



 Latest updates: <https://dl.acm.org/doi/10.1145/3706598.3713716>

RESEARCH-ARTICLE

## Media Bias Detector: Designing and Implementing a Tool for Real-Time Selection and Framing Bias Analysis in News Coverage

JENNY S WANG, Harvard Business School, Boston, MA, United States

SAMAR HAIDER, University of Pennsylvania, Philadelphia, PA, United States

AMIR TOHIDI, University of Pennsylvania, Philadelphia, PA, United States

ANUSHKAA GUPTA, University of Pennsylvania, Philadelphia, PA, United States

YUXUAN ZHANG, University of Pennsylvania, Philadelphia, PA, United States

CHRIS CALLISON-BURCH, University of Pennsylvania, Philadelphia, PA, United States

[View all](#)

Open Access Support provided by:

[University of Pennsylvania](#)

[Harvard Business School](#)

[Microsoft Research](#)



PDF Download  
3706598.3713716.pdf  
18 December 2025  
Total Citations: 1  
Total Downloads: 3300

Published: 26 April 2025

[Citation in BibTeX format](#)

CHI 2025: CHI Conference on Human Factors in Computing Systems  
April 26 - May 1, 2025  
Yokohama, Japan

Conference Sponsors:  
SIGCHI

# Media Bias Detector: Designing and Implementing a Tool for Real-Time Selection and Framing Bias Analysis in News Coverage

Jenny S Wang\*

Harvard Business School  
Boston, Massachusetts, USA  
jewang@hbs.edu

Anushkaa Gupta

University of Pennsylvania  
Philadelphia, Pennsylvania, USA  
agupta08@sas.upenn.edu

David Rothschild

Microsoft Research  
New York, New York, USA  
david@researchdmr.com

Samar Haider\*

University of Pennsylvania  
Philadelphia, Pennsylvania, USA  
samarh@seas.upenn.edu

Yuxuan Zhang

University of Pennsylvania  
Philadelphia, Pennsylvania, USA  
yuxuanzh@seas.upenn.edu

Amir Tohidi

University of Pennsylvania  
Philadelphia, Pennsylvania, USA  
atohidi@seas.upenn.edu

Chris Callison-Burch

University of Pennsylvania  
Philadelphia, Pennsylvania, USA  
ccb@upenn.edu

## Abstract

Mainstream media, through their decisions on what to cover and how to frame the stories they cover, can mislead readers without using outright falsehoods. Therefore, it is crucial to have tools that expose these editorial choices underlying media bias. In this paper, we introduce the MEDIA BIAS DETECTOR, a tool for researchers, journalists, and news consumers. By integrating large language models, we provide near real-time granular insights into the topics, tone, political lean, and facts of news articles aggregated to the publisher level. We assessed the tool's impact by interviewing 13 experts from journalism, communications, and political science, revealing key insights into usability and functionality, practical applications, and AI's role in powering media bias tools. We explored this in more depth with a follow-up survey of 150 news consumers. This work highlights opportunities for AI-driven tools that empower users to critically engage with media content, particularly in politically charged environments.

## CCS Concepts

- Human-centered computing → User studies; Empirical studies in HCI;
- Information systems → Web applications.

## Keywords

media bias, news analysis, large language models (LLMs), LLM-driven tools

## ACM Reference Format:

Jenny S Wang, Samar Haider, Amir Tohidi, Anushkaa Gupta, Yuxuan Zhang, Chris Callison-Burch, David Rothschild, and Duncan J Watts. 2025. Media Bias Detector: Designing and Implementing a Tool for Real-Time Selection

\*Both authors contributed equally to this research.



This work is licensed under a Creative Commons Attribution 4.0 International License.  
CHI '25, Yokohama, Japan

© 2025 Copyright held by the owner/author(s).  
ACM ISBN 979-8-4007-1394-1/25/04  
<https://doi.org/10.1145/3706598.3713716>

and Framing Bias Analysis in News Coverage. In *CHI Conference on Human Factors in Computing Systems (CHI '25)*, April 26–May 01, 2025, Yokohama, Japan. ACM, New York, NY, USA, 27 pages. <https://doi.org/10.1145/3706598.3713716>

## 1 Introduction

Since the 2016 election, there has been growing concern about the pervasive impact of fake news [5, 52, 79, 97]. However, in practice, most people consume news from mainstream sources, where stories, though factually accurate, can still be *biased* [6]. This bias arises from the considerable control that journalists and editors exert over *selection* (choosing to emphasize some issues, events, or people over others) as well as their *framing* (choosing the tone, perspective, or facts of a story to present an issue in a particular way) [28, 33, 88]. Unlike outright falsehoods, selection and framing can subtly mislead readers, making these forms of bias difficult for traditional fact-checking tools to detect [77]. Experimental evidence has shown that it is remarkably easy to mislead a reader without making any explicitly false statements [4, 54]. Moreover, because factually accurate but biased news coverage tends to be consumed by much larger populations than categorically fake news, its negative impact can be considerably larger in aggregate [7].

Measuring and exposing media bias, which we define as the preferential selection of some stories, facts, people, events, or perspectives over others, is critical to challenging the illusion of objectivity promoted by many news organizations. By revealing how editorial choices (decisions about which facts, voices, and perspectives to emphasize or downplay) shape coverage, readers can better understand both the context of the information they consume and the perspectives they are missing. A lack of diverse perspectives can contribute to echo chambers, where individuals relying on a single source or a limited set of ideologically similar sources are more likely to encounter “separate realities,” fostering disagreement on basic facts or priorities [15, 46, 69].

Communications researchers have long emphasized the importance of media bias [11, 27, 28, 94, 101]; however, methodological limitations have until recently prevented its identification at scale.

Traditional approaches, like manually quantifying topics, viewpoints, or facts across publishers, are costly and time-consuming [21, 55, 66], while natural language processing (NLP) methods such as n-gram counting typically fail to capture the full context of an article [29, 38]. Now, with advancements in large language models (LLMs), it has become possible to annotate large documents faster and more efficiently, while maintaining high accuracy [41]. Specifically, LLMs can efficiently ingest and parse tens of thousands of articles, extracting key features such as topics, subtopics, and facts as well as labeling them according to subjective quantities such as tone (positive vs. negative) and partisanship (Democrat vs. Republican) [78, 102]. This ability to extract nuanced information from large datasets makes LLMs a valuable tool for identifying and analyzing media bias.

In this paper, we present the **MEDIA BIAS DETECTOR**, an LLM-driven tool designed to dynamically track and analyze major news coverage from a diverse collection of prominent publishers across the political spectrum.<sup>1</sup> We operationalize media bias in two distinct ways, corresponding to the aforementioned distinction between selection and framing bias. First, we quantify selection bias by measuring the differential attention paid to different news categories (e.g. politics vs. business) as well as topics (e.g. the 2024 election within politics) and subtopics (e.g. the presidential horse race within the 2024 election). Second, we quantify framing bias along two key dimensions: a) political lean, which we define as the extent to which an article aligns with the viewpoints, policies, or concerns of Republicans versus Democrats, either explicitly or implicitly; and b) tone, defined as the emotional slant of coverage (i.e. positive, neutral, negative). Unlike existing tools, which typically consider subsets of news stories and label publishers broadly as left, right, or center, the **MEDIA BIAS DETECTOR** annotates individual articles produced every day. These annotations are then aggregated to provide a dynamic, data-driven view of each publisher, reflecting the diversity of their coverage rather than assigning static ideological labels. Our inclusion of tone also addresses a gap in media bias detection, which often focuses on political lean but has paid less attention to sentiment as a source of bias [87]. This added granularity allows readers to explore the distribution of lean and tone across articles published by a news outlet. For instance, a user might find that while the Wall Street Journal appears generally neutral overall, its coverage leans more Republican on immigration and slightly more Democratic on the environment and reproductive rights.

The **MEDIA BIAS DETECTOR** allows users—whether researchers, journalists, or everyday news consumers—to interactively explore and quantify media bias by examining the coverage focus, article type (i.e., report, analysis, or opinion), tone, and political lean of the top news stories. With this tool, users can examine how different publishers emphasize or ignore different topics, such as whether they focus more on Joe Biden's or Donald Trump's age, or whether they frame the election as an entertaining horse race versus a discussion of key policy differences. It also allows users to explore topic-specific questions, such as quantifying how much inflation is discussed compared to wages, or how crime is framed: does it focus on the perspective of law enforcement, economic impact, or the

<sup>1</sup>We limit our initial investigation to ten mainstream news publishers to maintain data quality and keep labeling costs manageable. But, the tool is easily adaptable to additional publishers.

effects of incarceration? Furthermore, the **MEDIA BIAS DETECTOR** helps users keep up with the fast pace of the news cycle by summarizing the top events of the day and highlighting key snippets of information from these stories. To evaluate the effectiveness of the **MEDIA BIAS DETECTOR** and inform future improvements, we conducted semi-structured interviews with 13 experts in media, communications, and politics, as well as a survey of 150 everyday news consumers. The goal of collecting feedback from these user groups was to address the following research questions:

- RQ1.** Can we develop a tool that effectively conveys media bias to users while remaining easy to use?
- RQ2.** How does the use of LLMs in the **MEDIA BIAS DETECTOR** influence users' trust in the tool compared to human-generated ratings?
- RQ3.** Which audiences are most likely to benefit from the **MEDIA BIAS DETECTOR**, and how can the tool be optimized to better serve these groups?

In our evaluation of the **MEDIA BIAS DETECTOR**, several key themes emerged. Experts appreciated the **MEDIA BIAS DETECTOR**'s ability to consolidate complex information and provide multiple forms of evaluation, emphasizing its value for media literacy education and qualitative research. They also stressed the importance of transparency when using LLMs to classify information related to complex topics like media bias. Everyday users found the tool useful and accessible, often discovering new insights they had not known before, and generally placed more trust in AI tools than experts. Both groups, however, valued the reassurance of human oversight. Overall, the **MEDIA BIAS DETECTOR** proved to be a versatile tool that appealed to a wide range of audiences for education, research, and daily bias checks.

To summarize, the main contributions of our work are as follows:

- (1) The implementation of the **MEDIA BIAS DETECTOR**, an interactive online tool that aggregates key information from news articles and guides users to explore different facets of media bias.
- (2) The results of a user study with 13 media experts and a survey of 150 news consumers, which demonstrate the **MEDIA BIAS DETECTOR**'s broad applicability and value in helping users understand and discover media bias in online news.

## 2 Background and Related Work

Bias, in statistical theory, implies the presence of a "true" value from which the biased estimate systematically differs. However, in practice, we generally lack a ground truth, making it infeasible to directly assess bias for individual articles or even entire publishers. This assertion might seem counterintuitive, as we often perceive articles as biased when we read them. This perception arises because we implicitly treat our subjective opinions as the "truth," even though they are likely biased as well [81]. In the absence of a universally accepted ground truth, we instead need to identify bias by observing patterns and differences in how news is reported between publishers, or within publishers over time. In this section, we anchor our discussion of media bias in the concepts of selection and framing biases and consider their relevance to designing media bias tools grounded in HCI principles. We also situate the **MEDIA BIAS DETECTOR** among existing media bias detection tools

and discuss their limitations. Finally, we review current literature on LLMs for annotation tasks and explain why these models are well-suited for analyzing media bias.

## 2.1 Agenda-Setting, Framing, and the Role of HCI in Designing Tools for Media Bias

Media bias presents major implications for the HCI community as online platforms increasingly become the primary medium for news consumption. Within the communications field, media bias is often analyzed through two key theories: agenda-setting and framing. Agenda-setting theory argues that the media doesn't tell us what to think but rather what to think about [60, 61]. The process of *selection bias*—choosing certain stories or facts to report while omitting others—directly contributes to the media's agenda-setting power, shaping public discourse by directing attention to specific topics over others [30]. For example, a news outlet might publish many stories on a frontrunner's campaign rallies while giving limited coverage to other candidates. Agenda-setting theory is highly relevant to HCI because digital platforms increasingly mediate what content users see [71]. Algorithms, interface design, and personalized recommendation systems often act as agenda-setters, determining what content is presented or excluded from users' screens [34, 59]. This phenomenon raises critical questions in HCI about how systems shape user awareness and understanding, as well as how tools can be designed to challenge these embedded biases.

Beyond selecting which issues to emphasize, the media also influences public perception via the manner in which information is presented—a process known as framing [100]. *Framing bias* builds on selection bias by not just determining what is covered, but how it is covered. This involves the inclusion or omission of specific details and perspectives, the tone or language used, the context provided, or the omission of background information, all of which can significantly alter how an issue is perceived by the audience [28, 33, 37]. For instance, framing a protest as a "riot" in one article versus a "peaceful demonstration" in another can significantly influence how readers perceive the event [18, 95]. Similarly, framing immigration as an "opportunity for economic growth" versus an "immigration crisis" can sway public perception on the issue [48]. Framing effects relate to HCI through a shared focus on how the presentation and emphasis of information influences its importance or impact [45]. Media bias tools informed by HCI principles can play a critical role in revealing framing biases [31]. For example, tools that allow users to analyze data from multiple perspectives [3], explore specific points in more detail [91], or compare diverse sources [12, 16] can foster a deeper understanding of how the same event is portrayed in different ways.

