

How to Effectively Identify and Communicate Person-Targeting Media Bias in Daily News Consumption?

Felix Hamborg^{1,2,3}, Timo Spinde¹, Kim Heinser¹, Karsten Donnay^{2,3} and Bela Gipp^{2,4}

¹University of Konstanz, Konstanz, Germany

²Heidelberg Academy of Sciences and Humanities, Heidelberg, Germany

³University of Zurich, Zurich, Switzerland

⁴University of Wuppertal, Wuppertal, Germany

Abstract

Slanted news coverage strongly affects public opinion. This is especially true for coverage on politics and related issues, where studies have shown that bias in the news may influence elections and other collective decisions. Due to its viable importance, news coverage has long been studied in the social sciences, resulting in comprehensive models to describe it and effective yet costly methods to analyze it, such as content analysis. We present an in-progress system for news recommendation that is the first to automate the manual procedure of content analysis to reveal person-targeting biases in news articles reporting on policy issues. In a large-scale user study, we find very promising results regarding this interdisciplinary research direction. Our recommender detects and reveals substantial frames that are actually present in individual news articles. In contrast, prior work rather only facilitates the visibility of biases, e.g., by distinguishing left- and right-wing outlets. Further, our study shows that recommending news articles that differently frame an event significantly improves respondents' awareness of bias.

Keywords

news bias, conjoint experiment, Google News, news aggregator, bias perception, polarization

1. Introduction

How topics are covered in the news frames public debates and profoundly impacts collective decision-making, such as during elections [1, 2]. News may be subtly biased through various forms, such as word choice, framing, intentional omission or misrepresentation of specific details [3]. In extreme cases, “fake news” may present entirely fabricated facts to intentionally manipulate public opinion toward a given topic. A rich diversity of opinions is desirable, but systematically biased information can be problematic as a basis for decision-making if not


INRA'21: 9th International Workshop on News Recommendation and Analytics, September 25, 2021, Amsterdam, Netherlands


✉ felix.hamborg@uni-konstanz.de (F. Hamborg); timo.spinde@uni-konstanz.de (T. Spinde);

kim.heinser@uni-konstanz.de (K. Heinser); donnay@ipz.uzh.ch (K. Donnay); gipp@uni-wuppertal.de (B. Gipp)

🌐 <https://felix.hamborg.eu/> (F. Hamborg); <https://www.karstendonnay.net/> (K. Donnay); <https://www.gipp.com/> (B. Gipp)

🆔 0000-0003-2444-8056 (F. Hamborg); 0000-0002-9080-6539 (K. Donnay); 0000-0001-6522-3019 (B. Gipp)

 © 2021 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

 CEUR Workshop Proceedings (CEUR-WS.org)

recognized as such. Therefore, it is crucial to empower newsreaders in recognizing relative biases in coverage.

In this paper, we thus seek to answer the following research question. “How can we effectively communicate instances of media bias in a set of news articles reporting on the same political event?” Instead of identifying and then communicating each of the various forms of bias individually, we focus on person-oriented polarity, which is a fundamental effect resulting from various bias forms [4]. While the problem statement misses bias not related to persons, coverage on policy issues is largely person-oriented, e.g., because decisions are made by politicians or affect individuals in society. Our key contributions—especially when comparing to prior work—are: (1) We present the first system that is able to identify even subtle forms of media bias affecting the perception of persons by imitating the manual content analysis procedure from the social science. (2) We present modular visualizations to communicate media bias during daily news consumption. (3) We conduct a large scale user study using conjoint analysis to measure effectiveness of our analysis, visualizations, and individual components therein.

We publish the survey materials, including questionnaires and anonymized answers, articles, and visualizations freely at: <https://doi.org/10.5281/zenodo.5517401>

2. Related Work

Media bias has been long studied in the social sciences, resulting in a comprehensive set of models to describe it, such as political framing [5] and the news production process defining causes, forms, and effects of bias [3], and effective methods to analyze it [3]. Established methods, such as content analysis and frame analysis [6], typically include systematic reading and labeling of texts. Despite their high effectiveness and reliability, they are largely conducted manually and do not scale with the vast amount of news. In contrast, many methods in computer science concerned with media bias employ automated and thus more efficient approaches but yield non-optimal results [3], e.g., because they treat bias as only vaguely defined “topic diversity” [7] or “differences in [news] coverage” [8]. Though, methods for the identification of biased words exist for other domains, such as Wikipedia articles [9].

