second-order Gaussian bias distributions to collect biased statements from news items to improve detection effectiveness. Ruan et al. [53] leveraged pseudo-labelling frameworks to filter samples from noisy distant supervision datasets to enhance the performance of bias detection models.

C. ALGORITHMS AMPLIFYING MEDIA BIAS

Recent studies have investigated the complex relationship between news recommendation algorithms and the propagation of media bias, highlighting how digital platforms can influence political polarization. Bakshy et al. [11] revealed that on Facebook, users predominantly engage with news that aligns with their pre-existing beliefs, a phenomenon that contributes to political polarization. Similarly, Barberá et al. [12] observed on Twitter that individuals tend to interact with those sharing similar political ideologies, noting a higher propensity for liberals compared to conservatives to engage in cross-ideological exchanges. Flaxman et al. [13] further explored the role of online platforms in ideological segregation, finding that the algorithms powering social networks and search engines can exacerbate ideological divides, thereby fostering polarization.

Simulation studies [14], [15], [16], [17] have extended this analysis, showing that modern news recommendation algorithms often reinforce users' biases, suggesting more polarized content to those who already exhibit a preference for such information. For example, Liu et al. [14] conducted simulation studies using synthetic users and a curated dataset of over 900K news articles to explore the formation of "filter bubbles" by political news recommendation algorithms. They found that such algorithms can reinforce users' existing beliefs, especially for those with extreme preferences. Ruan et al. [16] utilized a novel simulation framework to generate synthetic user reading histories with varied interests and media bias levels, exploring how personalized news recommendation algorithms respond to, and propagate, media bias over time.

Research Gap: Despite these insights, a direct link between news recommendation algorithms built on top of real-world recommendation datasets and their role in amplifying media bias remains under-explored. Simulation analyses, while valuable, do not fully capture the intricacies of user interaction with news content in actual digital environments. Addressing this research gap, our work delves into the combined impact of news recommendation algorithms and media bias based on a real-world news recommendation dataset.

III. METHODOLOGY

To investigate the influence of media bias on news recommendation systems, we construct observational experiments that examine the proportion of biased articles recommended by recommendation algorithms for user groups with different characteristics. This section describes the methodology used in these experiments.

A. GENERATING NEWS RECOMMENDATIONS

In this section, we briefly define the task of news recommendation, and introduce six news recommendation algorithms used to recommend news articles in our experiments.

Given a user, u, and a set of candidate news articles, N, the task of a news recommendation system is to first rank the items in the candidate news set N according to the probability that the user u will click on them. The top k news articles that are most likely to be clicked upon by user u are then typically presented to the user as a set of recommendations.

Modern news recommender systems are typically based on neural network models that generate representations of users and articles in the candidate set, and use these to generate a set of recommendations. Six state-of-the-art news recommendation algorithms are used in our experiments:

- **NPA [25]**: A personalised attention-based neural news recommendation model that uses both a news article representation and a user representation. NPA uses a convolutional neural network to learn a representation of news titles in a news encoder, and captures user information in a user encoder by modelling user click behaviours through an attention mechanism. During the prediction phase, the network applies the maximum likelihood method to minimise the log-likelihood of the news items clicked by the user.
- **NAML [24]**: NAML is an information representation integration model that utilises different kinds of news information. The main components of NAML are a news module and a user module. The news module employs an attentive multi-view learning method to learn unified representations of news articles from news titles, and news categories and subcategories. The user module applies an attention mechanism to capture behaviour information for user representation learning. The inner product of the vectors generated by the user module and the news module represents the likelihood of the user clicking on a news article.
- **LSTUR [23]**: LSTUR uses a gated recurrent unit (GRU) network that maintains both short-term and long-term user representations. The first step of the LSTUR model uses a long-term representation of the user to initialise the hidden state of the GRU network, and the second step of the model uses the GRU network to capture the short-term behaviour of the user from a click sequence. The news representation learns from news titles. The inner product of the user representation vector and the news representation vector is employed to express the likelihood of the user clicking on the news article.
- **NRMS [26]**: NRMS is a neural news recommendation model that applies multi-head self-attention. The NRMS model includes a news encoder and a user encoder. The news encoder in NRMS uses multi-head self-attention to model news representations from news titles. The user encoder also uses multi-head self-attention to capture user click behaviours. The likelihood of a user clicking on a news item is calculated from the inner product of the user representation and the news representation.
- **FIM [28]**: A neural news recommendation model that matches multiple interests from users' historical behaviour information, FIM uses a hierarchical dilated convolution for learning news representations from news titles, and a stacked dilated convolution to construct multi-level user representations from reading records. FIM then uses a cross-interaction module to output the integrated matching vectors of candidate news and users. The final predicted click score results from the integrated matching vectors transformed by a linear layer.

- **PLM-empowered** [30]: PLM-empowered enhances news representations and user representations by introducing external knowledge from pre-trained language models (PLMs). The result of the inner product between the news representation and the user representation is used to predict the likelihood of users clicking a news article.

These state-of-the-art news recommendation algorithms all use historical behaviour sequence modelling to achieve quality recommendations. In our experiments, all algorithms are trained using the same dataset (described in Section IV) and their recommendation performance is evaluated with standard ranking evaluation metrics used in recommender systems research: area under the curve (*AUC*), mean reciprocal rank (*MRR*), and normalized discounted cumulative gain (*NDCG*).

B. DETECTING MEDIA BIAS

We employ the definition from Kiesel et al. [4] that biased news articles "mimic the form of regular news articles, but are one-sided in the sense that opposing views are either ignored or fiercely attacked". The task of a media bias detector for news articles is to output a judgment on whether or not an article is biased. The bias detector employed in our experiments extends the work from Jiang et al. [49] (the winning solution of Task 4 from SemEval 2019 [4]). Jiang et al. trained their model entirely on a manually annotated dataset. We follow the recommendation of Ruan et al. [53] to augment training by selecting suitable samples from a distant supervision dataset to improve the accuracy of the bias detector. Ruan et al. [53] also establish the generalizability of their proposed approach to unseen datasets [54]. They show that a bias detection model trained on the SemEval-2019 Task 4 Hyperpartisan Dataset [4] can effectively identify bias in the Annotated Data Dataset [54]. We rely on this result to support our use of a bias detection model trained using the SemEval-2019 Task 4 Hyperpartisan Dataset [4] to identify bias in articles in the **MI**crosoft **N**ews **D**ataset [55] (the datasets used are described in Section IV-A).

C. GENERATING USER GROUPS

Users' historical behaviour patterns, which represent historical decisions made by users about which articles to read, are essential when building news recommendation systems. To investigate whether the recommendation algorithms described in Section III-A perform differently depending on the amount of media bias in different users' news reading histories, we divide the users in the study into groups based on their historical behaviour patterns. We use two strategies for this: *bias proportion* and *interest distribution*.

- **Bias Proportion**: For each news article in a user's historical reading record, we leverage the bias detection framework discussed in Section III-B to determine whether it is biased or not. For each user, we then calculate the proportion of articles in their reading history that are biased:
$$\text{prop}_u = \frac{|H_{\text{biased}}|}{|H|} \tag{1}$$
where u is a user and $|H_{\text{biased}}|$ is the number of biased news articles in the user's reading history which contains $|H|$ news articles. We divide users into groups based on this proportion. Figure 1(a) illustrates this.
- **Interest Distribution**: This strategy is designed to remove the influence of the category of news that a reader is consuming from our analysis. It could be the case that some categories are more prone to bias than others and that any influence of bias that we see in the groups created using the bias proportion strategy are in fact due to different interests of users in those groups. The interest distribution strategy leverages category labels that accompany news articles to form user groups. The intuition is that the categories of news read by users reflect their interests. For each user, we calculate a vector of the distribution of article categories across their reading history to represent their interests. Figure 1(b) illustrates this. After generating an interest vector for each user, we use k-means clustering [56] to generate k groups, each of which represents users with similar reading interests.