Although the concepts of selection and framing bias are well-known in media studies and communications [29, 58, 75], their impact has become even more pronounced in today's digital age [14, 43]. The rapid dissemination and broad reach of online news allow information consumers to quickly access content that reinforces their pre-existing opinions. Research indicates that online news consumption exhibits a polarized pattern, with users spending significantly more time on news sources that align with their political leanings compared to those that do not [36]. Tools that

make media bias visible—especially in terms of selection and framing biases—can help users recognize these patterns and engage more critically with their news consumption [51]. By applying HCI principles to these tools, we can design interfaces that promote media literacy and encourage balanced engagement with digital news platforms.

## 2.2 Examining Current Approaches and Limitations in Assessing Media Bias

Scholars have proposed numerous taxonomies to understand media bias, yet no universally accepted set of media bias metrics or standard measurements exist [39, 42, 47, 67, 82]. Current methods often reduce news publishers to a single metric by assigning a political lean rating from "Left" to "Right" [8, 26, 72]. These tools simplify bias labels to make it easier for users to digest, such as using 2-dimensional axes with political lean and reliability [63] or a 1-dimensional categorization of publishers across the political spectrum [8]. Although useful, these static classifications fail to capture within-publisher differences in bias over time, across topics, or even across individual articles.

Evaluating bias at the article level is less common and more challenging. Earlier tools evaluating article level bias relied on basic NLP techniques, such as keyword frequency analysis, to cluster similar stories in the media [76]. However, these methods are limited in their ability to account for context and the evolving nature of news stories. Recent advancements in LLMs offer improved contextual understanding capabilities which can more accurately help identify media bias in news coverage. Even though other proprietary tools that incorporate AI features have emerged [8, 13], these platforms generally do not update to reflect media bias in recent news coverage. Furthermore, the lack of transparency around their AI systems makes it unclear which models they use and how they arrive at their conclusions.

The MEDIA BIAS DETECTOR addresses these limitations by moving beyond oversimplified metrics and publisher level analyses. Rather than simplifying media bias to one or two dimensions, the MEDIA BIAS DETECTOR captures framing and selection biases through a multidimensional analysis of publisher coverage over time. This approach allows users to compare tone, political lean, and content focus between publishers, and across topics, resulting in a more comprehensive and dynamic approach to media bias. The MEDIA BIAS DETECTOR is intentionally designed to provide near real-time updates, allowing users to explore how media coverage—and the biases within it—shift dynamically as new information emerges. Furthermore, we make our underlying methodology, including our model usage, prompt phrasing, and human in the loop verification framework, readily available to all users of our tool (see Appendix A.4).

## 2.3 Incorporating Large Language Models (LLMs) for Media Bias Detection

LLMs, such as GPT-4, utilize deep contextual embeddings to capture subtle semantic relationships and nuances within text [1, 19, 83]. Studies have shown that LLMs perform well in generative tasks such as summarization [56, 57, 85], as well as discriminative tasks like sentiment analysis and topic classification [25, 78, 107]. More recent

studies have shown that state-of-the-art LLMs exhibit increasingly complex reasoning and problem-solving capabilities [20]. These findings suggest that LLMs can be especially useful for the complex task of media bias detection, where understanding context, tone, and subtle language is crucial.

While traditional machine learning models used in annotation tasks require task-specific training, instruction-tuned pre-trained LLMs can generalize to different tasks in a zero-shot setting [98]. Prior work has shown that using LLMs to automate the initial stages of data annotation enables researchers to quickly process large volumes of content that would be unmanageable with manual annotation alone [109]. This scalability is critical given the speed with which major news outlets publish new articles, which must be analyzed in near real-time to capture evolving narratives and help readers interpret the news. By incorporating LLMs into the MEDIA BIAS DETECTOR’s news analysis pipeline, we can efficiently extract detailed information like sentence-level codings and classifications by topic, subtopic, article type, tone, and political lean, while significantly increasing the volume of articles we can annotate. However, it is important to acknowledge potential limitations of LLMs in zero-shot settings, such as political biases in their outputs and the influence of this bias on the tool’s results. These challenges are explored in greater depth in Section 9.

The integration of LLMs into traditionally human processes is informed by a growing body of research that shows that human-AI collaboration enhances the accuracy and reliability of automated systems [41, 96]. Although LLMs can quickly process large datasets and perform initial classifications, human oversight is crucial for providing context that LLMs do not have for current events and verifying the model’s outputs for consistent standards of accuracy [10, 98]. In the MEDIA BIAS DETECTOR, humans play an important role at every step, from generating a targeted list of relevant news topics for LLMs to classify, to continuously monitoring the model’s classifications over time.

We intentionally chose to evaluate the responses generated by LLMs through a process of human annotation that emphasizes validation over independent labeling. Rather than having annotators blindly label the data, we engage them in a validation task where they read GPT-generated responses to assess their coherence and soundness. Although having humans independently label content and then comparing it to GPT’s output could yield interesting insights, it is important to recognize that human disagreements often occur, even among reasonable individuals. Particularly in complex tasks, such as reading an entire article and labeling its lean and tone, as well as extracting facts, achieving high agreement is challenging [65]. Our methodology accounts for this inherent subjectivity and aims to ensure that topics presented on the MEDIA BIAS DETECTOR are coherent, reasonable, and free from overt inaccuracies. By involving humans in the validation process, we impose a layer of quality control that acknowledges the absence of a singular ground truth.

### 3 Design Considerations

The overarching design goal of our tool is not to determine the media bias of each publisher for the user, but to provide affordances

that allow users to break down information and compare mainstream media publications. These affordances enable users to see the aggregated raw data broken down by topic and subtopic, and to compare the overall lean and tone within each topic and subtopic to understand the nature of each publisher’s views. This section outlines our key design considerations, summarized in Table 1.

#### D1: Enable *broad* exploration of the same data.

Media bias is a multidimensional problem [32] and cannot accurately be captured by just one metric. A central goal of our tool is to support *breadth* of exploration, providing users with affordances that enable them to examine the data from multiple perspectives and draw their own conclusions. This design is informed by principles of learning with multiple representations, which enhance cognitive processes in learning by allowing users to approach complex data from different angles [2, 3]. To operationalize these principles, we use faceted categories to structure the interface, enabling users to dynamically toggle between variables and uncover relationships that might otherwise be difficult to identify [53, 103]. For example, users can switch between metrics such as coverage volume, political lean, and tone across publishers and topics. They can also adjust the date range to compare publishers over any time frame and filter by article type, allowing for a comprehensive comparison of how various content categories reflect media bias over selected time periods. Figures 1 and 2 showcase how the MEDIA BIAS DETECTOR incorporates faceted categories to help users process information through multiple representations.

#### D2: Enable *deep* exploration of specific data.

In addition to being a multidimensional problem, media bias is also a *multilevel* one. While overall news coverage on a particular category may not differ much between two publishers, how they choose to cover specific topics, subtopics, and events within that category can reveal a hidden bias. A major affordance of our tool is the ability to engage in *deep* exploration of the data after seeing an overview, allowing users to uncover more detailed biases without being overwhelmed at the very first stage. This approach, informed by the principle of *progressive disclosure*, affords a gradual reveal of complexity as users build familiarity with the system [23, 70, 90, 91]. From a media bias standpoint, we apply progressive disclosure by allowing users to begin with broad categories and gradually narrow their focus to more specific areas. As users engage with a particular category, they can drill down into increasingly detailed topics and subtopics, enabling a deeper exploration of media coverage. This also extends beyond topic-level analysis, as our tool affords sentence-level text annotation which helps users examine how different news publishers write about the same events. Users can begin by exploring a news event cluster, dive into the top facts reported about that event, and then compare how the same facts are presented across different publications. Figure 3 illustrates how a user can progressively refine their analysis at the category, topic, and subtopic levels.

**Table 1: Design considerations for the MEDIA BIAS DETECTOR, summarizing the goals, descriptions, and key strategies employed to address media bias.**

Design Consideration	Description	Key Features and Strategies
<b>D1:</b> Broad exploration of the same data	Media bias is a multidimensional problem that requires users to explore data from various perspectives.	<ul style="list-style-type: none"> <li>- Facilitates multiple representations of data to enhance cognitive understanding.</li> <li>- Implements faceted categories to allow toggling between metrics such as coverage volume, political lean, and tone.</li> <li>- Uses progressive disclosure to reduce cognitive load and gradually reveal complexity.</li> <li>- Allows drilling down into detailed topics, subtopics, and events within broader categories.</li> <li>- Provides sentence-level text annotation for event-level comparisons.</li> </ul>
<b>D2:</b> Deep exploration of specific data	Media bias is a multilevel problem, requiring detailed examination beyond overviews to reveal nuanced biases.	<ul style="list-style-type: none"> <li>- Provides sentence-level text annotation for event-level comparisons.</li> <li>- Designed to expose key elements of news stories across the political spectrum using side-by-side comparisons.</li> <li>- Integrates breadth (D1) and depth (D2) principles for cohesive comparative analysis.</li> </ul>
<b>D3:</b> Enable easy comparison across publishers and topics	Media bias can be more clearly understood by comparing similarities and differences between publishers and topics.	<ul style="list-style-type: none"> <li>- News, The Guardian, The Huffington Post, The New York Times, USA Today, The Wall Street Journal, and The Washington Post.</li> </ul>

### D3: Enable easy comparison across publishers and topics.

News readers often have limited time and tend to rely on a single source or ideologically aligned sources to learn about current events [49]. The affordance of comparative analysis within our tool allows users to easily identify similarities and differences between multiple publishers in one place, which may help them recognize discrepancies between perceived and actual coverage [12, 16, 92]. To reduce cognitive load, our tool provides side-by-side comparisons of publishers and topics with the goal of better contextualizing the broader news ecosystem without users needing to switch between sources.

While D1 emphasizes breadth by encouraging users to explore data across multiple dimensions and D2 enables deep, focused analysis, D3 integrates these principles into a cohesive framework for comparison. Table 1 provides a summary of these design considerations and key features addressed by each of them. We used these considerations to design and implement the MEDIA BIAS DETECTOR, an interactive dashboard that exposes and aggregates key elements of top news stories across the political spectrum in close to real time. Throughout Section 4, we illustrate how the MEDIA BIAS DETECTOR facilitates comparative analysis through its various dashboards and visualization formats, helping users to critically evaluate similarities and differences across publishers and topics.

## 4 MEDIA BIAS DETECTOR: An Interactive Near-Real-Time News Tracking Tool

This section introduces the design of the MEDIA BIAS DETECTOR, a tool built to help users navigate the complexities of media bias through three guiding principles outlined in Section 3: broad exploration (**D1**), deep and multilevel exploration (**D2**), and comparative analysis (**D3**).<sup>2</sup> The dashboard currently focuses on ten prominent online news publishers, selected for their mix of reach and agenda-setting influence: Associated Press News, Breitbart News, CNN, Fox

News, The Guardian, The Huffington Post, The New York Times, USA Today, The Wall Street Journal, and The Washington Post. Although this focus ensures a balanced representation of diverse viewpoints, it is important to note the inherent limitation of restricting the news articles to only ten sources, which may exclude other influential or regional publishers. We discuss this limitation further in Section 9.