Helping news consumers to become aware of media bias is an effective means to mitigate the negative effects of slanted news coverage, such as polarization [8, 10]. Moreover, most studies find that automated approaches concerned with the communication of biases in the news can successfully increase bias-awareness in news consumers [11, 12, 13, 14]. However, previous approaches suffer from at least one of the following shortcomings. First, researchers cannot quantitatively pinpoint which individual components facilitate bias-awareness [11, 12, 13, 14]. Instead, studies measure overall effectiveness of analysis or visualizations. Second, approaches are concerned with the identification or communication of biases only in titles [11], on the article-level [8, 14], or outlet-level [15]. Considering only the overall article or even properties of its publisher, may lead to incorrect classification or missing instances of bias, e.g., that affect readers’ perception on the sentence-level.

In sum, many approaches effectively communicate biases to users and most studies indicate how society benefits from doing so. However, many approaches identify only vaguely defined or superficial biases, e.g., because they do not use the established, effective models and analyses.

Further, none of the reviewed approaches narrows down effectiveness regarding change in bias-awareness to individual analysis and visualization components. Lastly, to our knowledge, no approach identifies biases on the sentence level in news articles.

3. System

Given a set of news articles reporting on the same political event, our system’s analysis aims to find groups of articles that frame the event similarly using four analysis tasks. These groups are later visualized (Section 4) to enable non-expert news consumers to quickly get a bias-sensitive synopsis of a given news event.

For (1) *article gathering*, we extract news articles reporting on one event [16], currently for a set of user-defined URLs, or by providing texts to the system. We then perform state-of-the-art NLP (2) *preprocessing* using Stanford CoreNLP. (3) *Target concept analysis* finds and resolves person mentions across the topic’s articles, including also broadly defined and event-specific coreferences that are otherwise non-coreferential or even opposing, such as “freedom fighters” and “terrorists” [3].

(4) *Frame identification* determines how articles portray persons and then groups those articles that similarly portray (or frame) the persons. This task centers around political framing [5], where a frame represents a specific perspective on an issue, e.g., which aspects are highlighted when reporting on the issue. While identifying frames would approximate content analyses as conducted in social science research on media bias more closely, it would yield lower classification performance [17, 18] or require infeasible effort since frames are typically created for a specific research-question [5]. Our system, however, is meant to analyze media bias caused by framing on any coverage reporting on policy issues. Thus, we seek to determine a fundamental effect resulting from framing: polarity of individual persons, which we identify on sentence- and aggregate to article-level. To achieve state-of-the-art performance in target-dependent sentiment classification (TSC) on news articles, we use a fine-tuned RoBERTa-based neural model ($F1_m = 83.1$) [19].

The last step of frame identification is to determine groups of articles that similarly frame the event, i.e., the persons involved in the event. We currently use a simple, polarity-based method that first determines the person that occurs most frequently across all articles, named most frequent actor (*MFA*). Then, the method assigns each article to one of three groups, depending on whether the article’s MFA mentions are mostly positive, ambivalent, or negative. We also calculate each article’s relevance to the (1) event and the (2) article’s group using simple word-embedding scoring.

4. Visualizations

The overall workflow follows typical online news consumption, i.e., users see first an overview of news events and then view individual news articles. To measure effectiveness not only of our visualizations but also their constituents, we design them so that their components can be altered. To more precisely measure the change in bias-awareness concerning only the textual