D. BIAS PROPORTION ON TOP-K

To assess the prevalence of biased news articles in the recommendations made to users from a specific group, we measure the average proportion of biased articles in the top-k recommendations provided to a specific user group. Assume a user group g, a news recommendation algorithm a and a candidate news set N, where $g \in G$, $a \in A$, and G and A denote the set of user groups and news recommendation algorithms. For the i-th user, u_i, in the user group g, we leverage the recommendation algorithm a to select k items from the news candidate set N, and then calculate the proportion of biased articles,prop_{u_i}, in the k selected items. The average recommendation bias proportion in the recommendations generated by algorithm a for the users in user group g is:
$$\text{prop}_{g_a} = \frac{\sum_{i=1}^{|g|} \text{prop}_{u_i}}{|g|} \tag{2}$$
where $|g|$ is the number of users in the user group g.

IV. EXPERIMENTS

In this section we describe the datasets used in our experiments, and the three experiments performed: (1) to evaluate the performance of the news recommendation algorithms (Section IV-B); (2) to evaluate the accuracy of the news bias detector (Section IV-C); and (3) to investigate whether the level of biased media articles promoted by different recommendation algorithms is different for different user groups. The results of these experiments are presented in Section IV-D.

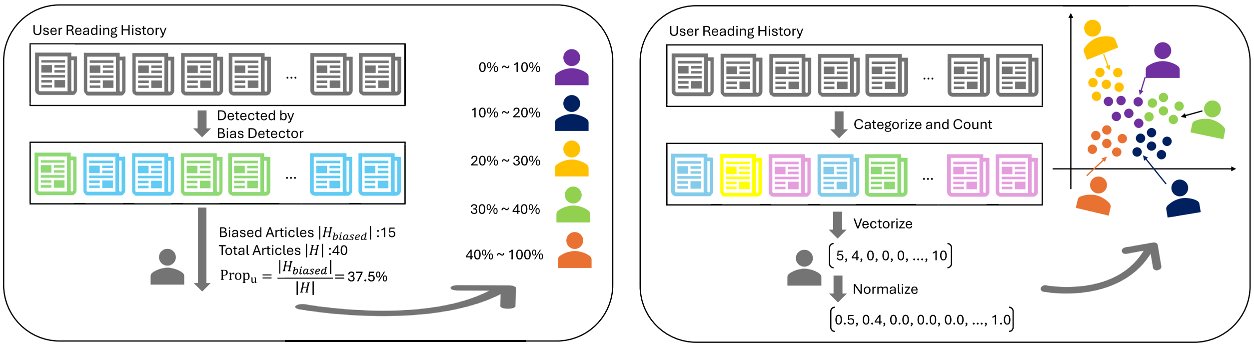

(a) Bias Proportion Strategy (b) Interest Distribution Strategy

FIGURE 1. An illustration of the (a) Bias Proportion and (b) Interest Distribution strategies for dividing users into groups.

A. DATASETS

Our experiments leverage two datasets: the **MI**crosoft **N**ews **D**ataset (MIND) [55] and the SemEval-2019 Task 4 Hyperpartisan Dataset [4].

MIND [55] is a large-scale, English-based news recommendation dataset that includes six weeks of real-world behavioural data from users of the Microsoft News platform.[3] MIND includes two versions, MIND-large and MIND-small. Each version is divided into training, validation, and test sets. We utilised the MIND-large version,[4] which contains 2,232,748 samples in the training set, 376,471 samples in the validation set, and 2,370,727 samples in the test set. Each sample consists of a user's interaction with the recommendation system, including a user ID, the timestamp of interaction, the user's reading history at that time, and a list of news items recommended to the user at that moment with indicators showing which items were clicked. The training set is based on 711,222 users and 101,527 news items, whereas the validation set is based on 255,990 users and 72,023 news items. The test set is based on 702,005 users and 120,961 news items. The training set is employed to train news recommendation algorithms, the validation set is employed to evaluate the performance of news recommendation algorithms, and the test set is employed to select user groups and candidate news articles for our final experiments.

The SemEval-2019 Task 4 Hyperpartisan Dataset [4] was released along with the SemEval-2019 hyperpartisan news detection task. The task is to train machine learning models to automatically detect whether news is biased.[5] The SemEval-2019 Task 4 Hyperpartisan Dataset includes a *by-article* portion of the dataset containing 1,273 manually labelled items and a *by-publisher* portion of the dataset containing 754,000 items labelled through distant supervision (articles are considered biased if their publisher is considered biased). We employ all samples from the by-article portion of the dataset for training the media bias detector, and a small set of news articles from the by-publisher portion of the dataset in a data augmentation strategy (following the approach described by Ruan et al. [53]).

B. EVALUATING NEWS RECOMMENDATION SYSTEMS

In this section, we describe an experiment to evaluate the performance of the news recommendation algorithms described in Section III-A.

1) EXPERIMENTAL SETTINGS

We train six news recommendation models using the algorithms described in Section III-A on the MIND dataset. NPA, NAML, LSTUR and NRMS are implemented using the Microsoft Recommenders open source repository,[6] and the rest are re-implemented using the deep learning framework PyTorch.[7] Hyperparameter tuning of the news recommendation algorithms is based on optimal results on the validation dataset. To evaluate the performance of the algorithms we use AUC, MRR, nDCG@5, and nDCG@10, which are standard ranking metrics for top-k recommendations and are also employed by the MIND [55] for recommendation evaluations.

2) RESULTS

We show the performance of the different news recommendation algorithms in Table 1. Our results show essentially the same, or slightly better, performance as results using the MIND dataset published by other researchers [28], [30], [55], indicating the reliability of the reproduced results. In addition, these well-trained news-specific recommendation algorithms are sufficient to represent the current state-of-the-art in the news recommendation domain—for example, Wu et al. [30] reported that PLM-empowered had been deployed on the Microsoft News platform.

C. EVALUATING MEDIA BIAS DETECTION

In this section, we describe an experiment to evaluate the performance of the media bias detection model described in Section III-B.

[3] https://microsoftnews.msn.com
[4] MIND-large is publicly available at: https://msnews.github.io/
[5] https://pan.webis.de/semeval19/semeval19-web/
[6] https://github.com/microsoft/recommenders
[7] https://pytorch.org/

TABLE 1. Performance of different news recommendation algorithms on the validation set of the MIND dataset.

Algorithms	AUC	MRR	nDCG@5	nDCG@10
NPA	0.6705	0.3182	0.3487	0.4127
NAML	0.6859	0.3313	0.3681	0.4307
LSTUR	0.6802	0.3289	0.3624	0.4258
NRMS	0.6737	0.3230	0.3553	0.4227
FIM	0.6353	0.2970	0.3268	0.3891
PLM-empowered	0.6972	0.3425	0.3793	0.4443

1) EXPERIMENTAL SETTINGS

We employ the model proposed by Jiang et al. [49] as the backbone, and all the hyperparameter values are set using their recommendations. In addition, we follow the overlap-checking approach of Ruan et al. [53] to randomly select pseudo-samples from the samples whose prediction results are consistent with the remote supervision data, which provides more training samples to the backbone. For our empirical analysis, we train this model using the training partition of the SemEval-2019 Task 4 Hyperpartisan dataset, and evaluate its performance using the validation partition. To evaluate the performance of the detection models we use accuracy, precision, recall, and f1 score, which are also the metrics employed in the SemEval-2019 Task 4 Hyperpartisan Detection task [4].

2) RESULTS

Table 2 compares the performance achieved by both the backbone and the approach integrating the data augmentation method. The best performance on the detection framework achieves 0.867 accuracy, 0.015 better than the backbone, indicating that the detection model is capable of accurately detecting media bias in news articles. In addition, the detector's generalisation ability is further tested by Ruan et al. [53] on another human-annotated dataset [54] not used for training the bias detection model, by measuring the correlation between the model outputs and the aggregated human bias scores. They demonstrated a strong correlation between the bias detection model outputs and human-annotated bias scores, indicating the strong generalisation capability of the bias detection model.