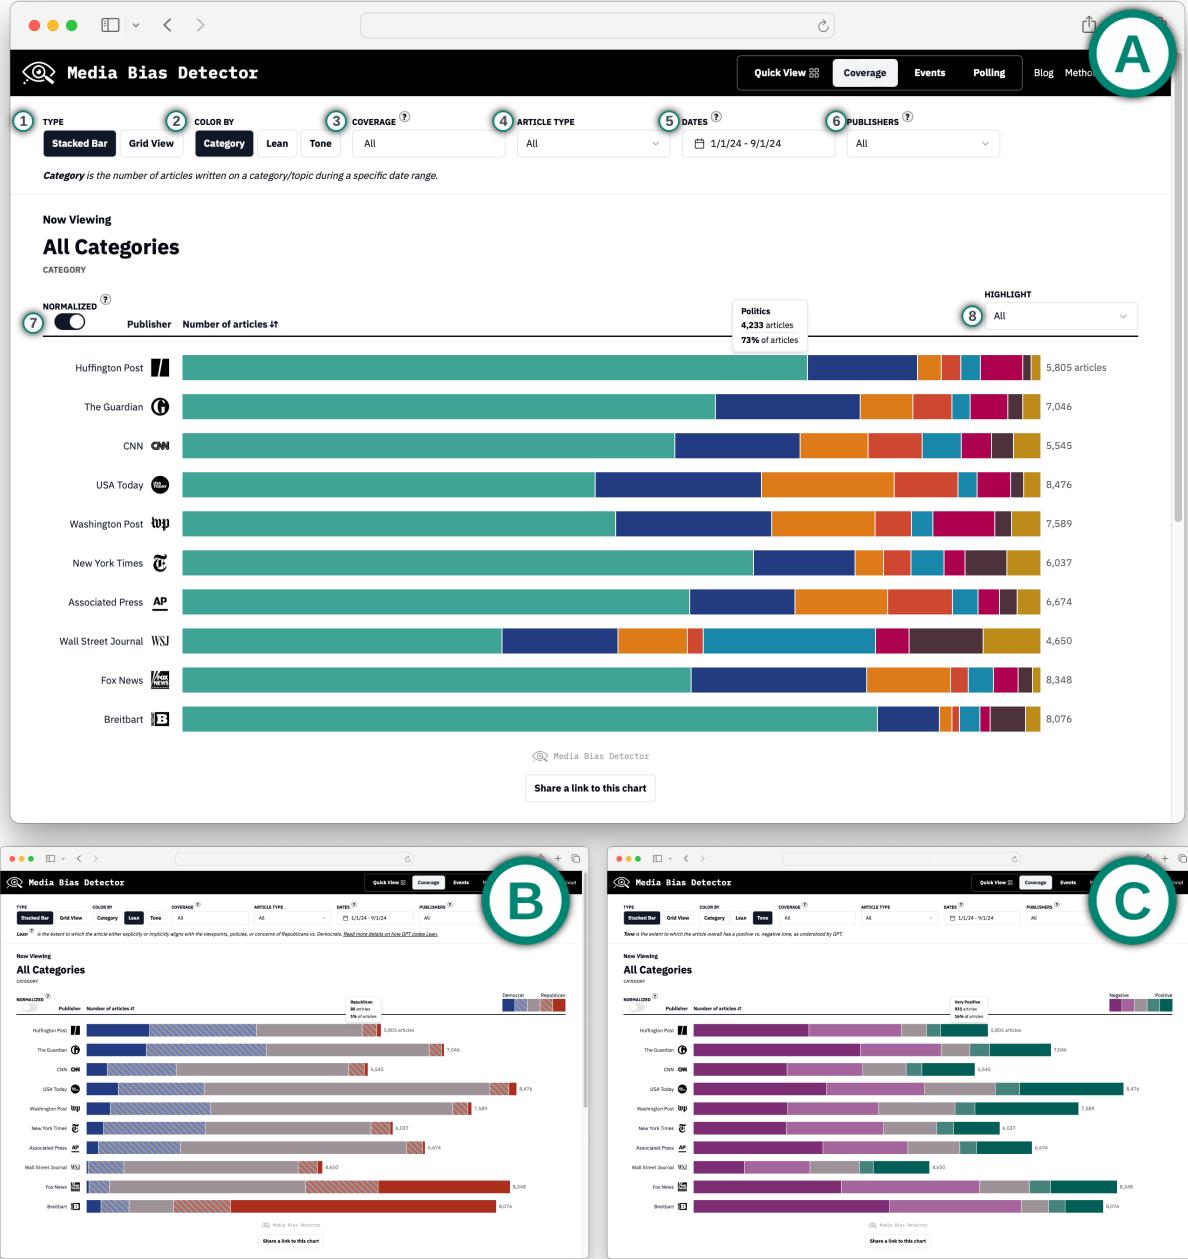
The user interface (UI) features two primary views: the *Coverage dashboard*, which visualizes the volume, lean, and tone of articles published on different topics; and the *Events dashboard*, which tracks the top events covered by publishers in the preceding three days and highlights their top facts. Events are defined as significant incidents that generate a large number of articles within a short time frame. We detect them algorithmically by clustering similar articles published within a one-day window.<sup>3</sup> The underlying mechanics of the MEDIA BIAS DETECTOR rely on large-scale automated analysis, using OpenAI's GPT-4o with human-in-the-loop review to classify articles by topic, subtopic, article type, tone, and political lean. To accommodate the scale of the data and the need to efficiently run analyses every day, we use AWS S3 for storage and utilize OpenAI's asynchronous API to process requests in parallel and achieve maximum throughput within existing rate limits. Additional methodological details are provided in Appendix A.4. Here, we examine how users interact and engage with the design of an LLM-driven media bias tool by focusing on the features and affordances users find most valuable (RQ1), how the awareness of LLMs being used affects their trust in the tool (RQ2), and which user groups benefit most from using the tool's more sophisticated set of features (RQ3).

### 4.1 Coverage Dashboard

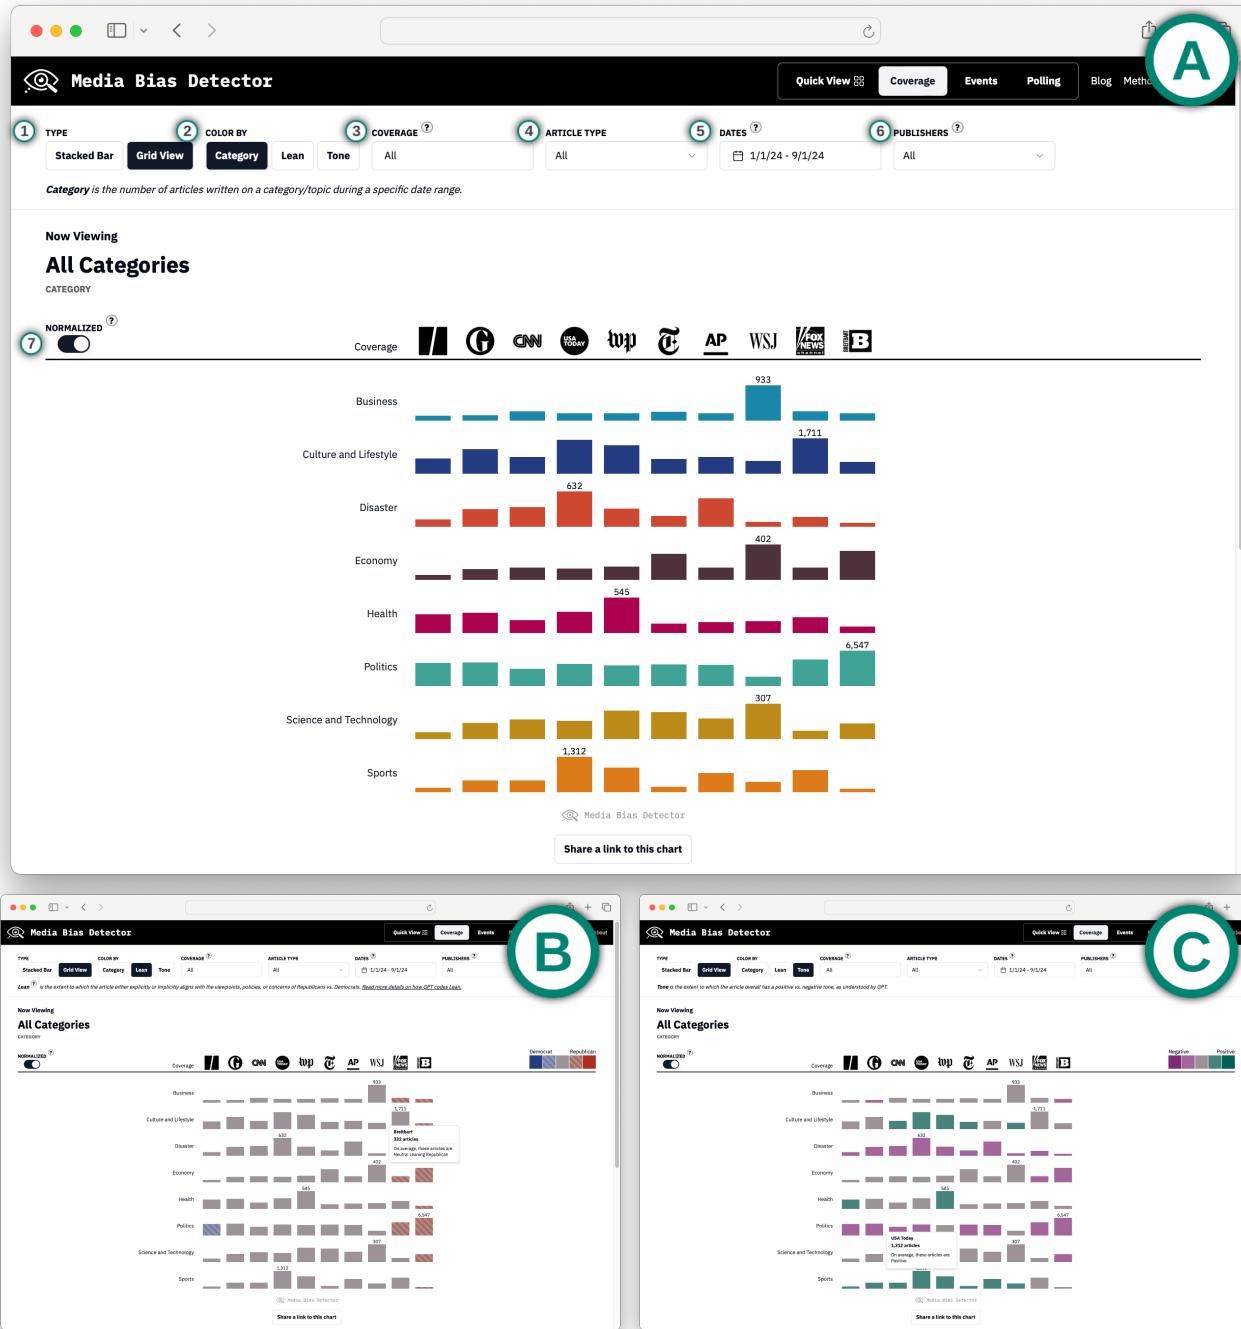
The Coverage view forms the default landing page of our dashboard and presents an overview of the topic-level and subtopic-level news coverage since the beginning of our data collection (January 1st, 2024). Figure 1 shows this view which takes the form of a stacked

<sup>2</sup>From this point forward, we refer to these design considerations as D1, D2, and D3 for clarity and consistency.

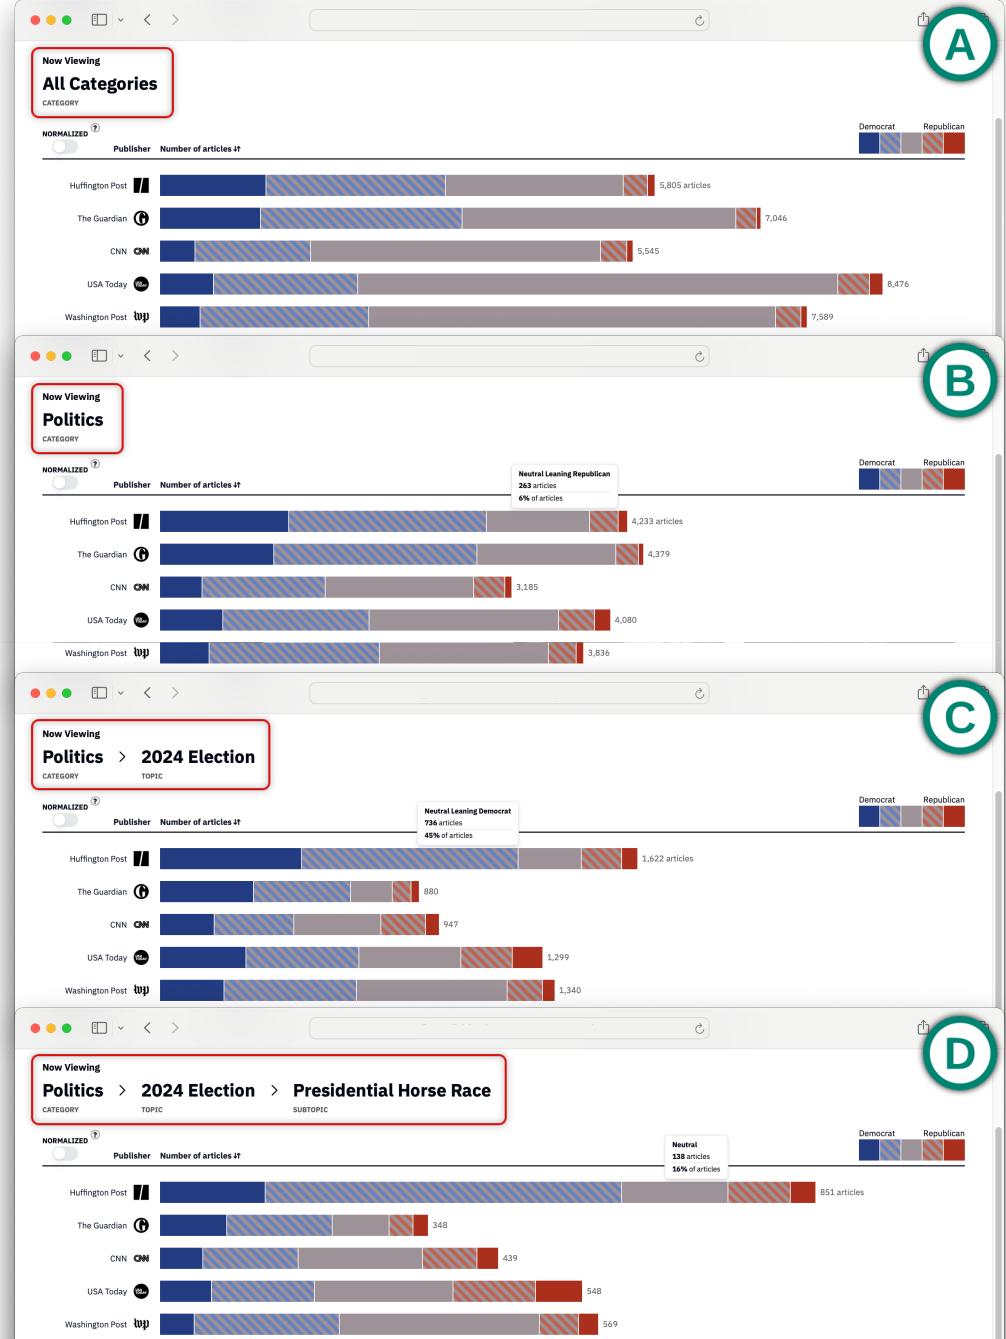
<sup>3</sup>Initially, a three-day window was used, but this was later adjusted to one day to better capture emerging events in today's fast-paced news environment.