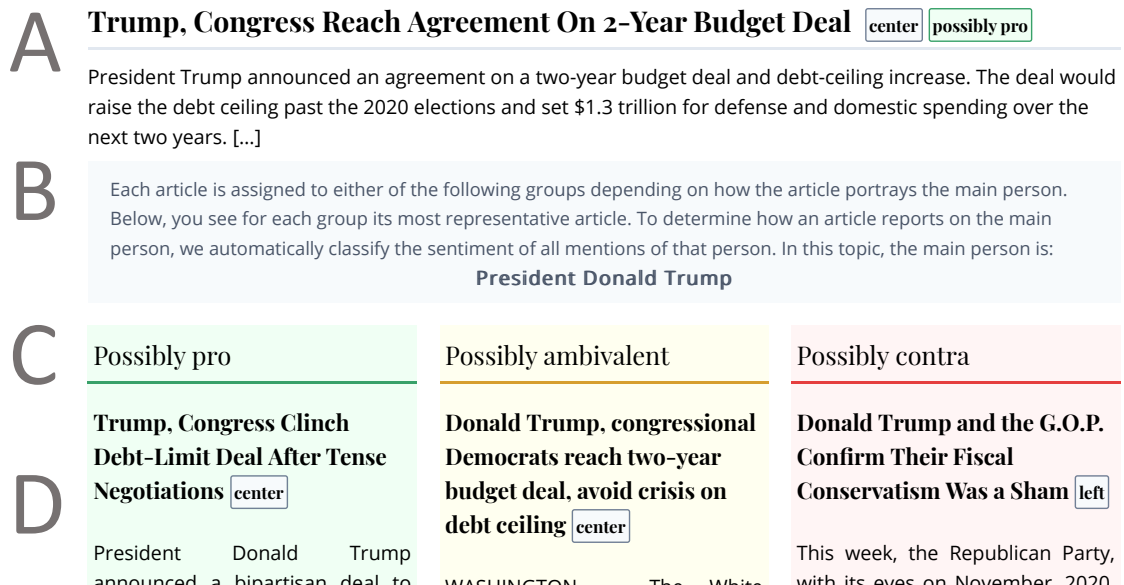


Figure 1: Excerpt of the MFAP overview showing three primary frames present in news coverage on a debt-ceiling event.

content, the visualizations show texts of articles (and information about biases in the texts) but no other content, e.g., photos and outlet name.

4.1. Overview

The overview aims to enable users to quickly get a synopsis of a news event. We devise three visualizations. (1) *Plain* represents popular news aggregators. Using a bias-agnostic design similar to Google News, this baseline shows article headlines and excerpts in a list sorted by their relevance to the event (Section 3). (2) *PolSides*, which represents a bias-aware news aggregator [15], and (3) *MFAP* share a bias-aware, comparative layout but use different methods to determine which frames or biases are present in event coverage.

The layout of *PolSides* and *MFAP* is vertically divided in three parts, two of which are shown in Figure 1. The event’s *main article* (part A) shows the event’s most representative article. The comparative *bias-groups* part (C) shows up to three frames present in event coverage, by showcasing each frame’s most representative article. *PolSides* yields these frames by grouping articles depending on their political orientation (left, center, and right) [15]. For *MFAP*, we use our polarity-based grouping (Section 3) so that the resulting groups represent frames that are primarily in favor, against, or ambivalent regarding the event’s MFA. Conceptually, *PolSides* employs the left-right dichotomy, which is a simple yet often effective means to partition the media into distinctive slants. However, this dichotomy is determined only on the outlet-level and thus may incorrectly classify event-specific framing, e.g., articles with different perspectives having supposedly identical perspectives (and vice versa). Finally, a list shows the headlines of *further articles* reporting on the event (bottom, not shown in Figure 1).

In each overview, further components can be enabled depending on the conjoint profile (cf.




Figure 2: Polarity context bar showing the current and other articles polarity regarding the MFA.

Section 5). *PolSides tags* (shown close to D in Figure 1) and/or *MFAP tags* are shown next to each article headline, and indicate the political orientation of the article’s outlet and the article’s overall polarity regarding the MFA, respectively.

4.2. Article View

The article view shows an article’s text and optionally the following visual clues to communicate bias information: (1) *in-text polarity highlights*, (2) *polarity context bar*, (3) *PolSides tags*, and (4) *MFAP tags* (with identical function to those in the overview). These clues are enabled, disabled, or altered depending on the conjoint profile.

In-text polarity highlights aim to enable users to identify person-targeting sentiment on the sentence-level. We test the effectiveness of the following modes: *single-color* (visually marking a person mention using a neutral color, i.e., gray, if the respective sentence mentions the person positively or negatively), *two-color* (using green and red colors for positive and negative mentions, respectively), *three-color* (same as two-color and additionally showing neutral polarity as gray), and *disabled* (no highlights are shown).