We employ the detection framework to calculate the proportion of biased articles in different subsets of the MIND dataset. Although it is true that the outputs of the bias detection model will have some errors, we are confident that its high performance and good generalisation capacity provide reliable indications of bias level when bias proportions are aggregated across sets of articles. While it would be better to train the bias detector on data from the MIND dataset, MIND does not include bias labels so this is not easily achieved. On the other hand, there are no large publicly available datasets that include bias labels that also include user behaviour data. Therefore, we believe that using recommender systems trained using MIND in combination with a bias detector trained using the SemEval-2019 Task 4 Hyperpartisan Dataset is a good compromise in the absence of a news recommendation dataset containing human-annotated bias labels and user behaviour data.

TABLE 2. Media bias detection performance on the SemEval-2019 Task 4 Hyperpartisan Dataset.

Approach	Accuracy	Precision	Recall	F1
Backbone	0.852	0.824	**0.767**	0.780
Overlap-checking	**0.867**	**0.863**	0.760	**0.803**

The number of biased and unbiased articles in each news category in the MIND dataset is shown in Figure 2(a). Biased news accounted for 8.88% of the total news, most of which are concentrated in the news and sports categories, while videos, music, and TV have almost no biased articles. Figure 2(b) illustrates the frequent words in articles from the news and sports categories respectively, presented as word clouds where more frequent words are larger. We can observe that "Trump", "democrat", and "impeachment" are the words that appear most often in the news category. It is unsurprising that "Trump", as a highly topical political figure, commonly appears in biased articles. In the sports category, we observe that the words "game", "patriot", "defense", and "offense" were frequently associated with biased articles. In fierce competitive sports, multiplayer confrontations are inevitable, which makes it easy to understand the biased nature of sports coverage.

D. ASSESSING ALGORITHMIC MEDIA BIAS

In this section we describe experiments that assess the promotion of biased media by different news recommendation algorithms and report the results of each step.

1) SELECTING USER GROUPS

We first extract all users that appear in the MIND test set. The bias detector requires a complete news article (including title and body) as input, and the recommender systems require a reasonable reading history. So we remove incomplete articles[8] and users with less than five complete articles in their reading history. We then calculate the proportion of biased articles in the reading history of each user.

For user groups based on the bias proportion strategy, we group each user according to the proportion of biased articles in their reading history using intervals of (0%, 10%), [10%, 20%), [20%, 30%), [30%, 40%) and [40%, 100%]. We randomly select 2,000 users within each group from the users in the MIND dataset. The average proportion of biased articles in the reading histories of each group are: 6.92%, 13.99%, 23.06%, 33.58%, and 47.16% respectively.

For user groups based on the interest distribution strategy, we first use k-means clustering [56] to find interest clusters of users from the training and validation sets of MIND. Based on preliminary experiments we set $k = 5$. Then we use the trained k-means model to summarize the interests of the users from the MIND test set. The distribution of articles in different categories in the reading histories of users in each of these groups is shown in Figure 3. To illustrate

[8]The body of news articles in the MIND dataset is in the form of a URL, which requires an additional network request to fetch, and some of the fetch requests fail.

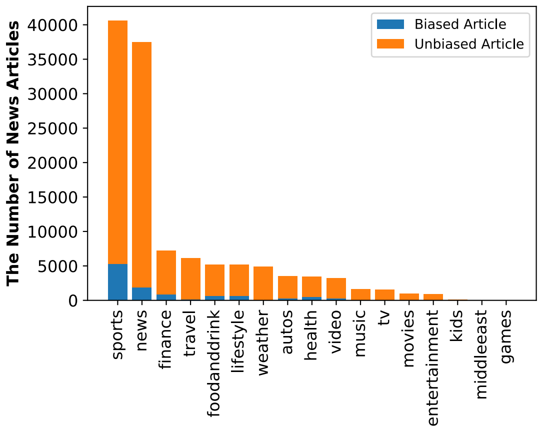

(a) The number of biased and unbiased articles in each category

(b) Word clouds for biased articles in the news and sports categories

FIGURE 2. (a) The number of biased and unbiased articles in each category in the MIND dataset. (b) The upper graph shows words that appear in biased articles in the news category, and the lower graph shows words that appear in biased articles in the sports category (bigger words appear more often).

that these interest groups indeed differentiate users based on news category interests, we visualize the news category embedding and labels predicted by the k-means model using t-SNE [57] in Figure 4. We find that Interest Clusters 2, 3, 4 and 5 can be easily distinguished from each other, with Interest Cluster 2 occupying the lower-left quadrant of the vector space, Interest Cluster 3 appearing in the upper-left quadrant, Interest Cluster 4 concentrated in the upper-right, and Interest Cluster 5 densely occupying the lower-right quadrant. Interest Cluster 1 is less well separated. This is perhaps because, unlike other interest clusters with aggregated reading categories, the users in Interest Cluster 1 read more broadly across multiple categories as shown in Figure 3.

The distributions of bias proportion for users in each interest group are shown in Figures. 5(a) to 5(e). The median of the proportion of biased articles in the historical reading records of these five user groups are 16.67%, 7.14%, 10.00%, 10.00%, and 12.50% respectively.

For each interest cluster, we use the median split method [58] to divide the users into two sub-groups—one with high bias proportion and one with low bias proportion. We randomly select 2,000 users for each interest group in the low bias user sub-group and the high bias user sub-group. The average proportion of biased articles in the historical reading records of users in the low bias interest sub-groups are 5.74%, 0.80%, 2.23%, 2.39%, 4.31% respectively. The average proportion of biased articles in the historical reading records of users in the high bias interest sub-groups are 27.68%, 16.69%, 20.57%, 20.45%, 22.69% respectively.

2) SELECTING CANDIDATE NEWS ARTICLES

To facilitate bias detection, we filtered out news articles that are missing body text from the test set, which meant a total of 113,984 articles remained in the test set. To ensure fairness across all users, we further selected those news items that were completely new to all users in the system for this part of our study. This step helps ensure that our analysis of user responses to different news items is not biased by any prior exposure those items might have had. Finally, we formed a candidate set consisting of 22,283 news items, with an average bias proportion of 9.66%. This set is presented to users as the candidate news set.

3) RESULTS

The purpose of this experiment is to investigate the impact of bias in users' reading histories on the recommendations provided by different recommendation algorithms. To this end, we calculate the proportion of biased news articles recommended to users in the top-k recommendation sets for different recommendation algorithms. We record the average proportion of biased articles read respectively for each of the grouping strategies: *bias proportion-based user groups* and *interest distribution-based user groups*. The average proportion of biased news items in the recommendation sets generated by the NPA, NAML, LSTUR, NRMS, FIM, and PLM-empowered news recommendation algorithms for the five bias proportion user groups for the Top-20, Top-50, and Top-100 recommendation sets are shown in Table 3. The same results for user groups based on the interest grouping strategy are presented in Table 4.

V. DISCUSSION

This section uses the results from the experiments described in the previous section to address the two research questions outlined in Section I.