**Figure 1:** The default view of the Coverage dashboard which allows *broad exploration* of the data (D1). Every user lands at a screen showing the category-wise coverage of the news publishers (A). Each publisher is represented with a stacked bar where a segment represents the number of articles published by them in a given category, allowing them to directly compare the proportion of attention they give to those topics (D3). Users are provided with a variety of controls to adjust the chart type and color, and to filter on publishers, article type, and date range (D1). Users can also toggle normalization off to allow direct comparisons of absolute numbers instead of the proportion of coverage. Coloring by Lean (B) shows the distribution of articles that are more aligned with a given political viewpoint, and coloring by Tone (C) shows the variation in sentiment across the same articles. Hovering on a segment displays a tooltip explaining what it represents and presenting the count and proportion of articles that fall within it.



**Figure 2: The grid view of the Coverage dashboard presents an alternative visualization to the stacked bar in Figure 1 by giving each bar segment its own cell in a grid. This allows similarly broad exploration (D1) by enabling more direct comparisons between different categories and publishers (D3) without the need to hover or click on a news category to highlight it. Each cell shows how many articles were published on a given topic by a particular publisher, with the highest-publishing publisher highlighted by its article count. The news category color map for this view matches that of the stacked bar chart. Similar to that view, the bars can be colored by Lean (B) or Tone (C), where the cell color represents the average political lean or tone of the articles in that category. Hovering on a cell displays a tooltip explaining what it represents and presenting the count of articles that fall within it.**



**Figure 3: To enable deep exploration of specific data (D2), we allow users to click through a hierarchy of news topics and subtopics to zoom into news of interest to them and be able to compare the volume, tone, and lean across publishers and date ranges. In this sequence of images, we show the user interacting with the dashboard to focus on the ‘Presidential Horse Race’ subtopic (D) after starting with the default all-category view (A) and then clicking on the ‘Politics’ category (B), the ‘2024 Election’ topic within that, and finally the horse race subtopic within it (D). At each level of interaction, the user is provided the same controls to color by Category, Lean, or Tone, to filter by publishers and article type, and to select a date range to focus on (D1). Hovering on a segment displays a tooltip explaining what it represents and presenting the count and proportion of articles in it.**

bar chart coupled with a number of controls that enable users to slice the data along different dimensions (**D1**) or filter it to focus on a specific publisher/topic/time frame (**D2**). Each bar in this view represents the number of articles published by a particular news source on the selected topic(s) within a given time frame. The array of controls at the top gives users the ability to change the visualization parameters in various ways. The first two toggles (① and ②) allow users to vary the type of visualization (stacked bar or grid view, see Figure 2) and toggle between different types of article labels (category, tone, or lean), respectively (**D1**). Figure 1B shows the lean and Figure 1C shows the tone labels for the same set of articles.

Control ③ allows users to select a particular category, topic, or subtopic to focus on and can be used to zoom into specific news articles (**D2**). This functionality is shown as an interactive sequence in Figure 3 which begins with all news categories, then zooms into ‘Politics’ to ‘2024 Election’ to ‘Presidential Horse Race’. Similarly, control ④ allows users to pick the article type (news report, news analysis, or opinion) and the final two drop-down menus let users select a time frame and subset of publishers to focus on. The *Normalized* toggle (control ⑦) allows users to directly compare the proportion of coverage on a given topic between different publishers, as shown in Figure 1A. Disabling this option allows users to see the absolute number of articles published on a given topic (Figure 1B and 1C).

In this stacked bar chart view, we use five colors to represent five levels of political lean (‘Democrat’, ‘Neutral Leaning Democrat’, ‘Neutral’, ‘Neutral Leaning Republican’, and ‘Republican’) and tone (‘Very Negative’, ‘Negative’, ‘Neutral’, ‘Positive’, ‘Very Positive’). The length of each color segment is proportional to the number of articles with that label, allowing for a comprehensive view of the label distribution. Users can hover on any segment to see the number of articles that fall within that label, as can be seen in Figures 1B and 1C.

Figure 2 shows the grid view of the Coverage dashboard, accessible via control ②. Instead of using stacked bars, this view separates each segment into individual rows, with the bar size in each cell indicating the number of articles a news source has published on that topic. The view provides a broad overview of the news coverage and allows users to easily see at a glance what publishers prioritize which topics over others (**D3**). The publisher with the highest number of articles on a given topic is represented by the tallest bar in the row, with the exact article count displayed at the top. This enables users to compare and interpret the relative heights of other bars accordingly. Similar to the stacked bar view, users can also see the lean and tone versions of this grid as in Figure 2B and 2C. Here, the colors indicate the average political lean or tone of each publisher’s articles within a given news topic or subtopic. The grid view’s summary statistics provide a complementary perspective to the full distribution of labels shown in the stacked bar chart.

To enable *deep* exploration of the data, users can select from controls ③, ④, and ⑤ (see Figures 1A or 2A) to filter the full set of articles in our database to the topics/subtopics, publishers, and time frame they are interested in (**D2**). The sequence in Figure 3 shows what interacting with control ③ looks like. Clicking on ‘Politics’ takes a user from the view in Figure 3A to the one shown in Figure 3B. Subsequently selecting the ‘2024 Election’ topic within that

category, and then selecting the Presidential Horse Race subtopic within that, takes the user through the views shown in subfigures 3C and 3D. At each level, users can choose to filter according to their desired time frame and subset of publishers, and also choose to color the data by category or tone instead of lean (**D1**).

## 4.2 Events Dashboard

The Events view of our dashboard is intended to capture the fast pace of the news cycle, providing users with an overview of major events and key information highlights. Figure 4 shows this view where rows represent different events, sorted in descending order of the number of articles written about them in the past day. If an event is covered by a particular publisher, its corresponding square will be colored turquoise; otherwise, it will be empty. At the top of the page, we provide users with controls to select the date range (control ①) and set of publishers to display (control ②).

To provide a more in-depth view of these events (**D2**), users can click on the arrow at the end of each row (control ③) to expand the event and see a more detailed description about it, as shown in Figure 4B. These descriptions include a longer title which more accurately summarizes the event, as well as overall statistics on the types of sentences found in articles about it (facts, quotes, or opinions). This view also displays the top facts about the event that have been mentioned by different publishers in their coverage of it. As shown in Figure 4C, clicking on the arrow next to a top fact (control ④) displays various phrasings of it from different articles, and clicking on any article listed takes the user directly to the story on the publisher’s website. This functionality, grounded in comparative analysis (**D3**), allows a user to compare and contrast not only how much coverage (selection bias) has been given to an event but also how this coverage has been presented (framing bias).

## 5 Evaluating the MEDIA BIAS DETECTOR

To understand how the MEDIA BIAS DETECTOR could benefit experts who study news and media, we conducted interviews with participants from academia, industry, and newsrooms, all of whom work closely with news content, but approach it from different disciplines. In addition to need-finding, we also conducted a within-subject evaluation study where we compared using our tool to baseline media bias detection strategies. This evaluation helped us understand how the MEDIA BIAS DETECTOR complements existing tools, and we used feedback gathered from this study to inform improvements for the tool’s next iteration.

### 5.1 Participants

We recruited 13 participants (7 women, 6 men, aged 24–63).<sup>4</sup> Potential candidates were identified through personal networks, including referrals from colleagues, friends, and word of mouth. To ensure unbiased evaluations, participants were selected specifically for their lack of prior exposure to the tool. This design decision allowed us to collect immediate feedback on the tool’s ability to engage and effectively communicate its purpose during a user’s first

<sup>4</sup>Two participants, P10 and P12, were unable to complete the second task evaluation and post-task interview due to time constraints, and we were unable to reschedule their interviews prior to the submission. Their evaluation results are dropped from the quantitative analysis, but we still discuss insights from their interviews in our findings.

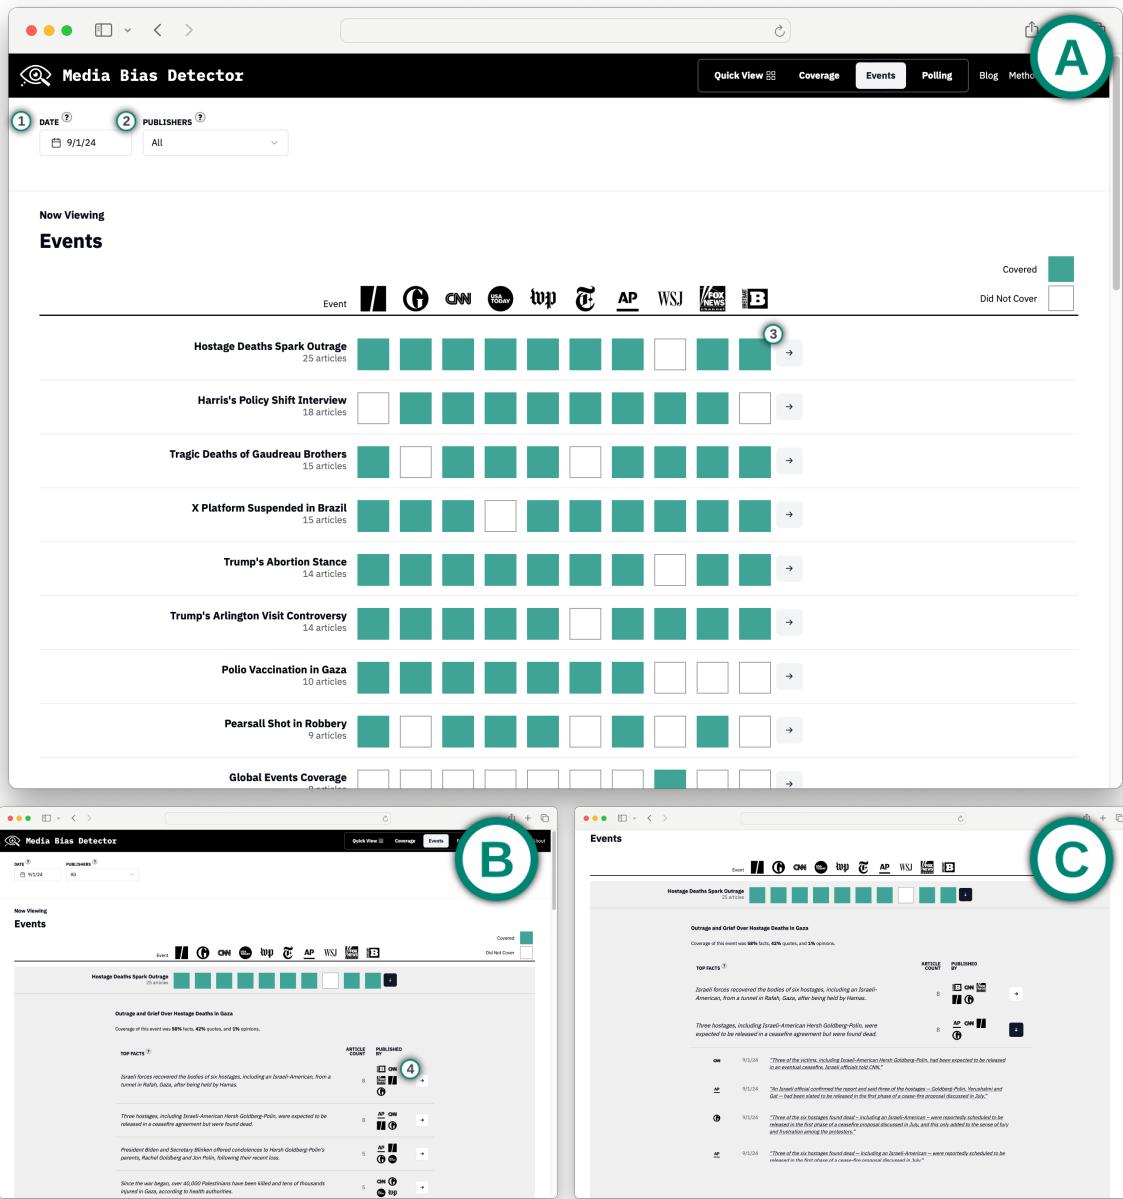


Figure 4: Complementary to the Coverage view, the Events dashboard shown here presents an event-level view of the news that focuses on the fast pace of the news cycle and caters to both *broad* (D1) and *deep* (D2) exploration of the news. It offers a quick overview of happenings from the past three days but allows users to dig deeper into each event and its major facts and compare selection and framing bias across publishers (D3). Each row is a news event and each cell represents whether it was covered or not by the respective publisher. The events are sorted by importance (measured as their amount of coverage) and we provide a summarized title for the event and a count of the number of articles about it across all publishers. Clicking on an event using the button on the right displays a detailed view (B) which contains the full event description as well as a summary of its sentence-level composition in terms of facts, quotes, and opinions. This view also lists the top facts about the event and shows which publishers mentioned or omitted certain statements. Clicking on any one of these top facts shows different variations of it (C) as they were written in the original news articles allowing users to compare their framing. Clicking on any one of these variations takes the user to the original article on the publisher’s website.

interaction. By focusing on first-time users, we were able to identify potential areas for improvement that might be less apparent to users more familiar with the tool. All participants had undergone rigorous training as either researchers or writers for at least three years, enabling them to provide well-thought-out perspectives on news media. Their areas of expertise spanned various domains, including communications, political science, and journalism, with differing levels of familiarity with media bias. Detailed information about the participants is provided in Table 2.