The *polarity context bar* aims to enable users to quickly contrast how the current article and others portray the MFA. The 1D scatter plot depicted in Figure 2 places articles as circles depending on their overall polarity regarding the MFA.

5. Experiments

To evaluate the effectiveness of our system in supporting non-expert users to become aware of biases in news coverage, we conducted a user study consisting of two conjoint experiments. The study seeks to answer two research questions. *What are effective means to communicate biases to non-expert news consumers when viewing an overview of a news topic (RQ1) and when reading a single news article (RQ2)?* The first experiment (E1) focuses on improving the general design of the overview (RQ1), while the second experiment (E2) focuses on answering both RQ1 with improved visualizations and RQ2. All survey data including questionnaires and anonymized respondents’ information is available freely (Section 1).

5.1. Methodology

In both experiments, we used a conjoint design to “separately identify [...] component-specific causal effects by randomly manipulating multiple attributes of alternatives simultaneously”

[20]. Respondents are asked to rate so-called “profiles,” which consist of multiple “attributes,” which are for example the overview, which topic it shows (or which article is shown in the article view), and if or which tags or in-text color highlights are shown. In conjoint design, these attributes are chosen randomly and independently of another for each respondent, which allows an estimation of the relative influence of each component on the bias-awareness (called Average Marginal Component Effects (AMCE)) [20].

We selected three news topics to ensure varying degrees of expected polarization as an indicator for biased coverage: gun control (high polarization), debt ceiling (high-mid), and Australian bushfires (low). We selected a single event for each topic. To ensure heterogeneity in content and writing styles, we manually retrieved a balanced selection of ten articles from left-, center, and right-wing US online outlets as self-identified by them.

We conducted both experiments on Amazon Mechanical Turk. Respondents had to be located in the US, have a history of successfully completed, high quality work, and were compensated 1-2\$ depending on the study duration. In E1, we used data of 260 (of 308) respondents, which satisfied our quality measures, i.e., we discarded 48 respondents that, e.g., were unrealistically fast or answered test questions incorrectly. To keep cognitive load low, respondents were shown only a single topic in the overview, which was randomly drawn from the aforementioned selection. In E2, we used data of 98 (of 110) respondents. To increase cost efficiency, we only showed a selection of overview variants that exhibited positive trends in E1 (instead of fully randomly varying all attributes as in E1). Further, we showed respondents three tasks (resulting in 294 tasks in total), where each task consisted of a single overview and article view.

Our study consists of seven steps. A (1) *pre-study questionnaire* asks demographic data [21]. (2) *Overview* (as described in Section 4.1). A (3) *post-overview questionnaire* operationalizes the bias-awareness in respondents by asking about their perception of the diversity and disagreement in viewpoints, whether the visualization encourages contrasting the individual headlines, and how many perspectives of the public discourse were shown. While it is “intrinsically difficult to objectively define what bias is” [12], on a high level we expect bias to be *perceived* in the form of differences in and opposition of the slant of articles; hence, we operationalize bias-awareness as the motivation and skill of a person to compare and contrast perspectives and information presented in the news using a set of 10-point Likert scaled questions, such as “When shown the overview, did this encourage you to compare and contrast the different articles?” (4) *Article view* (as described in Section 4.2). A (5) *post-article questionnaire* operationalizes bias-awareness in respondents [21]. In a (6) *post-study questionnaire*, users give feedback on the study, i.e., what they (dis)liked. E1 consisted of steps (1–3, 6). E2 consisted of all steps, where (2, 3) may be skipped depending on the conjoint profile and (2–5) were repeated three times since three topics were shown.

5.2. Results and Discussion

We found positive, significant effects on the change in bias-awareness when using the overview variants PolSides or MFAP (RQ1) in E2. Tags also increased bias-awareness significantly. Regarding the article view (RQ2), E2 yielded insignificant results.

Prior to E2, we focused our evaluation of E1 on qualitatively identifying flaws in the design of visualizations and the study, due to the lack of significant trends in E1. For example, 35% users

Table 1

Effects in E2 (excerpt). Names in parentheses indicate enabled tags (except for “random”). Baselines are italic.