A. INFLUENCE OF BIASED READING HISTORY ON RECOMMENDER SYSTEMS

RQ1: *Are news recommendation systems, influenced by users' historical reading biases, leading to recommendations of more biased articles?*

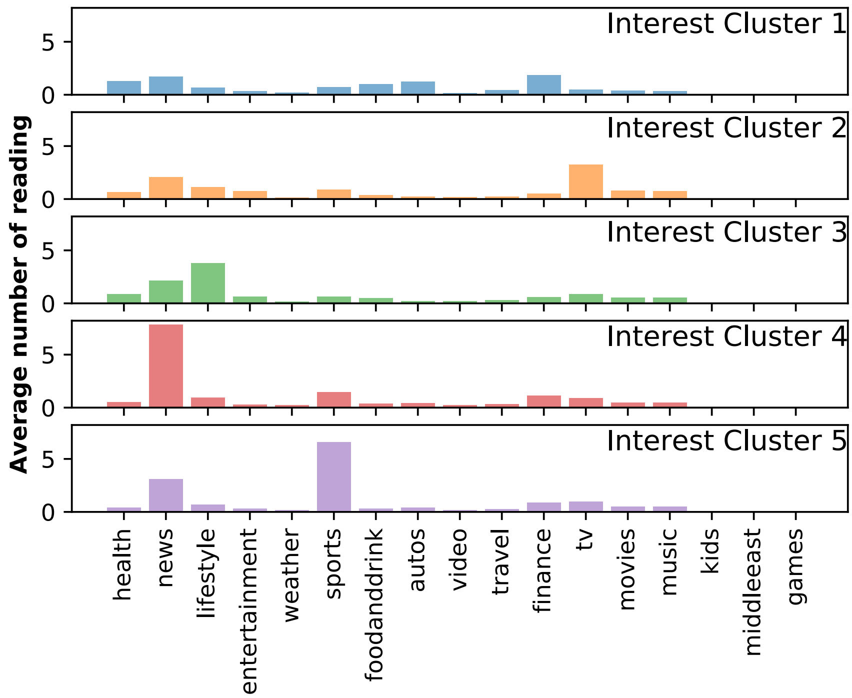

FIGURE 3. The average number of articles read in each news category for each cluster.

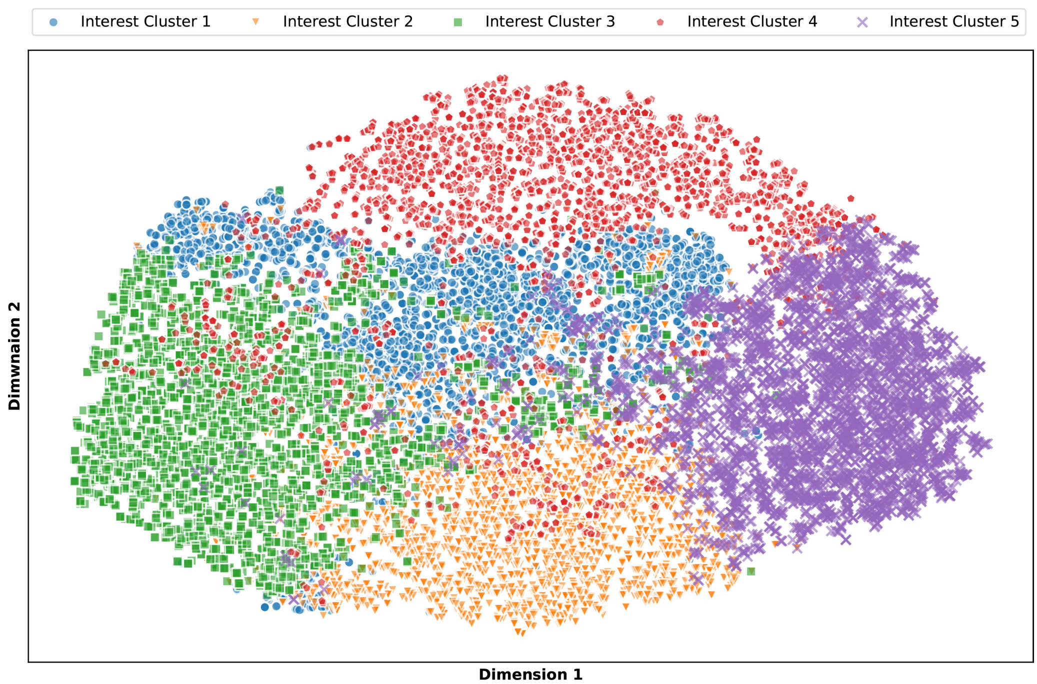

FIGURE 4. A visualisation of the users in the MIND dataset where colours the Interest Cluster to which a user belongs and position is determined by applying to the t-SNE algorithm to the interest distribution for each user.

Based on the results in Tables 3 and 4, we argue that news recommendation algorithms are indeed affected by the amount of bias in users' historical reading records. As users read more biased news articles, news recommendation algorithms will recommend more biased news articles to them.

In Table 3, users are divided into five groups based on the amount of biased news in their reading histories. We can observe that the proportion of biased articles in the recommended article sets generated by different recommendation algorithms (*NPA*, *NAML*, *LSTUR*, *NRMS*, *FIM*, *PLM*) for the Top 20, Top 50 and Top 100 recommendation sets increases