## 5.2 Task Design

Each participant engaged in two consecutive tasks, where the first task required the use of a baseline online search and the second used the MEDIA BIAS DETECTOR. These tasks were designed to simulate a realistic scenario in which participants evaluate media bias across different news publishers, setting up a practical comparison between traditional methods and our tool.

In the first task, participants were presented with a scenario where they imagined discussing with friends the news sources they followed for the 2024 election. Participants selected two publishers to analyze from a predetermined list of ten publishers, corresponding to the ten publishers on the dashboard, described in Section 4. To allow for a more personalized experience, participants were given the flexibility to choose two publishers from the list that aligned with their own interests or questions they wanted to explore. This approach prevented potential biases that might arise from pre-selecting specific publishers, which might have influenced the types of insights participants could generate. It also allowed us to capture variations in themes, editorial choices, and general nuances across publishers. Participants were asked to explore how selection and framing biases manifested in these outlets and influenced their news coverage. They were given 10 minutes to evaluate the overall bias of the two newspapers using any available tools or platforms known for assessing media bias and prepare a brief oral summary comparing the two. For the second task, participants engaged in the same scenario, evaluating the same two newspapers, but this time using our tool. We provided participants with a URL to access the MEDIA BIAS DETECTOR and gave them 10 minutes to assess the overall bias of the publishers (see Appendix A.1 for task scenarios).

After each task, participants answered a set of structured Likert-scale questions designed to gather quantitative feedback on their experience, focusing on the affordances of the tools, how effectively they supported users in recognizing bias, making comparisons across publishers, and engaging with media content (outlined in Section 5.3). The evaluation after Task 2 also included Likert-scale questions aimed at capturing participants' reflections on the MEDIA BIAS DETECTOR's usability and intuitiveness (see Appendix A.1 for survey questions). This two-task design ensured consistency in the evaluation process while providing participants an opportunity to explore and assess the affordances of both an existing tool and our proposed tool for detecting media bias.

## 5.3 Quantitative Measures

We used a mixed-methods approach to evaluate the MEDIA BIAS DETECTOR. In addition to collecting qualitative feedback on how

participants used our tool, we also measured key dimensions of the tool's impact: bias identification and awareness (RQ1), comparative analysis (RQ1), and user engagement (RQ3). To quantify these aspects, we asked participants to rate their agreement on a 7-point Likert-scale (1 = Strongly Disagree, 7 = Strongly Agree), similar to prior work [86, 105]. Furthermore, to address RQ1, we assessed subjective workload using the NASA-TLX procedure with weighting [44], a method commonly used in HCI research to measure a user's cognitive demand when performing a task.

### (1) Bias Identification and Awareness

- **Selection Bias:** The news gives disproportionate attention to specific topics.
- **Framing Bias:** The way news stories are presented, including the choice of language, influences how people perceive events being reported.
- **Bias Awareness:** This method sharpened my awareness of bias in news content.
- **Critical Thinking:** This method encouraged me to think critically about the news articles I read.
- **Effectiveness:** This method was effective at helping me identify media bias.

### (2) Comparative Analysis

- **Qualitative Comparison:** This method offered useful descriptive information for comparing bias across different publishers.
- **Quantitative Comparison:** This method offered useful quantitative measures for comparing bias across different publishers.

### (3) User Engagement and Application

- **Proactive Sharing:** I am interested in sharing and discussing the information I've learned about media bias with my family and friends.
- **Practical Application:** I am interested in applying the information I've learned about media bias in my daily news consumption.

## 5.4 Procedure

The study sessions began after participants signed a consent form and completed a survey that collected demographic data, professional experience, and self-perceived familiarity with media bias. The study was conducted via one-on-one Zoom calls, with participants granting permission to record the session. Each interview lasted between 45 and 75 minutes. The session started with a brief introduction, during which the interviewer informed the participant that their feedback would be used to improve a tool designed to help news consumers navigate the vast amount of news content, particularly during the 2024 election cycle.

During the pre-task interview, we explored participants' general media consumption habits, preferred news sources, and their engagement with news and politics. We also asked for their views on how news outlets may emphasize certain topics while overlooking others (selection bias) and how the tone and language of articles can influence readers' perceptions (framing bias). Next, participants undertook the sequential 10-minute media bias comparative analysis tasks. Participants were encouraged to think aloud as they performed both tasks [73, 84]. Before the MEDIA BIAS DETECTOR task,

**Table 2: Participant Characteristics in the User Study with Experts.** *Current Role and Area of Expertise* are self-reported descriptions. *Exp. (Yrs.)* refers to the participant's experience working in their domain. *Bias Familiarity* refers to the participants' self-perceived media bias familiarity.

ID	Domain	Current Role	Area of Expertise	Exp. (Yrs.)	Bias Familiarity
1	Communications	Pre-doctoral Research Assistant and PhD Student	Digital Political Communication	3-5	High
2	Communications	Pre-doctoral Research Assistant	Environmental Communication	3-5	High
3	Communications	PhD Student	Tech Innovation for Cultural Protection	3-5	High
4	Journalism	Writer	Political Journalism	5-10	High
5	Communications	PhD Candidate	Political Communication, Journalism, AI	10+	High
6	Political Science	PhD Candidate	Political Behavior, N. American Politics, Misinformation	3-5	Medium
7	Communications	PhD Candidate	Social Media, Health Communication	3-5	Medium
8	Journalism	Director of News & Media	Editorial Leadership, Tech in Journalism	10+	High
9	Political Science	PhD Candidate	Racial and Ethnic Politics, Black Politics	5-10	Low
10	Journalism	Retired; Former Director of Content and Managing Editor	Digital News, Social Media, Leadership	10+	High
11	Journalism	CEO of Journalism/Content Consulting Company	Journalism and Content Strategy	10+	High
12	Journalism	Guideline Architect; News & AI Strategist	AI & Content Strategy, Policy Development	10+	High
13	Journalism	Media Reporter	Media Studies	5-10	High

we showed a brief walkthrough video highlighting its main features. At the end of each task, we provided participants with a survey link to complete the evaluation questions, including Likert-scale and NASA-TLX questions.

The session concluded with a semi-structured discussion about their experiences using the tool. The discussion covered topics such as the tool's potential impact on the participant's work or research, its usefulness in the context of the 2024 election, and how it compared to other tools they had previously used. Participants were also asked to suggest additional features and provide feedback on improvements that could enhance their experience with the tool (see Appendix A.1 for guiding questions).

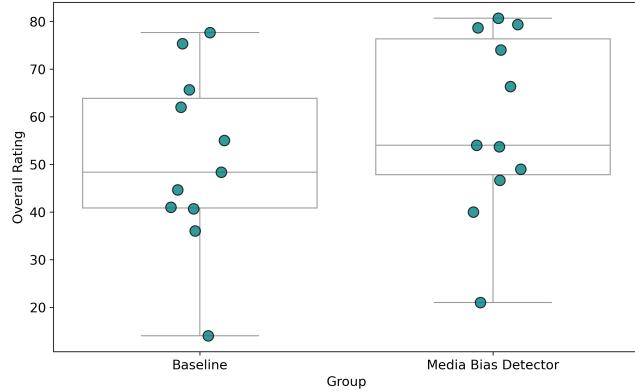
## 5.5 Analysis

We collected interview transcripts, observation notes, and a set of quantitative responses for each task. We coded and analyzed the qualitative data using Miro [64], which allowed us to collaboratively code the data using digital sticky notes [22, 40, 50]. Four team members manually reviewed the interview transcripts noting key points. From these notes, we used an inductive coding approach to generate initial descriptive codes and identify relationships between themes discussed in participant quotes [17]. Subsequently, our team reviewed the codes and preliminary themes with the assistance of an AI summarization tool and had multiple group discussions to note similarities and differences before agreeing on the main broad themes we learned from our expert interviews [62].

After analyzing the data, we reviewed the quantitative measures collected, including the NASA-TLX scores and Likert-scale

responses. The boxplot in Figure 5 shows the overall results of the NASA-TLX ratings for subjective workload, comparing the baseline tool with the MEDIA BIAS DETECTOR. While the MEDIA BIAS DETECTOR received slightly higher overall ratings, indicating that participants found it somewhat more challenging to use, the difference is not substantial. The overlap between the distributions suggests that the MEDIA BIAS DETECTOR's task workload is comparable to the baseline tool, and several participants rated both tools similarly. These results are encouraging for the MEDIA BIAS DETECTOR, as the tool is inherently more complex and offers more features than typical baseline tools, including those selected by participants. Despite this added complexity, participants did not find it overwhelmingly harder to use. This suggests that with continued improvement, such as providing more tutorials or onboarding, we can reduce the tool's perceived difficulty and help users take full advantage of its advanced features without significantly increasing the cognitive load.

The results of the Likert-scale responses in Figure 6 show minimal variation between the baseline tool and the MEDIA BIAS DETECTOR for most metrics, suggesting that short-term use of the tool may not lead to significant shifts in users' overall perceptions of media bias. The subplot describing *M7: Quantitative Comparison* in Figure 6 shows that only responses to "this method offered useful quantitative measures for comparing bias across different publishers" had a notable increase in the MEDIA BIAS DETECTOR group. Given that media bias perceptions are often deeply ingrained and difficult to shift in a short time frame, these findings align with our expectations. Our small sample size also may have contributed to



**Figure 5: Overall results on the NASA-Task Load Index (NASA-TLX)’s measures of subjective workload for the Baseline and Media Bias Detector task evaluations in the user study with experts.**



**Figure 6: Average responses to Likert-scale measures on bias identification and awareness (M1-5), comparative analysis (M6, M7), and user engagement (M8, M9) in the user study with experts.**

the lack of significant differences across most metrics. However, the results generally reinforce the idea that the true value of the MEDIA BIAS DETECTOR may emerge over long-term use, as users engage more deeply with its features and gain more exposure to its capabilities. We explore this further in our follow-up survey (see Section 7) that focuses on other dimensions, including users’ long-term beliefs about the tool’s potential impact, rather than just the measures evaluated in this initial study.

## 6 Findings

The qualitative data from our interviews provided rich insights, addressing each of our main research questions while also uncovering additional findings. Common themes such as usability and educational potential offered a deeper understanding of the tool’s strengths, contributing to RQ1. Other themes, such as trust in AI-driven tools and identifying the most suitable audience for the

MEDIA BIAS DETECTOR, aligned closely with RQ2 and RQ3, respectively. In this section, we explore the major themes that emerged from our expert interviews.

### 6.1 RQ1: Functionality and Impact of the MEDIA BIAS DETECTOR

*Interactive Presentation of Complex Information.* Overall, participants were impressed that the MEDIA BIAS DETECTOR transformed a complex dataset into an easy-to-navigate interface. In particular, the tool’s emphasis on progressive disclosure allowed participants to engage with the content at varying levels of depth (**D2**). P4 expressed that they *“appreciate the effort to make the subtopics super granular”* after diving deeper from the *Politics* category, to the *2024 Election* topic, to the *Presidential Horse Race* subtopic on the Coverage dashboard. Others expressed that the visualizations effectively communicated information when needed, making it easier to focus on coverage they were interested in. P3 especially liked how *“the tool can be as interactive as you need it to be”*.