Attribute	Level	Est.	SE	z	Pr(> z)
Overview	PolS. (PolS.)	7.83	2.06	3.79	<.001***
	MFAP (both)	5.87	1.73	3.39	<.001***
	MFAP (none)	6.13	2.00	3.05	.002**
	MFAP (MFAP)	5.67	2.10	2.69	.007**
	MFAP (PolS.)	5.60	1.85	3.02	.002**
	MFAP (random)	5.76	2.40	2.39	.016*
	Plain (both)	3.66	1.90	1.92	.053
	Plain (MFAP)	1.28	2.24	0.57	.568
	Plain (PolS.)	3.04	2.07	1.46	.142
	<i>Plain (none)</i>	0	-	-	-
Topic	Debt Ceiling	-1.52	0.78	-1.93	.052
	Gun control	-1.14	0.84	-1.34	.179
	<i>Bushfire</i>	0	-	-	-

experienced a lack of clarity and transparency, e.g., how the visualized information was derived. This weakness was exaggerated if respondents felt the shown information was incorrect (6% PolSides, 11% MFAP), e.g., an article that seemed negative from its headlines was labeled as ambivalent. Prior to E2, we addressed all the major lines of criticism, e.g., by adding brief explanations about all visual clues and in particular about the bias grouping (see B in Figure 1).

The goal of E2 was to test the set of overviews that had positive trends in E1 (to answer RQ1) and to test components in the article view (RQ2). E2 showed positive, significant effects of both bias-sensitive overviews, where the best overviews were PolSides (with PolSides tags) and MFAP (without tags). The AMCEs in Table 1 show that both overviews have very high and strongly significant effectiveness (PolSides $Est = 7.83$ and MFAP $Est = 6.13$, which are not significantly different to another albeit one being slightly higher [22]). Quantitative (by analyzing the effects on the individual questions composing the overall post-overview score) and qualitative analysis of both overviews suggests that MFAP reveals biases that are *actually present* in the news articles, whereas PolSides *only facilitates* the visibility of biases—a typical issue of prior approaches for automated bias detection [14]. By imitating the content analysis, our system yielded substantial frames as shown in Figure 1 whereas PolSides showed, e.g., the following headlines of rather “artificial” frames: “Trump Announces Deal On Debt Limit, Spending Caps” and “Trump, Congress Clinch Debt-Limit Deal After Tense Negotiations.” MFAP (random), an overview where articles were randomly assigned to one bias-group, yielded $Est = 5.76$, indicating that only pointing out possible biases already increased bias-awareness.

In MFAP, showing no tags yielded the highest effectiveness (6.13 compared to 5.87 when both tags were shown). This indicates that the MFAP design reveals “enough” bias information and further visual clues may yield too complex visualizations. None of the tags alone have significant effects when combined with the Plain version, indicating that a bias-group layout is necessary for bias-awareness.

In the article view, only showing the PolSides tags had a significant positive effect on bias awareness (2.45). There were no significant effects for the MFAP tags and polarity context bar. Analyzing respondents' criticism in the post-study questions, we attribute this to two shortcomings. First, too few in-text highlights to have a consistent effect (21% of the article views had ≤ 5 highlights, 7% had none). When controlling for the number of highlights, they had a significant, positive effect on bias-awareness. Second, there was a strong influence of individual topics on the effectiveness, e.g., respondents reported the debt ceiling topic was "too complicated" or "boring" to follow. We plan to address this by conducting a study with more respondents and a wider range of topics, to ensure a better representation of the public discourses. Doing so will also strengthen the generalizability of the results and allow to investigate the effects of users' demographic data on their bias-awareness and change thereof [23]. Due to the small sample size in E2, our current analysis was inconclusive regarding demographic effects.

Although we did not filter for a representative sample of the US population, the distributions of our samples are approximately similar to the distributions of the US population in key dimensions such as age and political education.¹ However, we propose to verify the generalizability of the study's findings to the entire US population or other countries using a larger respondent sample. Further, in E1 and E2 we assumed that MTurk workers are mostly non-expert news consumers. To verify this, we propose to explicitly ask for participants' degree of media literacy.