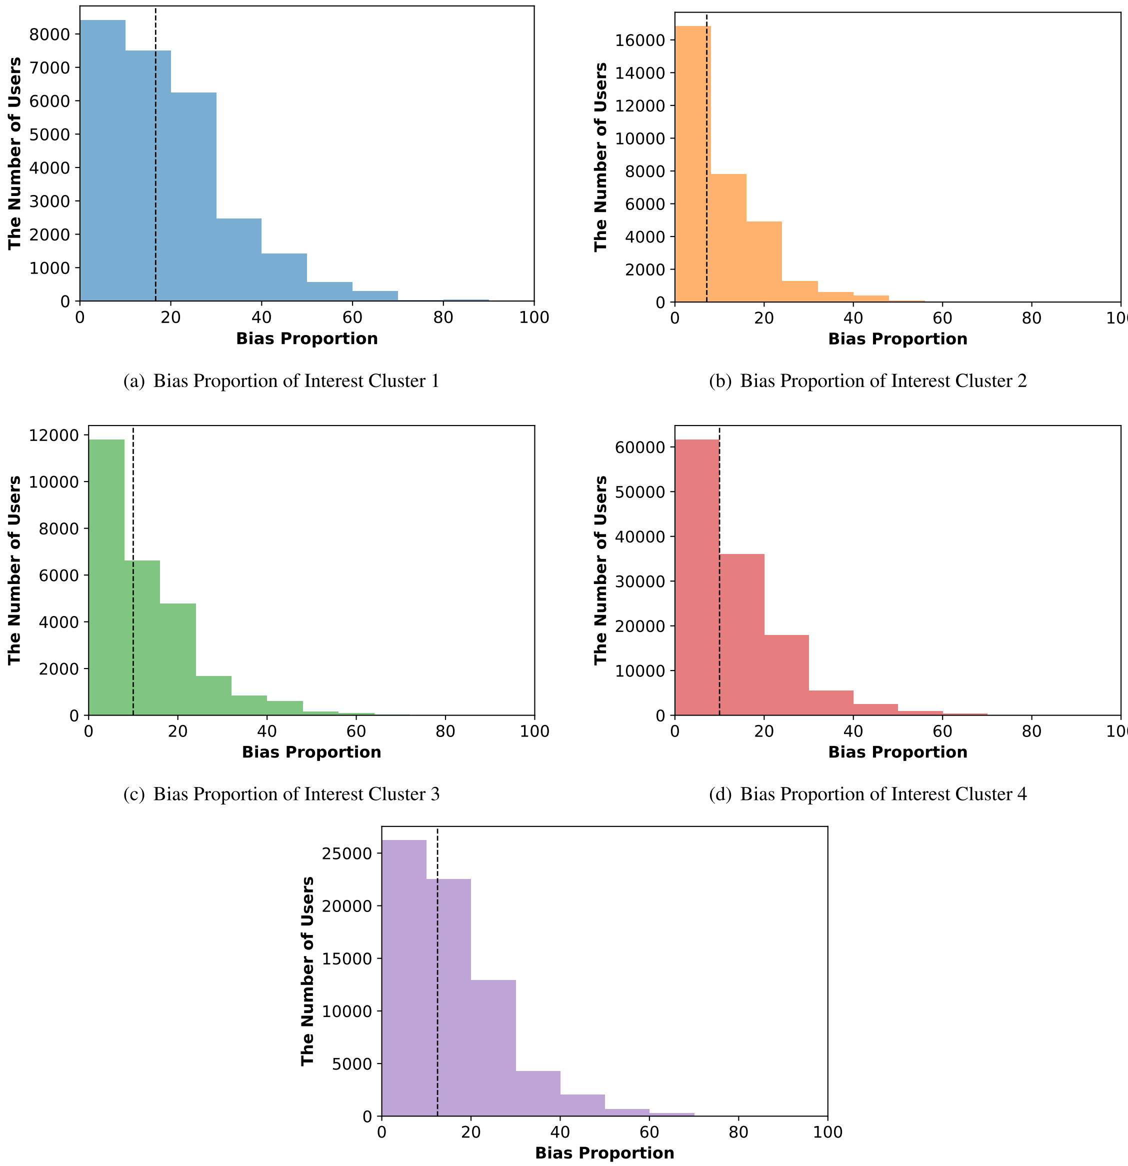

(a) Bias Proportion of Interest Cluster 1

(b) Bias Proportion of Interest Cluster 2

(c) Bias Proportion of Interest Cluster 3

(d) Bias Proportion of Interest Cluster 4

(e) Bias Proportion of Interest Cluster 5

FIGURE 5. The distribution of bias proportion for users in the five interests groups. The black dashed line represents the median bias proportion for each interest cluster.

with the proportion of biased articles in users' reading histories. This data shows a clear trend that users with less bias in their historical reading records are recommended less biased news articles by the news recommendation algorithms than those with more biased articles in their historical reading records.

To remove the influence of the category of news that a reader is consuming from this analysis, we repeat the analysis for each interest group (users have the same reading interests within a group). This is shown in Table 4. In this analysis users are further sub-divided according to the number of biased news articles in their reading histories into a high-bias subgroup and a low-bias subgroup. We can observe that, for the same interest group, all news recommendation algorithms tend to recommend more biased news items to the high-bias subgroup than the low-bias subgroup. This indicates that these recommendation algorithms consider biased attributes of articles when making recommendations, recommending more biased articles to users who have read more biased articles in the past.

In summary, we have strong evidence that sequence-based news recommendation algorithms are influenced by the proportion of biased news in users' reading history and tend to recommend more biased news articles to users who read a lot of biased news articles.

B. RECOMMENDATION ALGORITHM SENSITIVITY TO BIAS

RQ2: *If news recommendation systems are indeed influenced by users' historical reading biases, is the level of influence different for different recommendation algorithms?*

TABLE 3. The proportion of biased articles in the top-k (k = 20, 50, 100) recommendations for different news recommendation algorithms on bias proportion based user groups.

Bias Prop (%).	Top 20						Top 50						Top 100					
	NPA	NAML	LSTUR	NRMS	FIM	PLM.	NPA	NAML	LSTUR	NRMS	FIM	PLM.	NPA	NAML	LSTUR	NRMS	FIM	PLM.
Group 1 (0%, 10%)	8.04	6.64	6.71	19.32	8.84	9.68	7.17	7.02	6.86	14.03	8.00	8.03	6.88	7.17	7.06	11.47	7.60	7.19
Group 2 [10%, 20%)	11.01	9.40	8.67	23.60	13.18	14.40	10.18	9.57	8.87	17.56	11.73	11.87	9.75	9.45	9.06	14.47	10.88	10.65
Group 3 [20%, 30%)	13.51	12.00	10.63	25.75	16.68	18.76	13.07	12.14	10.95	19.48	14.31	15.72	12.64	11.94	11.06	16.31	13.29	14.03
Group 4 [30%, 40%)	16.81	15.71	13.29	26.82	20.27	23.09	16.45	15.51	13.49	20.85	17.39	19.82	16.02	15.04	13.57	17.93	15.75	18.00
Group 5 [40%, 100%)	20.21	18.46	15.49	28.29	22.65	26.01	19.73	18.18	15.69	22.47	19.28	22.64	19.11	17.66	15.38	19.45	17.35	20.66

Note: PLM. stands for PLM-empowered

TABLE 4. The proportion of biased articles in the top-k (k = 20, 50, 100) recommendations for different news recommendation algorithms on interest distribution based user groups.

Bias Prop (%).		Top 20						Top 50						Top 100					
		NPA	NAML	LSTUR	NRMS	FIM	PLM.	NPA	NAML	LSTUR	NRMS	FIM	PLM.	NPA	NAML	LSTUR	NRMS	FIM	PLM.
Interest Cluster 1	Low	9.86	10.14	9.98	28.67	13.76	12.84	10.24	10.37	10.23	22.07	12.30	11.11	10.03	10.10	10.18	17.74	11.59	10.00
	High	15.56	18.71	16.24	32.15	18.99	20.41	15.94	18.07	16.32	25.97	16.66	17.96	15.62	17.38	16.10	22.03	15.20	16.22
Interest Cluster 2	Low	7.62	6.55	7.03	22.21	10.16	8.48	7.84	6.69	6.68	18.11	9.84	7.20	7.35	6.63	6.34	14.93	9.35	6.39
	High	10.94	10.29	10.09	24.66	12.72	12.75	10.58	9.96	9.70	20.50	12.18	10.58	9.76	9.52	9.29	17.37	11.20	9.42
Interest Cluster 3	Low	5.56	6.42	6.09	18.32	9.87	8.53	6.16	6.73	6.97	15.47	10.14	8.39	6.75	7.10	7.37	13.12	10.14	7.73
	High	9.46	14.24	11.57	24.05	14.41	15.88	10.33	13.19	12.16	20.34	13.59	13.91	10.67	12.71	12.30	17.33	12.90	12.53
Interest Cluster 4	Low	6.48	5.92	6.45	20.92	8.37	8.73	5.74	5.23	5.48	15.64	7.08	6.84	5.29	4.83	5.06	12.27	6.61	5.93
	High	13.90	13.89	14.33	26.54	16.65	18.57	13.15	12.25	12.42	21.51	13.74	15.27	12.22	11.36	11.52	18.11	12.34	13.72
Interest Cluster 5	Low	10.22	10.45	5.26	34.88	9.79	12.78	8.42	9.74	5.38	23.93	9.03	10.50	8.27	9.02	5.69	18.70	8.82	9.33
	High	11.71	14.43	6.65	37.46	13.77	17.32	10.27	12.68	7.01	26.31	12.51	14.13	10.46	11.68	7.69	21.24	11.92	12.67

Note: PLM. stands for PLM-empowered

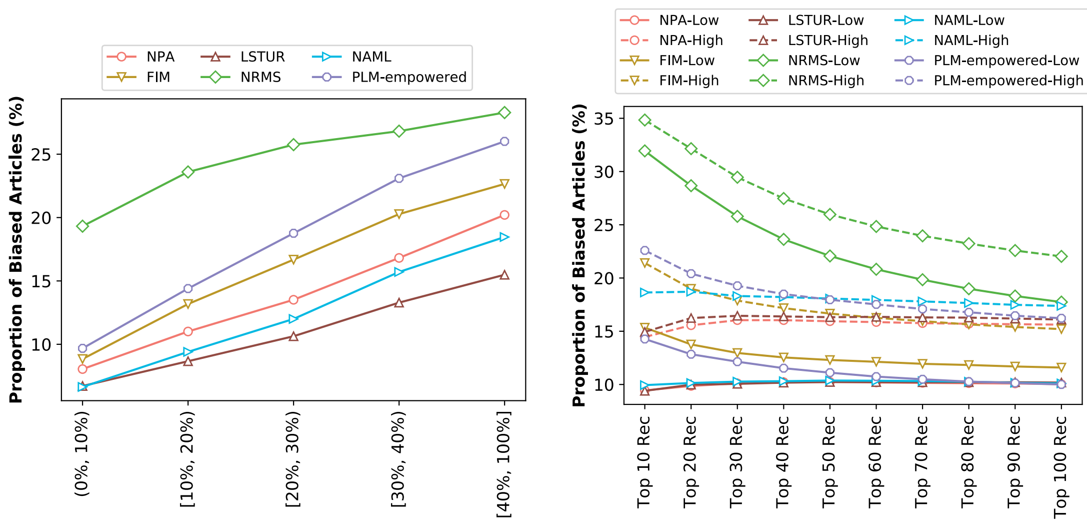

(a) Bias Proportion on Top 20 (b) Bias Proportion Trending from Top 10 to Top 100

FIGURE 6. (a) Biased articles as a percentage of the top 20 recommended set for user groups based on bias proportions for six news recommendation algorithms. (b) Biased articles as a percentage of top k (k from 10 to 100 in increments of 10 on the horizontal axis) recommended set for six news recommendation algorithms, for low-bias and high-bias user clusters from Interest Cluster 1.