Regarding **D1**, multiple participants mentioned the value of being able to examine the data from both the lean and tone perspectives. While traditional media bias tools focus on average political lean scores, P3 highlighted the unique distinction between lean and tone, emphasizing the importance of going beyond political lean. P8 expanded on this, pointing out that tone is a less talked about form of media bias: *“[The bias is] not about negativity, but rather the lack of solutions journalism. For decades, people have wanted more context”*. Furthermore, P4 stated outright that the MEDIA BIAS DETECTOR’s ability to group data by both tone and lean offers a more comprehensive way of assessing media bias: *“It’s always hard to know on what basis [people building those tool] are making those decisions... This is a huge improvement because it allows you to select different forms of evaluation”*, contrasting it with tools like Allsides, which typically use charts to map news outlets’ political lean from left to right.

Participants also highlighted opportunities for improvement. P4, for instance, valued the range of features but remarked, *“there was a lot of information in front of me, and that was overwhelming”*. Others saw potential in enhancing the user experience by introducing a search bar feature which allows users to more precisely search for specific topics by keyword (P1, P6). P6 suggested that this feature could reduce some users’ cognitive burden by eliminating the need to actively scan through topic and subtopic options.

*Interactions Shaping Media Bias Perceptions.* Being able to customize the data not only personalized the experience with the tool but also influenced how participants perceived bias in news coverage. For example, P3 described how the process of focusing on particular topics and within specific date ranges helped them challenge their preconceived notions of some publishers’ biases: *“I come in [with] really strong opinions on what the biases are for each of these outlets... because I got to choose the topics and choose the time frame and choose looking across the lean and tone for different topics, ... It helped me to ground my thinking in real world evidence”*. Similarly, P9 praised **D3**, noting that even without explicit bias ratings, the Events dashboard allowed them to easily compare different news sources and independently analyze how they framed the same events: *“It’s nice to see how they’re covering them differently in one*

*place... It's still helping me assess bias because it's putting everything in a place where I can look at it and seek out broader coverage, look, and assess it for myself.*

While some participants mentioned ways the MEDIA BIAS DETECTOR could help challenge and expand their current understanding of media bias, not all participants described a significant shift in their perceptions. Instead, many participants shifted the discussion towards how the tool might help others learn about media bias, particularly those less experienced than themselves.

## 6.2 RQ1 and RQ3: Potential Value in Education and Research

**RQ1: Applications in Education.** Several participants, particularly those studying journalism and communications, highlighted the tool's potential as a resource in media literacy classrooms (P1, P2, P3, P7). P3 expressed that the MEDIA BIAS DETECTOR would be very useful for media literacy classes, whether at the college level or in more public-facing settings. Furthermore, P1 and P2, both communications students, expressed interest in sharing the tool with journalism professors at their respective institutions. Beyond the classroom, participants also highlighted the MEDIA BIAS DETECTOR's ability to educate users to recognize blind spots in news coverage. P2 emphasized that the tool provided a straightforward way to demonstrate that one news organization covers an issue more or less than another. On the Events dashboard, P9 highlighted an example of what an everyday user could learn from the MEDIA BIAS DETECTOR; during the Gaza ceasefire negotiations, only seven major articles were published, while most coverage focused on the Democratic National Convention. P9 found this quantitative measure useful in showing how a significant event was being overshadowed, stating that *"Americans are kind of distracted by the DNC and aren't able to pay attention to something that's really important, and that is reflected quantitatively here"*.<sup>5</sup>

In addition to recognizing the tool's potential for media literacy education, some participants noted that its full benefit depends on users having a baseline understanding of media bias. P7 suggested the tool would be most impactful when paired with education on bias, while P13 emphasized that users would have a better understanding of the MEDIA BIAS DETECTOR if the experience could start with their own personal beliefs. More specifically, P13 recommended integrating a feature that allows users to reflect on how they currently view a topic, and then explore the news organizations that are covering it in a way that fits their own understanding. These participants emphasized a key point: *"education about what media bias is and how it can affect them actually is the most important step for people to actually use these tools, because if they don't know what that does to them, they wouldn't really care"* (P7). For everyday users, simply providing access to the MEDIA BIAS DETECTOR may not be enough; they need foundational knowledge to fully engage with its features, which we reaffirmed in our follow-up survey (Section 7).

**RQ3: Applications in Research.** The MEDIA BIAS DETECTOR was also recognized as a valuable tool for research, particularly for

<sup>5</sup>P9 was referring to the Events page on Wednesday, August 21, where 33 articles on the top news event were related to the 2024 DNC, while only 7 articles were related to the Gaza Ceasefire Negotiations.

supporting mixed-methods communications research. P1 and P2 highlighted its strength in providing quantitative comparisons that can help set the stage for more in-depth qualitative analysis (D3). P2 noted that the MEDIA BIAS DETECTOR offered a straightforward solution for obtaining quantitative evidence, saving time compared to relying on specialized organizations like Media Matters or searching through existing studies for relevant data. *"[Need] to justify a news analysis?"* P2 declared, *"here's results from a tool that justify it"*. In political science research, P6 and P9 proposed that the MEDIA BIAS DETECTOR could be used as an experimental stimulus in studies on political behavior. P9 envisioned using the tool as an intervention in behavioral experiments that study how media exposure influences political attitudes, such as assessing *"how [the tool's information] shapes your view of the American government's seriousness about ending the war in Gaza"*. Finally, to conduct high-quality research, P2 and P5 also pushed for a CSV download feature.

## 6.3 RQ2: Challenges Surrounding Trust and Transparency in LLM-Driven Bias Detection

Participants recognized the MEDIA BIAS DETECTOR's potential for shaping education and research but also expressed reservations about fully trusting the LLM-based classifications. For some, ensuring that humans were reviewing the data was the most important factor in building trust. P7, for example, immediately asked if human annotators were cross-checking the LLM's classifications and emphasized that it was important for them to know that real people were involved in the process.

Their skepticism was rooted in doubts about the LLM's ability to capture the nuances of media bias, particularly in sensitive topics, such as child tax credits or international conflicts (P1, P9). P5 pointed out that while they felt comfortable believing that LLMs can easily automate formulaic news stories like financial reports, they needed more evidence to believe that the models can properly classify subjective news content. Similarly, P9 wanted explanations for or article examples of the LLM's labels for "Strong Democrat" or "Strong Republican" on topics with varying levels of polarization, noting that political distinctions may be harder to define for more nuanced issues.

Several participants suggested that adding features to increase transparency could help build trust (P1, P2, P9). P1 expressed, *"For me, someone who is a bit more skeptical of AI research in general, being able to see the articles [labeled under a topic] would make me feel better"*. P2 also suggested that having access to article texts would help them evaluate how the LLM processes articles with mixed tones, particularly in climate change coverage, where optimistic, positive outlooks may exist within negative stories [74]. This feature would allow them to determine if the LLM captures subtle shifts in tone or oversimplifies it as neutral.

Nevertheless, many participants noted that spending more time reviewing the Methodology section might help alleviate their concerns (P2, P6, P7, P9, P11, P13). Others saw LLM-driven bias detection not as a concern but as a practical solution for scalability. P11 shared the desire to understand the MEDIA BIAS DETECTOR's inner workings but recognized the importance of AI in handling large-scale, dynamic data. During the task, they repeatedly "calibrated" themselves, using their prior knowledge to compare the

tool's lean and tone ratings across the political spectrum to feel more confident in the MEDIA BIAS DETECTOR's trustworthiness.

Overall, increased visibility into the MEDIA BIAS DETECTOR's decision-making process could mitigate the skepticism expressed by participants. Comments from multiple participants suggested that their trust in the tool was closely linked to their own ability to see, interpret, and agree with the AI model's decisions. Thus, providing more transparency about the article texts, human-in-the-loop's weekly analysis, and the LLM's classification process could help users feel more comfortable using the MEDIA BIAS DETECTOR.

#### 6.4 RQ3: Who is the User? Clarifying the MEDIA BIAS DETECTOR's Audience

A recurring theme among some participants (P8, P10, P11, P13) was an encouraging attitude toward the MEDIA BIAS DETECTOR's development. However, many did not see themselves as the primary users. P10 mentioned that they had already developed methods for evaluating bias independently because their role at a news aggregator had required assessing media bias and deciding what content to publish. Meanwhile, P8 provided a reflective observation, remarking, *"Can you teach media bias [sic] about media bias? Because a lot of media people don't necessarily think that they are biased in certain ways."* This highlights the challenge of introducing media bias tools to professionals who, confident in their expertise, may be less inclined to fully engage with a tool designed to reveal what they feel they already know.

For other experts, there was no clear consensus on the tool's target audience. Some, like P3, found the tool useful for encouraging self-reflection, noting that it helped them *"think through my own thoughts—what I thought were the biases versus what actually are the biases"*. Others, like P5, believed the tool was well-suited for researchers with a nuanced understanding of media studies but felt it might be too detailed for regular news consumers. In contrast, P1 suggested that the tool could resonate with everyday people who are seeking balanced news, especially during the 2024 election.

### 7 Follow-Up Survey: Evaluating the Broader Population

Through our 13 semi-structured interviews, we gained valuable feedback on the MEDIA BIAS DETECTOR's usability, its educational and analytical value, and trust concerns associated with LLM-driven media bias detection. It also became clear that while some participants saw value in the tool for their own use, many perceived it as being more relevant for other users rather than themselves. In particular, the feedback suggested that understanding the tool's impact on everyday news consumers would be crucial for further refining its design and assessing its broader applicability. To address this, we conducted a follow-up survey targeting everyday news consumers.

#### 7.1 Procedure

We recruited 150 participants (51% male, 46% female, 3% other<sup>6</sup>) through the crowdsourcing platform, Prolific, with the sample stratified by political party to balance Democrats, Republicans, and Independents (see Appendix A.2, Figure A2). Participants first completed a demographic survey, followed by questions about their media consumption habits, familiarity with media bias, and prior use of bias detection tools. The main survey was divided into three stages: pre-tool exposure, training and exploration, and post-tool exposure.

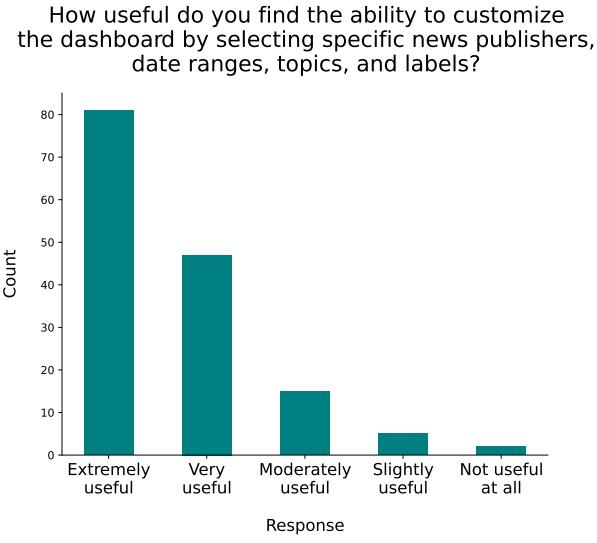
In the pre-tool exposure stage, participants answered general questions about their perceptions of selection and framing bias in the media, as well as a set of more targeted questions related to specific issues covered in the upcoming training task. Participants were asked to give their best guesses to the following questions. The questions were selected to illuminate selection biases in coverage by various news outlets on important and timely topics during the 2024 election campaign. We chose questions that might challenge people's existing misconceptions about media bias, so that observing the actual coverage on the dashboard would better capture their interest.

- (1) Which news categories tend to get mostly negative coverage?
- (2) Which media outlets covered the topic of Biden's age more frequently?
- (3) What is the political lean of The Wall Street Journal's economic coverage?
- (4) What is the most prominent event covered by the media this past week?

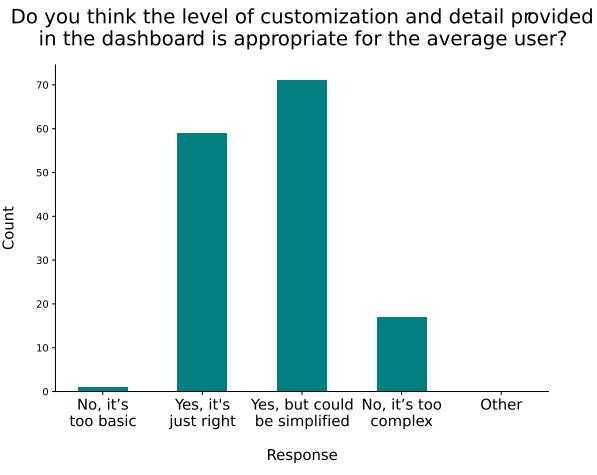
During the training and exploration stage, participants first watched a brief walk-through video of the MEDIA BIAS DETECTOR. Afterward, they followed guidelines to explore the tool and answer the same questions asked in the pre-tool exposure stage. This part of the task design aimed to teach participants how to use different features of the tool. Following this guided experience, participants were given five minutes to use the MEDIA BIAS DETECTOR to explore media bias related to a topic of their choice and report their most interesting findings. This exploration was intended to encourage deep engagement with the tool, giving participants the freedom to find patterns in the data and consider how they might use it on their own.