6. Conclusion

We present the first system to automatically identify and then communicate person-targeting forms of bias in news articles reporting on policy events. Earlier, these biases could only be identified using content analyses, which—despite their effectiveness in capturing also subtle yet powerful biases—could only be conducted for few topics in the past due to their high cost, manual effort, and required expertise. In a large-scale user study, we employ a conjoint design to measure the effectiveness of visualizations and individual components. We find that our overviews significantly increase bias-awareness in respondents. In particular and in contrast to prior work, our bias-identification method seems to reveal biases that emerge from the content of news coverage and individual articles. In practical terms, our results suggest that the biases found and communicated by our method are actually *present* in the news articles, whereas the reviewed prior work only *facilitates* detection of biases, e.g., by distinguishing between left- and right-wing outlets. In sum, our exploratory work indicates the effectiveness of bias-sensitive news recommendation as a promising line of research for future work.

Acknowledgments

This work is funded by the WIN program of the Heidelberg Academy of Sciences and Humanities, financed by the Ministry of Science, Research and the Arts of the State of Baden-Württemberg, Germany. The authors thank the anonymous reviewers for their valuable comments that helped to improve this paper.

¹For statistics of the sample please refer to the online resources, see Section 1.

References

- [1] S. DellaVigna, E. Kaplan, The Fox News Effect: Media Bias and Voting, Technical Report 3, National Bureau of Economic Research, Cambridge, MA, 2006. URL: <http://www.nber.org/papers/w12169.pdf>. doi:10.3386/w12169.
- [2] C. Budak, What happened? The Spread of Fake News Publisher Content During the 2016 U.S. Presidential Election, in: The World Wide Web Conference on - WWW '19, ACM Press, New York, New York, USA, 2019, pp. 139–150. URL: <http://dl.acm.org/citation.cfm?doid=3308558.3313721>. doi:10.1145/3308558.3313721.
- [3] F. Hamborg, K. Donnay, B. Gipp, Automated identification of media bias in news articles: an interdisciplinary literature review, International Journal on Digital Libraries 20 (2019) 391–415. URL: <https://doi.org/10.1007/s00799-018-0261-y>. doi:10.1007/s00799-018-0261-y.
- [4] F. Hamborg, Media Bias, the Social Sciences, and NLP: Automating Frame Analyses to Identify Bias by Word Choice and Labeling, in: Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: Student Research Workshop, Association for Computational Linguistics, Stroudsburg, PA, USA, 2020, pp. 79–87. URL: <https://www.aclweb.org/anthology/2020.acl-srw.12>. doi:10.18653/v1/2020.acl-srw.12.
- [5] R. M. Entman, Framing Bias: Media in the Distribution of Power, Journal of Communication 57 (2007) 163–173. URL: <https://academic.oup.com/joc/article/57/1/163-173/4102665>. doi:10.1111/j.1460-2466.2006.00336.x.
- [6] M. S. Davis, E. Goffman, Frame Analysis: An Essay on the Organization of Experience., Contemporary Sociology 4 (1975) 599. URL: <http://www.jstor.org/stable/2064021?origin=crossref>. doi:10.2307/2064021.
- [7] S. Park, M. Ko, J. Kim, Y. Liu, J. Song, The politics of comments, in: Proceedings of the ACM 2011 conference on Computer supported cooperative work - CSCW '11, ACM, ACM Press, New York, New York, USA, 2011, p. 113. URL: <http://portal.acm.org/citation.cfm?doid=1958824.1958842>. doi:10.1145/1958824.1958842.
- [8] S. A. Munson, D. X. Zhou, P. Resnick, Sidelines: An Algorithm for Increasing Diversity in News and Opinion Aggregators., in: ICWSM, 2009.
- [9] M. Recasens, C. Danescu-Niculescu-Mizil, D. Jurafsky, Linguistic Models for Analyzing and Detecting Biased Language, in: Proceedings of the 51st Annual Meeting on Association for Computational Linguistics, Association for Computational Linguistics, Sofia, BG, 2013, pp. 1650–1659. URL: <https://www.aclweb.org/anthology/P13-1162.pdf>.
- [10] S. Mullainathan, A. Shleifer, The market for news, American Economic Review (2005) 1031–1053.
- [11] H.-K. Kong, Z. Liu, K. Karahalios, Frames and Slants in Titles of Visualizations on Controversial Topics, in: Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems, ACM, New York, NY, USA, 2018, pp. 1–12. URL: <https://dl.acm.org/doi/10.1145/3173574.3174012>. doi:10.1145/3173574.3174012.
- [12] S. Park, S. Kang, S. Chung, J. Song, NewsCube: Delivering multiple aspects of news to mitigate media bias, in: Proceedings of the 27th international conference on Human factors in computing systems - CHI 09, ACM Press, New York, New York, USA, 2009, p. 443. URL: <http://dl.acm.org/citation.cfm?doid=1518701.1518772>. doi:10.1145/1518701.1518772.