Regarding the extent to which different news recommendation algorithms are sensitive to algorithmic media bias, we first observe that the news recommendation algorithms investigated are all sensitive to the number of biased articles in a user's reading records. Figure 6 shows the proportion of biased articles in the top-20 articles recommended for users in each bias proportion user group for different recommendation algorithms. All news recommendation algorithms are sensitive to the amount of bias in users' reading records: the more biased articles are read by users, the more biased articles are presented in the top-20 recommendation sets.

For the same user group, however, recommendation algorithms differ in their sensitivity to bias. The media bias sensitivity of NRMS is significantly higher than other algorithms among all user groups. The other algorithms vary in their sensitivity to the amount of bias in users' reading histories. For users with few biased articles in their reading history, the sensitivity of each algorithm to media bias is not much different. For users with many biased articles in their reading history, PLM-empowered and FIM are more sensitive than NPA, NAML, and LSTUR in capturing bias factors.

The second observation is that recommendation algorithms are sensitive to ranking biased articles from the candidate news set. Figure 6(b) shows the proportion of biased news recommended by news recommendation algorithms in the top k (where k ranges from 10 to 100 in steps of 10) news items that best match users in the high-bias and low-bias user groups from Interest Cluster 1. As k rises, we can see that the proportion of biased news articles recommended by the NRMS, FIM and PLM-empowered algorithms tends to decrease, with the NRMS algorithm showing the sharpest decrease. LSTUR and NPA algorithms have a small increase in the proportion of biased articles as the number of recommendations increases. The proportion of biased articles recommended by the NAML algorithm fluctuated slightly up and down, showing a trend of uniform change.

In the ranking-based recommendation task, the recommendation algorithm tries to give the set of articles from the candidate set that best matches the user's interests the highest rank. The results of this experiment show that the NRMS, FIM and PLM-empowered algorithms tend to emphasise the bias in articles as a user preference more than the other algorithms. When we show limited news to users, NRMS, FIM and PLM-empowered become even more likely to recommend biased articles than others.

C. EXPLAINING BIAS SENSITIVITY IN RECOMMENDATION ALGORITHMS

In this section, we provide some insights into the impact of media bias on news recommendation algorithms, specifically to understand why some algorithms are more sensitive to bias than others. NRMS, NAML, PLM-empowered, and LSTUR use the dot product to calculate the similarity between the candidate news vector and the user vector when making recommendations. The higher the similarity score between the two vectors, the more confidently the algorithm matches the candidate news to the user's interests. Among these recommendation algorithms, the candidate news selected by NRMS is most affected by the amount of biased information in the user sequence, followed by PLM-empowered and NAML. LSTUR is the least affected.

The reason that NRMS is most affected might be the characteristics of its network structure. It uses multi-head self-attention to find the relatedness of clicked news at different positions in the reading records, which promotes the generated user vectors to strengthen the modelling of similar parts in different news articles. When the reading records are full of heavily biased text, the user vector captures this part of the similarity and then promotes news articles containing biased text to be selected through the dot product calculation.

The operation of the PLM-empowered and NRMS algorithms is similar. The difference is that PLM-empowered introduces a pre-trained language model to enhance user vectors, which provides rich external knowledge to break the limitation of over-modelling the relatedness between reading records. This may be a reason for the reduced influence of bias on the recommendations made by this algorithm. The user vectors generated by NAML enhance information diversity by including article category and subcategory information in addition to text information. So, candidate news similarity to user vectors is due to article attributes rather than just text information, which reduces the dependence on text information.

The user sequence modelled by LSTUR uses a recurrent neural network to give more weight to the latest clicked news. The result of this is that generated user vectors place more emphasis on information from more recently clicked news articles than information from older articles. This has the impact of reducing the dependence on global bias information.

In the FIM and NPA algorithms, the likelihood that a user will click on an article is generated by a prediction module that takes a concatenation of the user representations and candidate news representations as input, rather than using a simple dot product. The results of our experiments suggest that the tendency of these algorithms to recommend more biased articles is jointly affected by user and candidate news modelling. This differs from the dot product prediction method discussed above, which is more influenced by user modelling.

VI. CONCLUSION

To address the research gap regarding the effects of media bias on news recommendation algorithms, in this work we investigate the impact of media bias on a variety of news recommendation algorithms (*NPA*, *NAML*, *LSTUR*, *NRMS*, *FIM*, and *PLM-empowered*) using state-of-the-art media bias detection technology and evaluation experiments based on the well-known news recommendation dataset MIND [55]. The results of our experiments show that today's news recommendation algorithms are indeed affected by media bias, i.e., the more biased news articles a user previously read, the more likely news recommendation algorithms are to suggest further biased articles. Our study also shows that differences in how they are implemented mean that different recommendation algorithms vary in the degree to which they are affected by media bias. Among the algorithms, we found that NRMS is the most sensitive to bias, whereas LSTUR is the least affected, despite all algorithms exhibiting similar recommendation performance.

The harmful effects of media bias on users have been extensively studied [6], [7], [39]. Our findings suggest that current state-of-the-art news recommendation algorithms, designed to offer personalised content, are exacerbating the propagation of media bias. This work can serve as a reference for analysing whether the proposed algorithms inadvertently strengthen the dissemination of media bias to users. In future work, apart from continuing to assess accuracy-based algorithms (the most common design paradigm), we will expand our assessments to include algorithms designed based on alternative metrics, such as fairness [59] and diversity [60]. This expansion will enable us to better understand how algorithms crafted to enhance user satisfaction affect the dissemination of biased news. To address the scarcity of publicly available datasets containing both user behaviours and bias annotations, we will annotate a subset of news articles from the MIND dataset for bias. Furthermore, We will also investigate potential adjustments to news recommendation algorithms aimed at bursting the media bias bubble, thereby providing users with a more comprehensive, fair, and diverse range of information to enhance user satisfaction and trust in the news recommendations.

ACKNOWLEDGMENT
This publication has emanated from research conducted with the financial support of Science Foundation Ireland under Grant number 18/CRT/6183 and 12/RC/2289_P2. For the purpose of Open Access, the author has applied a CC BY public copyright licence to any Author Accepted Manuscript version arising from this submission.