Finally, in the post-tool exposure stage, participants revisited the specific questions asked in the pre-tool exposure stage to measure whether the training task had any short-term impact on their responses. However, recognizing that a brief five-minute interaction is unlikely to change deeply held views about bias in the news, as discussed in Section 5.5, we also asked them directly whether they believe that using the MEDIA BIAS DETECTOR over time could change their view on media bias. Additionally, participants provided feedback on the tool's complexity, customization options, and their trust in AI-driven classifications (see Appendix A.2 for survey questions).

<sup>6</sup>The "other" category comprises two participants who preferred not to disclose their gender and two participants who self-identified as non-binary. This categorization does not significantly impact the broader analysis, as gender is not a primary parameter in our study.



(a) Customization usefulness.



(b) Customization appropriateness.

**Figure 7: In response to questions about the MEDIA BIAS DETECTOR's customization features, the vast majority of participants stated that they found these options to be 'extremely' or 'very' useful. However, when asked whether this level of customization would be appropriate for an average user, most recommended that it could be simplified, suggesting a steep learning curve for the tool's current design.**

## 7.2 Analysis

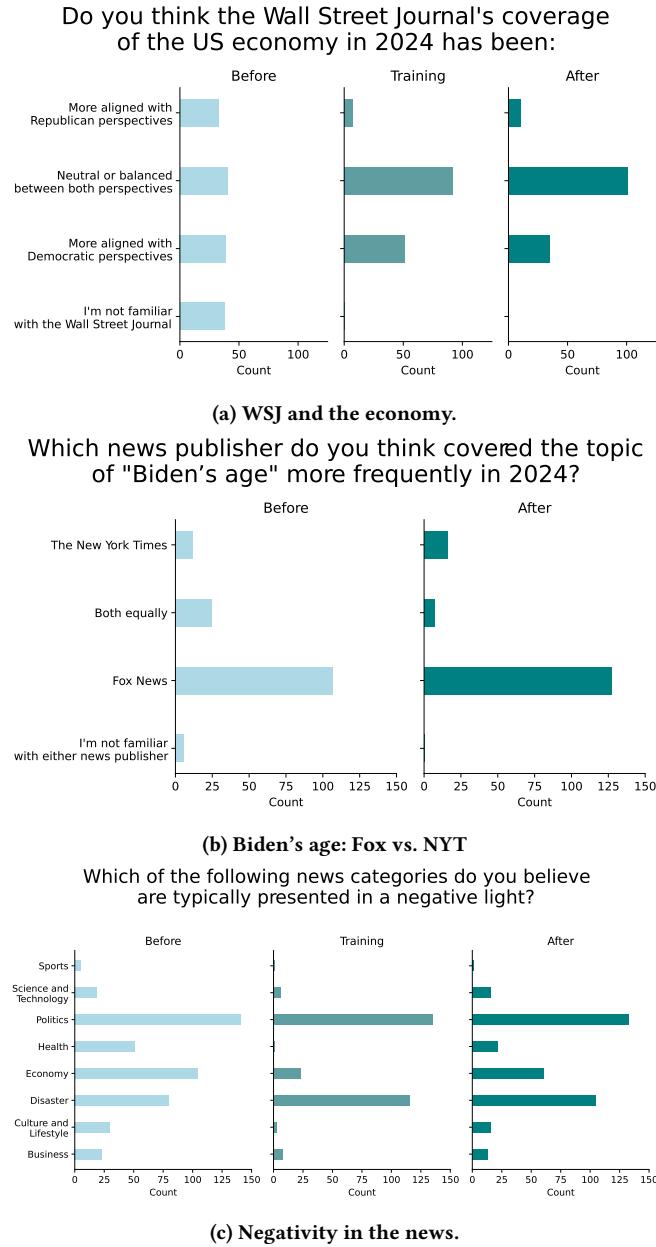
**User experience.** One of the primary features of the MEDIA BIAS DETECTOR is the ability to view the data from many different perspectives (D1). To understand whether ordinary users like this feature, we asked them if they found this ability to customize the visualizations useful, and whether they thought that this level of customization was appropriate for other ordinary users. Figure 7 shows their responses. A little over half the participants (81/150) said that they found the provided customization 'extremely useful',

while an additional one-third (47/150) said that they found it 'very useful'. Interestingly, when asked whether they thought this level of customization and detail would be appropriate for the average user, around 40% of respondents (59/150) said it was just right, but a slightly larger 47% of respondents (71/150) said that it could be simplified.

We also asked for respondents' feedback on any features they found unclear or challenging to use, as well as suggestions for improvements. This input further mirrored the variation between 'just right' and 'simplified' seen in Figure 7b. Some found it easy, with one saying, *"the truth is I think it is very easy to use and learn,"* while others felt overwhelmed, noting, *"there is so much to take in just in this one setting"* and *"Still absorbing what is here"*. Participants also provided a wide range of ideas, suggesting improvements to the UI (*"When I hover over different clickable options, it'd be nice to be reminded what each option means"*), data visualizations (*"average person will probably want something more visual... Perhaps a wheel or pie chart structure"*), and usability (*"I only want to do 2 or 3 clicks to get my information"*).

**Impact of using the MEDIA BIAS DETECTOR.** We also investigated whether and to what extent the MEDIA BIAS DETECTOR can change a user's perceptions of media bias. To do this, we asked participants about their views on news coverage for certain topics in the pre-tool exposure stage, had them answer the same questions during the training stage, and then asked again in the post-tool exposure stage. Figure 8 shows the impact using the MEDIA BIAS DETECTOR had on their responses. In Figure 8a, we asked users what they imagined the political leaning of the Wall Street Journal's coverage of the economy looked like in 2024. We received a fairly even response across the board: while a quarter of the respondents (38/150) were not familiar with the publisher, the rest were almost evenly split between pro-Democrat, pro-Republican, and neutral coverage. The training phase revealed that the MEDIA BIAS DETECTOR classified a majority of the Wall Street Journal's economic coverage as neutral, and most respondents arrived at the same conclusion after using the tool. Furthermore, this had a significant impact on their post-training responses, with over two thirds of respondents (101/150) correctly stating that the publisher's coverage was indeed neutral. Some participants noted their surprise at learning this in a free-response question later in the survey: *"I find it interesting that the Wall Street Journal really does seem to have the most balanced takes on business matters"*.

In some cases, however, the data found on the dashboard was less surprising and mostly confirmed people's pre-existing notions. Another question we asked survey participants was who they think covered the topic of President Joe Biden's age more: the New York Times or Fox News (D3). Figure 8b shows the result: 71% of respondents (107/150) answered Fox News, the correct answer, even before using the tool. During the training stage, participants were asked to provide the exact number of articles on Biden's age published by each of the two publishers, and 90% of the participants (135/150) provided the correct answer using the tool. When asked the same question in the final stage of the survey, some participants adjusted their responses accordingly, resulting in 85% of them (127/150) providing the correct answer in the end.



**Figure 8: Participants showed a significant shift in their responses after using the tool, either correcting their priors where the dashboard's results showed otherwise, or strengthening them when the results matched their predictions.**

Finally, using the MEDIA BIAS DETECTOR can help people tune their priors to be more in line with observed data. We asked participants which topics they thought were generally presented in a negative light in the news, where they could choose multiple topics at once. As shown in Figure 8c, most participants correctly chose topics such as Politics (141/150) and Disaster (80/150)<sup>7</sup>, but many

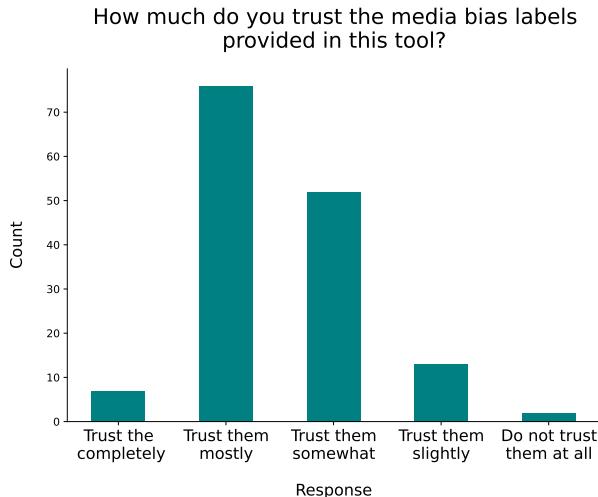
<sup>7</sup>Note that participants could select multiple topics.

also selected Health (51/150), Economy (104/150), and Culture and Lifestyle (30/150), which are generally more neutral or positive. After using the tool, they encountered labels indicating that only Politics and Disaster had an overall negative tone across publishers, while Economy was covered negatively by some of them. Post-training, participants updated their answers accordingly, resulting in an increased number of votes for Disaster (105/150) and decreased number for Economy (61/150), Health (22/150) and Culture and Lifestyle (16/150).

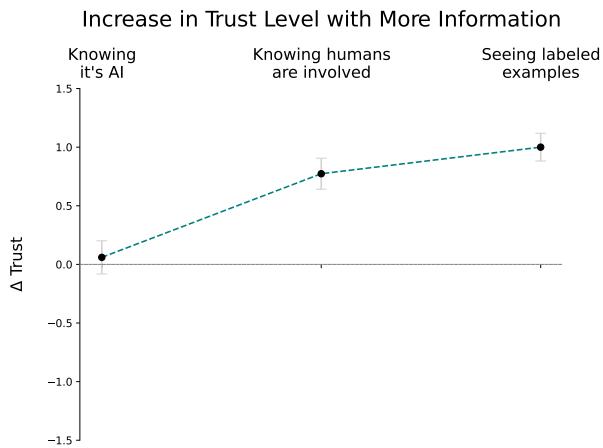
**AI Skepticism.** In the user study with experts, many participants voiced skepticism about the underlying methodology and expressed a desire for more details, along with examples of how the labels were chosen (see Section 6.3). To understand whether the general public exhibited a similar mistrust of AI, we asked them whether they trusted the MEDIA BIAS DETECTOR's labels before knowing it was powered by AI, after knowing it was powered by AI, and after being told that the labels had been validated by human annotators. Figure 9 shows how participants' responses varied. We found that most users exhibited a high degree of trust in the tool, with nearly half (76/150) stating that they trusted the labels 'mostly', and another third (52/150) stating that they trusted them 'somewhat'. Being told that these labels were generated by LLMs did not move the respondents' trust on average. Figure 9b shows the average of the responses to this question on a -2 to +2 discrete scale, while the error bars depict the associated 95% confidence interval. We did not observe any significant variation in these trends by age group, suggesting that old and young people are equally (un)skeptical about AI. However, trust in the tool rose after respondents were told that humans are involved in monitoring data samples to ensure the LLMs are performing consistently. Participants also indicated that seeing examples of specific articles rated by the models would increase their trust in the labels even further.

We also asked participants what additional information or features would increase their trust in the tool and received useful feedback. Some respondents wanted more details on the underlying methods ("An overview of how the AI system makes determinations would help me understand what's happening, which is likely to increase my trust in it.") while others wanted to dig into the data and be able to "drill down to the level of headlines or blurbs" and "take the time to look at specific examples and compare my judgements with the LLM". One of the most frequent pieces of feedback we received was a request for insight into the political leanings of the team that built the tool, with one participant asking for the tool to be "a collaboration of all diversities and political preferences" and another saying their trust would increase if "independent organizations regularly audit the tool for accuracy and allow users to compare bias findings with other tools for added validation".

**Target audience.** In the user study (see Section 6.4), we asked experts who they believed would benefit from the MEDIA BIAS DETECTOR. Most suggested that, given their own extensive experience with media bias, the target audience might be everyday news consumers who lack similar training and expertise. Here, we repeated this question to ordinary people and asked who they thought would be the right audience for this tool (Figure 10). The responses were evenly distributed across all options, with participants viewing researchers, educators, journalists, and the general public as equally



(a) Trust in the labels provided by the MEDIA BIAS DETECTOR.

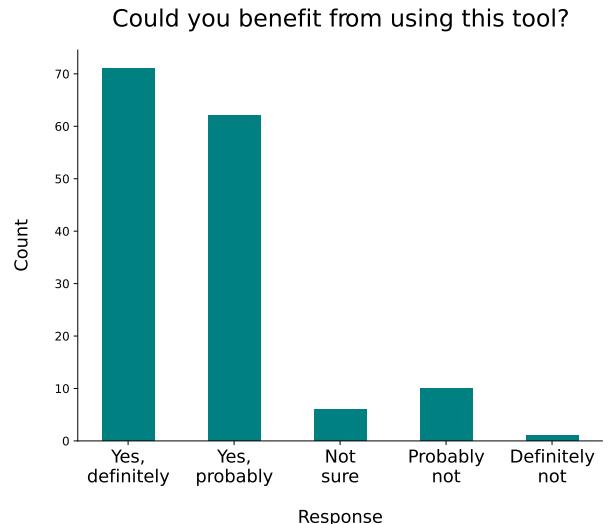


(b) Change in level of trust with more knowledge of the tool's inner workings.