- [13] S. Park, M. Ko, J. Kim, H. Choi, J. Song, NewsCube 2.0: An Exploratory Design of a Social News Website for Media Bias Mitigation, in: Workshop on Social Recommender Systems, 2011.
- [14] F. Hamborg, N. Meuschke, B. Gipp, Bias-aware news analysis using matrix-based news aggregation, *International Journal on Digital Libraries* 21 (2020) 129–147. URL: <http://link.springer.com/10.1007/s00799-018-0239-9>. doi:10.1007/s00799-018-0239-9.
- [15] AllSides.com, AllSides - balanced news, 2021.
- [16] F. Hamborg, N. Meuschke, C. Breitingner, B. Gipp, news-please: A Generic News Crawler and Extractor, in: Proceedings of the 15th International Symposium of Information Science, Verlag Werner Hülsbusch, 2017, pp. 218–223.
- [17] D. Card, A. E. Boydstun, J. H. Gross, P. Resnik, N. A. Smith, The Media Frames Corpus: Annotations of Frames Across Issues, in: Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers), Association for Computational Linguistics, Stroudsburg, PA, USA, 2015, pp. 438–444. URL: <http://aclweb.org/anthology/P15-2072>. doi:10.3115/v1/P15-2072.
- [18] D. Card, J. Gross, A. Boydstun, N. A. Smith, Analyzing Framing through the Casts of Characters in the News, in: Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, Association for Computational Linguistics, Stroudsburg, PA, USA, 2016, pp. 1410–1420. URL: <http://aclweb.org/anthology/D16-1148>. doi:10.18653/v1/D16-1148.
- [19] F. Hamborg, K. Donnay, Newsmtsc: (multi-)target-dependent sentiment classification in news articles, in: Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics (EACL 2021), 2021, pp. 1663–1675.
- [20] J. Hainmueller, D. J. Hopkins, T. Yamamoto, Causal Inference in Conjoint Analysis: Understanding Multidimensional Choices via Stated Preference Experiments, *Political Analysis* 22 (2014) 1–30. URL: https://www.cambridge.org/core/product/identifier/S1047198700013589/type/journal_article. doi:10.1093/pan/мпt024.
- [21] T. Spinde, F. Hamborg, K. Donnay, A. Becerra, B. Gipp, Enabling News Consumers to View and Understand Biased News Coverage: A Study on the Perception and Visualization of Media Bias, in: Proceedings of the ACM/IEEE Joint Conference on Digital Libraries in 2020, ACM, New York, NY, USA, 2020, pp. 389–392. URL: <https://dl.acm.org/doi/10.1145/3383583.3398619>. doi:10.1145/3383583.3398619.
- [22] A. Gelman, H. Stern, The Difference Between “Significant” and “Not Significant” is not Itself Statistically Significant, *The American Statistician* 60 (2006) 328–331. URL: <http://www.tandfonline.com/doi/abs/10.1198/000313006X152649>. doi:10.1198/000313006X152649.
- [23] K. Coe, D. Tewksbury, B. J. Bond, K. L. Drogos, R. W. Porter, A. Yahn, Y. Zhang, Hostile News: Partisan Use and Perceptions of Cable News Programming, *Journal of Communication* 58 (2008) 201–219. URL: <https://academic.oup.com/joc/article/58/2/201-219/4098517>. doi:10.1111/j.1460-2466.2008.00381.x.