REFERENCES
[1] S. J. Ward, "Inventing objectivity: New philosophical foundations," in *Journalism Ethics: A Philosophical Approach*. New York, NY, USA: Oxford Univ. Press, 2010.
[2] E. Sirotkina, "How biased media generate support for the ruling authorities: Causal mediation analysis of evidence from Russia," *Eur. J. Commun.*, vol. 36, no. 2, pp. 183–200, Apr. 2021.
[3] L. L. Kaid and C. Holtz-Bacha, *Encyclopedia of Political Communication*. Newbury Park, CA, USA: SAGE, 2007.
[4] J. Kiesel, M. Mestre, R. Shukla, E. Vincent, P. Adineh, D. Corney, B. Stein, and M. Potthast, "SemEval-2019 task 4: Hyperpartisan news detection," in *Proc. 13th Int. Workshop Semantic Eval.*, 2019, pp. 829–839.
[5] W.-F. Chen, K. Al Khatib, B. Stein, and H. Wachsmuth, "Detecting media bias in news articles using Gaussian bias distributions," in *Proc. Findings Assoc. Comput. Linguistics, EMNLP*, 2020, pp. 4290–4300.
[6] D. P. Baron, "Persistent media bias," *J. Public Econ.*, vol. 90, nos. 1–2, pp. 1–36, Jan. 2006.
[7] D. Bernhardt, S. Krasa, and M. Polborn, "Political polarization and the electoral effects of media bias," *J. Public Econ.*, vol. 92, nos. 5–6, pp. 1092–1104, Jun. 2008.
[8] R. Paul and L. Elder, *The Thinker's Guide for Conscientious Citizens on How to Detect Media Bias and Propaganda in National and World News: Based on Critical Thinking Concepts and Tools*. Lanham, MD, USA: Rowman & Littlefield, 2019.
[9] A. S. Das, M. Datar, A. Garg, and S. Rajaram, "Google news personalization: Scalable online collaborative filtering," in *Proc. 16th Int. Conf. World Wide Web*, May 2007, pp. 271–280.
[10] C. Wu, F. Wu, Y. Huang, and X. Xie, "Personalized news recommendation: Methods and challenges," *ACM Trans. Inf. Syst.*, vol. 41, no. 1, pp. 1–50, Jan. 2023.
[11] E. Bakshy, S. Messing, and L. A. Adamic, "Exposure to ideologically diverse news and opinion on Facebook," *Science*, vol. 348, no. 6239, pp. 1130–1132, Jun. 2015.
[12] P. Barberá, J. T. Jost, J. Nagler, J. A. Tucker, and R. Bonneau, "Tweeting from left to right: Is online political communication more than an echo chamber?" *Psychol. Sci.*, vol. 26, no. 10, pp. 1531–1542, Oct. 2015.
[13] S. Flaxman, S. Goel, and J. M. Rao, "Filter bubbles, echo chambers, and online news consumption," *Public Opinion Quart.*, vol. 80, no. S1, pp. 298–320, 2016.
[14] P. Liu, K. Shivaram, A. Culotta, M. A. Shapiro, and M. Bilgic, "The interaction between political typology and filter bubbles in news recommendation algorithms," in *Proc. Web Conf.*, Apr. 2021, pp. 3791–3801.
[15] H. Zhang, Z. Zhu, and J. Caverlee, "Evolution of filter bubbles and polarization in news recommendation," in *Proc. 45th Eur. Conf. Inf. Retr.*, Dublin, Ireland. Cham, Switzerland: Springer, 2023, pp. 685–693.
[16] Q. Ruan, B. Mac Namee, and R. Dong, "Unveiling the relationship between news recommendation algorithms and media bias: A simulation-based analysis of the evolution of bias prevalence," in *Proc. Int. Conf. Innov. Techn. Appl. Artif. Intell.* Cham, Switzerland: Springer, 2023, pp. 210–215.
[17] A. Kasirzadeh and C. Evans, "User tampering in reinforcement learning recommender systems," in *Proc. AAAI/ACM Conf. AI, Ethics, Soc.*, Aug. 2023, pp. 58–69.
[18] M. Potthast, J. Kiesel, K. Reinartz, J. Bevendorff, and B. Stein, "A stylometric inquiry into hyperpartisan and fake news," in *Proc. 56th Annu. Meeting Assoc. Comput. Linguistics*, 2018, pp. 231–240.
[19] Q. Chen, H. Zhao, W. Li, P. Huang, and W. Ou, "Behavior sequence transformer for e-commerce recommendation in Alibaba," in *Proc. 1st Int. Workshop Deep Learn. Pract. High-Dimensional Sparse Data*, Aug. 2019, pp. 1–4.
[20] W. Yuan, H. Wang, X. Yu, N. Liu, and Z. Li, "Attention-based context-aware sequential recommendation model," *Inf. Sci.*, vol. 510, pp. 122–134, Feb. 2020.
[21] B. Hidasi, A. Karatzoglou, L. Baltrunas, and D. Tikk, "Session-based recommendations with recurrent neural networks," in *Proc. 4th Int. Conf. Learn. Represent.*, San Juan, Puerto Rico, 2016. [Online]. Available: http://arxiv.org/abs/1511.06939
[22] H. Fang, D. Zhang, Y. Shu, and G. Guo, "Deep learning for sequential recommendation: Algorithms, influential factors, and evaluations," *ACM Trans. Inf. Syst.*, vol. 39, no. 1, pp. 1–42, Jan. 2021.
[23] M. An, F. Wu, C. Wu, K. Zhang, Z. Liu, and X. Xie, "Neural news recommendation with long- and short-term user representations," in *Proc. 57th Annu. Meeting Assoc. Comput. Linguistics*, 2019, pp. 336–345.
[24] C. Wu, F. Wu, M. An, J. Huang, Y. Huang, and X. Xie, "Neural news recommendation with attentive multi-view learning," in *Proc. 28th Int. Joint Conf. Artif. Intell.*, Jul. 2019, pp. 3863–3869.
[25] C. Wu, F. Wu, M. An, J. Huang, Y. Huang, and X. Xie, "NPA: Neural news recommendation with personalized attention," in *Proc. 25th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, Jul. 2019, pp. 2576–2584.
[26] C. Wu, F. Wu, S. Ge, T. Qi, Y. Huang, and X. Xie, "Neural news recommendation with multi-head self-attention," in *Proc. Conf. Empirical Methods Natural Lang. Process. 9th Int. Joint Conf. Natural Lang. Process. (EMNLP-IJCNLP)*, 2019, pp. 6389–6394.
[27] Q. Zhu, X. Zhou, Z. Song, J. Tan, and L. Guo, "DAN: Deep attention neural network for news recommendation," in *Proc. AAAI Conf. Artif. Intell.*, Jul. 2019, vol. 33, no. 1, pp. 5973–5980.
[28] H. Wang, F. Wu, Z. Liu, and X. Xie, "Fine-grained interest matching for neural news recommendation," in *Proc. 58th Annu. Meeting Assoc. Comput. Linguistics*, 2020, pp. 836–845.
[29] Q. Zhang, J. Li, Q. Jia, C. Wang, J. Zhu, Z. Wang, and X. He, "UNBERT: User-news matching BERT for news recommendation," in *Proc. 30th Int. Joint Conf. Artif. Intell. (IJCAI)*, Aug. 2021, pp. 3356–3362.
[30] C. Wu, F. Wu, T. Qi, and Y. Huang, "Empowering news recommendation with pre-trained language models," in *Proc. 44th Int. ACM SIGIR Conf. Res. Develop. Inf. Retr.*, Jul. 2021, pp. 1652–1656.
[31] R. Baly, G. Da San Martino, J. Glass, and P. Nakov, "We can detect your bias: Predicting the political ideology of news articles," in *Proc. Conf. Empirical Methods Natural Lang. Process. (EMNLP)*, 2020, pp. 4982–4991.
[32] G. D. S. Martino, S. Cresci, A. Barrón-Cedeño, S. Yu, R. D. Pietro, and P. Nakov, "A survey on computational propaganda detection," in *Proc. 29th Int. Joint Conf. Artif. Intell.*, Jul. 2020, pp. 4826–4832.
[33] F. Hamborg, K. Donnay, and B. Gipp, "Automated identification of media bias in news articles: An interdisciplinary literature review," *Int. J. Digit. Libraries*, vol. 20, no. 4, pp. 391–415, Dec. 2019.
[34] S. DellaVigna and E. Kaplan, "The fox news effect: Media bias and voting," *Quart. J. Econ.*, vol. 122, no. 3, pp. 1187–1234, Aug. 2007.

[35] S. Iyengar and K. S. Hahn, "Red media, blue media: Evidence of ideological selectivity in media use," *J. Commun.*, vol. 59, no. 1, pp. 19–39, Mar. 2009.
[36] S. Wolton, "Are biased media bad for democracy?" *Amer. J. Political Sci.*, vol. 63, no. 3, pp. 548–562, Jul. 2019.
[37] M. Prior, "Media and political polarization," *Annu. Rev. Political Sci.*, vol. 16, pp. 101–127, Jan. 2013.
[38] J. Dunaway and D. A. Graber, *Mass Media and American Politics*. Washington, DC, USA: CQ Press, 2022.
[39] T. Spinde, L. Rudnitckaia, J. Mitrović, F. Hamborg, M. Granitzer, B. Gipp, and K. Donnay, "Automated identification of bias inducing words in news articles using linguistic and context-oriented features," *Inf. Process. Manage.*, vol. 58, no. 3, May 2021, Art. no. 102505.
[40] T. Groseclose and J. Milyo, "A social-science perspective on media bias," *Crit. Rev.*, vol. 17, nos. 3–4, pp. 305–314, 2005.
[41] T. Groeling, "Media bias by the numbers: Challenges and opportunities in the empirical study of partisan news," *Annu. Rev. Political Sci.*, vol. 16, no. 1, pp. 129–151, May 2013.
[42] C. Budak, S. Goel, and J. M. Rao, "Fair and balanced? Quantifying media bias through crowdsourced content analysis," *Public Opinion Quart.*, vol. 80, no. S1, pp. 250–271, 2016.
[43] J. McCarthy, L. Titarenko, C. McPhail, P. Rafail, and B. Augustyn, "Assessing stability in the patterns of selection bias in newspaper coverage of protest during the transition from communism in Belarus," *Mobilization, Int. Quart.*, vol. 13, no. 2, pp. 127–146, Jun. 2008.
[44] R. M. Entman, "Framing: Towards clarification of a fractured paradigm," *J. Commun.*, vol. 43, pp. 51–58, 1993.
[45] D. D'Alessio, "Media bias in presidential elections: A meta-analysis," *J. Commun.*, vol. 50, no. 4, pp. 133–156, Dec. 2000.
[46] W.-H. Lin, T. Wilson, J. Wiebe, and A. Hauptmann, "Which side are you on? Identifying perspectives at the document and sentence levels," in *Proc. 10th Conf. Comput. Natural Lang. Learn.*, 2006, pp. 109–116.
[47] H. Rashkin, E. Choi, J. Y. Jang, S. Volkova, and Y. Choi, "Truth of varying shades: Analyzing language in fake news and political fact-checking," in *Proc. Conf. Empirical Methods Natural Lang. Process.*, 2017, pp. 2931–2937.
[48] M. Recasens, C. Danescu-Niculescu-Mizil, and D. Jurafsky, "Linguistic models for analyzing and detecting biased language," in *Proc. 51st Annu. Meeting Assoc. Comput. Linguistics*, 2013, pp. 1650–1659.
[49] Y. Jiang, J. Petrak, X. Song, K. Bontcheva, and D. Maynard, "Team Bertha von Suttner at SemEval-2019 task 4: Hyperpartisan news detection using ELMo sentence representation convolutional network," in *Proc. 13th Int. Workshop Semantic Eval.*, 2019, pp. 840–844.
[50] M. Peters, M. Neumann, M. Iyyer, M. Gardner, C. Clark, K. Lee, and L. Zettlemoyer, "Deep contextualized word representations," in *Proc. Conf. North Amer. Chapter Assoc. Comput. Linguistics, Hum. Lang. Technol.*, 2018, pp. 2227–2237.
[51] J. D. M.-W. C. Kenton and L. K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," in *Proc. NAACL-HLT*, 2019, pp. 4171–4186.
[52] T. Ogawa, Q. Ma, and M. Yoshikawa, "News bias analysis based on stakeholder mining," *IEICE Trans. Inf. Syst.*, vol. E94, no. 3, pp. 578–586, 2011.
[53] Q. Ruan, B. Mac Namee, and R. Dong, "Bias bubbles: Using semi-supervised learning to measure how many biased news articles are around us," in *Proc. AICS*, 2021, pp. 153–164.
[54] S. Lim, A. Jatowt, M. Färber, and M. Yoshikawa, "Annotating and analyzing biased sentences in news articles using crowdsourcing," in *Proc. 12th Lang. Resour. Eval. Conf.*, 2020, pp. 1478–1484.
[55] F. Wu, Y. Qiao, J.-H. Chen, C. Wu, T. Qi, J. Lian, D. Liu, X. Xie, J. Gao, W. Wu, and M. Zhou, "MIND: A large-scale dataset for news recommendation," in *Proc. 58th Annu. Meeting Assoc. Comput. Linguistics*, 2020, pp. 3597–3606.
[56] J. D. Kelleher, B. M. Namee, and A. D'Arcy, *Fundamentals of Machine Learning for Predictive Data Analytics: Algorithms, Worked Examples, and Case Studies*. Cambridge, MA, USA: MIT Press, 2020.
[57] L. Van der Maaten and G. Hinton, "Visualizing data using t-SNE," *J. Mach. Learn. Res.*, vol. 9, no. 11, pp. 2579–2605, 2008.
[58] D. Iacobucci, S. S. Posavac, F. R. Kardes, M. J. Schneider, and D. L. Popovich, "The median split: Robust, refined, and revived," *J. Consum. Psychol.*, vol. 25, no. 4, pp. 690–704, Oct. 2015.
[59] E. Pitoura, K. Stefanidis, and G. Koutrika, "Fairness in rankings and recommendations: An overview," *VLDB J.*, vol. 31, no. 3, pp. 431–458, May 2022.
[60] P. Castells, N. Hurley, and S. Vargas, "Novelty and diversity in recommender systems," in *Recommender Systems Handbook*. Cham, Switzerland: Springer, 2021, pp. 603–646.

QIN RUAN received the M.S. degree in computer science from University College Dublin, Dublin, Ireland, in 2019, where she is currently pursuing the Ph.D. degree in computer science. Her research interests include natural language processing and news recommender systems.

BRIAN MAC NAMEE received the Ph.D. degree in computer science from Trinity College Dublin, Ireland, in 2004. After a period working in industry and with Dublin Institute of Technology, he joined University College Dublin, Ireland, in 2015, where he is currently an Associate Professor and the Co-Director of the Science Foundation Ireland Centre for Research Training in Machine Learning and the Insight Science Foundation Ireland Centre for Data Analytics Research. He has coauthored the textbook *Fundamentals of Machine Learning for Predictive Data Analytics: Algorithms, Worked Examples and Case Studies* (MIT Press, 2015 and 2020). His research interests include machine learning and in particular novelty detection, human-in-the-loop machine learning, data visualization for machine learning, and machine learning applications in domains, such as medicine, agriculture, and space.

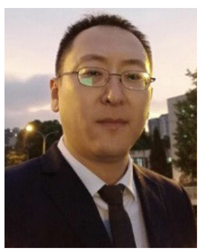

RUIHAI DONG (Member, IEEE) received the Ph.D. degree in computer science from University College Dublin, Dublin, Ireland, in 2015. He is currently an Assistant Professor with the School of Computer Science, University College Dublin. He has a track record of collaboration with industry and has worked with companies, including Eagle Alpha, SkillPages, and Samsung, as well as individually winning funding from Enterprise Ireland for commercialization of his research. He has published in top peer-reviewed journals and conferences, such as WWW, RECSYS, IUI, and IJCAI. His research interests include machine learning and deep learning and their applications in recommender systems and finance. He was awarded the Outstanding Research Award 2018 by the UCD School of Computer Science.