**Figure 9: In contrast to our expert participants, the majority of our general public respondents expressed a significant amount of trust in the MEDIA BIAS DETECTOR's labels. More interestingly, however, this trust did not waver after learning that the tool was powered by AI, and telling people about the involvement of humans in its labeling process increased this trust, as did the possibility of seeing labeled examples of articles so they could gauge its quality for themselves. Point estimates show the average of the responses on a -2 to +2 discrete scale, while the error bars depict the associated 95% confidence interval.**

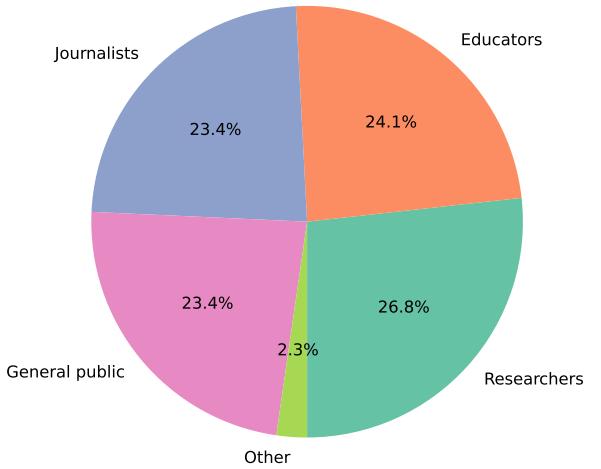
appropriate audiences for the tool. This highlights the broad appeal of the MEDIA BIAS DETECTOR and shows that people in general consider this to be a useful tool for a wide range of people.

One of the major design goals of the MEDIA BIAS DETECTOR is to *show*, not *tell* people about media bias. This is because we believe helping people recognize and understand the various ways media



(a) Personal benefit.

Which of the following audiences do you think could benefit from using this tool?



(b) Recommended audience.

**Figure 10: A vast majority of participants said that they would benefit from using this tool and considered both experts and the general public to be suitable audiences for it.**

can present a biased picture of the world is more valuable than simply assigning bias ratings to different publishers. Because of the nature of the problem this tool is addressing, we anticipated that the short-term impact of engagement with the tool would be limited. To evaluate this, we surveyed respondents about their views on selection and framing bias before and after using the tool. The results indicated that there was no significant change in users' perception of bias as a result of this brief engagement (see Appendix

A.2, Figure A1). However, the true impact of using the MEDIA BIAS DETECTOR will be apparent over time, as people continue to use it in tandem with their usual media consumption. When asked whether they thought this tool had the potential to change their views on media bias over time, the majority said that this would ‘definitely’ (50/150) or ‘possibly’ (59/150) be the case (Figure 11a). Furthermore, a fair number of participants (39/150) also believed that educational materials that explained different kinds of media bias and more information about using the tool would be helpful (Figure 11b), with one of them noting that “*understanding media bias thoroughly and knowing how to navigate the tool would enhance its usability and reliability*”.

## 8 Discussion

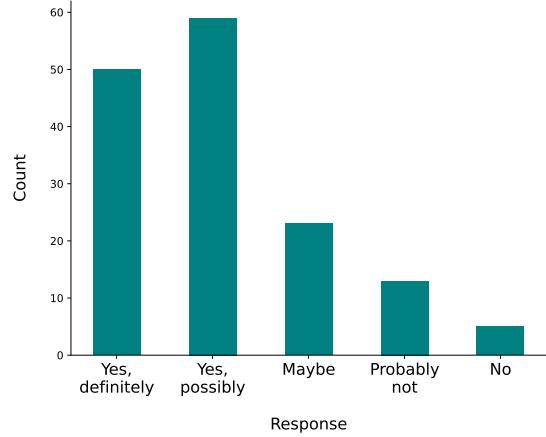
In this paper, we have introduced the MEDIA BIAS DETECTOR and described two mixed-methods studies in which we investigated users’ perceptions of the tool as well as its impact on their ability to explore and quantify media bias in news coverage. Here, we summarize the key findings for each of the research questions articulated above and characterize the MEDIA BIAS DETECTOR’s contribution to the existing state of media bias tools.

### RQ1: Conveying Media Bias Through Political Lean and Tone

The MEDIA BIAS DETECTOR aims to empower people to understand media bias themselves instead of simply telling them which publishers are biased and which aren’t—a value judgment that is subjective in and of itself. To achieve this, we designed the MEDIA BIAS DETECTOR to be easily accessible to researchers, journalists, and everyday news consumers because media bias is a universal problem that affects people from all spheres. In our evaluation with media experts, we found that participants valued the affordances provided by the tool to examine and evaluate data beyond political lean, an important design consideration highlighted in D1 (see Section 3). Participants also engaged in-depth with the topics they evaluated across lean and tone, as well as the events they viewed across publishers. Users dug deeper into the data to better understand common patterns across and within publishers, as described in D2 (see Section 3).

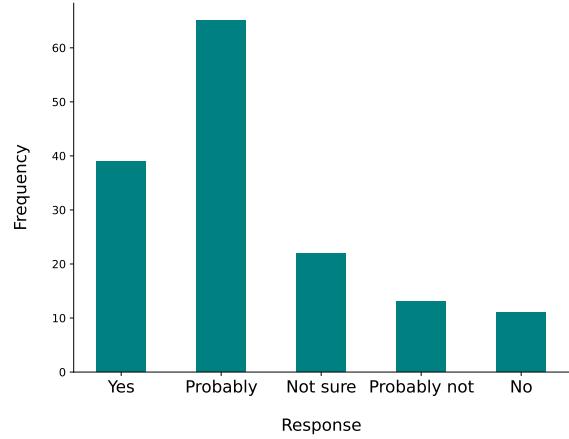
Other tools such as AllSides and Media Bias/Fact Check consider lean at the publisher level, assigning a single label to a publisher across the political spectrum. Although this approach simplifies the problem greatly, it fails to capture the heterogeneous patterns of bias that can vary within a single publisher across different topics, events, or article types. For instance, an article’s political lean or tone may shift depending on whether it is a news report, news analysis, or opinion piece, reflecting the organization’s priorities and context. It may also vary due to the evolving nature of publishers’ perspectives on emerging issues, such as elections or the COVID-19 pandemic. In contrast, the MEDIA BIAS DETECTOR assigns labels at the article level, allowing for a more nuanced understanding of these within-publisher variations. Then, the tool aggregates this data at topic, subtopic, and publisher levels to provide high-level assessments while accounting for heterogeneous effects between articles.

Do you think using this tool over time has the potential to change your view on media bias?



(a) Impact of tool use over time.

Would you be interested in additional educational materials that explain different types of media bias and how to use this tool effectively?



(b) Usefulness of education materials to aid in tool usage.

**Figure 11: A majority of participants said the MEDIA BIAS DETECTOR had the potential to change their views on media bias over time. They also stated that they would find educational materials explaining different kinds of media bias helpful for getting the most out of this tool.**

We also note that while some tools attempt to evaluate publisher level reliability, we have avoided making article level reliability judgments. Assessing reliability at the article level is particularly challenging because automating these metrics requires contextual understanding beyond the text, which remains difficult for both LLMs and humans alike. The dynamic nature of news reporting, including the ability to correct errors post-publication, further complicates reliability assessment at the article level and presents an intriguing avenue for future research.

From a user experience perspective, we chose to distill complex information into two salient outcome measures that vary significantly by article within publisher: lean and tone, where we emphasize that tone is unrepresented in existing bias tools but is increasingly considered important to audience perception [87]. The MEDIA BIAS DETECTOR focuses on these two dimensions to provide an accessible, user-friendly starting point before expanding to additional variables. Several participants noted that the tool's interactivity gave users more autonomy to form their own judgments. Participants appreciated the dual focus on lean and tone, as well as the affordances and customization options offered by the tool. However, some participants suggested that simplifying the tool could improve usability, which implied that widening the scope of the MEDIA BIAS DETECTOR to additional dimensions could introduce too much complexity in learning how to use it. These findings demonstrate that while deconstructing the multidimensional aspects of media bias is both feasible and impactful, user-centered design is crucial. The MEDIA BIAS DETECTOR contributes a design-centered approach to media bias tools, focusing on interactive evaluation across multiple perspectives, within categories, and while comparing sources to each other.

## RQ2: AI Skepticism and Human-in-the-Loop Design

Our findings reveal a divergence in trust toward AI-driven media bias tools, with experts largely expressing skepticism and everyday users showing greater initial confidence in LLM-generated outputs. This skepticism among experts highlights the need for transparency, particularly around the LLM's methodology and human-in-the-loop processes. Participants suggested adding transparency features, such as LLM explanations of article labels and examples of labeled articles, to help foster trust in the AI model's decision-making process.

Our approach includes humans in the loop to validate the model's outputs and ensure they align with human judgment. This hybrid system balances automation with accountability, maintaining consistency without placing sole reliance on automated processes. Encouragingly, both experts and everyday news consumers generally showed increased trust in the tool after learning about the human-in-the-loop process. This finding is consistent with prior research showing that trust in AI grows when people know it has higher accuracy [108]. In our case, the humans-in-the-loop act as accuracy validators by ensuring the model's outputs are reasonable and reliable.

As LLMs become increasingly popular, their role in shaping the information ecosystem grows more significant. Media bias tools like AllSides and Ground News are also incorporating AI into their features. While these websites are integrating AI to scale media bias assessments traditionally done by humans, their lack of transparency on both human and AI bias assessments leaves room for skepticism. Unlike these commercial tools, the MEDIA BIAS DETECTOR is free and open source, with transparent prompts, methodology, and weekly human validation. We implement these safeguards to minimize the risks associated with unchecked automation and ensure the LLM is performing consistently.

Beyond immediate trust, our findings highlight broader issues associated with human-AI tool development. While experts who work closely with language models and understand their faulty nature remain skeptical, regular users are quickly adjusting to the new normal with LLMs. Prior research has shown that trust in AI grows with familiarity [104]. Assuming the research community continues to validate these models and ensure they remain calibrated, the trust growing among users is encouraging. However, it also underscores the critical need to design tools that keep humans in the loop, prioritize transparency, and incorporate safeguards to address initial skepticism—preventing potential failures before it's too late.

## RQ3: Contributing a Practical Tool to the Media and Communications Research Community

Our evaluations across different user groups showed that the MEDIA BIAS DETECTOR effectively serves a diverse audience. While media experts predominantly suggested that the tool was more relevant to everyday news consumers, those news consumers felt that not only they, but also experts—journalists, educators, and researchers—could benefit from using this tool in the long run. Participants across user groups highlighted the MEDIA BIAS DETECTOR's value in revealing unexpected patterns, such as selectivity in covering major events and overall political lean of topics like the economy. The guided exploration tasks in the follow-up survey found that the tool not only confirmed prior beliefs but also introduced new insights, showing that users critically engaged with the tool.

We also discovered that news consumers saw the MEDIA BIAS DETECTOR as a tool that could shift their views on media bias over time, with many expressing interest in additional educational resources. This suggests that the tool not only revealed biases they hadn't noticed before but also sparked curiosity to learn more. In our follow-up study, a significant number of participants indicated they would benefit from extra guidance on using the tool effectively. This suggests that incorporating more resources on communication techniques like prebunking [9], accuracy nudges and fact-checking [80], and debunking [24] could supplement the MEDIA BIAS DETECTOR and further support the learning process. Although accuracy nudges and fact-checking are more pertinent to social media content and may not be directly applicable to mainstream media outlets—which typically do not contain outright factual inaccuracies—parallel strategies could be developed to prompt news consumers to recognize subtle biases, such as those involving framing and selection. The MEDIA BIAS DETECTOR is purposefully designed to be modular, prioritizing popular news publishers with the greatest reach [7]. By focusing on high-impact publishers, we can sustainably validate LLMs' outputs and implement necessary safeguards before expanding the tool further.

The true impact of the MEDIA BIAS DETECTOR will become apparent as people regularly view the data it presents and incorporate it into their daily lives. But beyond individual learning, the MEDIA BIAS DETECTOR contributes to democratic engagement by making media bias more transparent and providing people with resources to contextualize narratives and question editorial choices. Our hope is that equipping users with the tools and knowledge to recognize instances of bias will help the public benefit in the long run.

These studies have helped us more clearly recognize that the MEDIA BIAS DETECTOR's primary impact lies in improving users' understanding of media bias over time rather than simply showing them who is biased and who is not. In the wake of the 2024 election, where misinformation, one-sided narratives, and concerns about election integrity dominated public discourse, our findings show that interactive bias detection tools may play an important role in contextualizing what happens in the world during the digital age.

## 9 Limitations and Future Work

Our study helped identify areas for improvement and potential features to incorporate into future work. Below, we detail these limitations, structured across design, data, and methodological considerations.

### 9.1 Design-Related Limitations

Media experts and survey participants provided valuable feedback on the need for usability enhancements. Several participants requested a search bar to help users find specific topics and subtopics. Similarly, the Events dashboard requires users to know the exact date of an event to access related articles, which can be challenging when they only remember general time frames or keywords. Implementing a search feature for different events would improve usability by allowing users to find relevant content more reliably. Another recommended feature was the ability to click on URLs that go directly to the publisher's website to view the original articles. While this feature is currently implemented on the Events dashboard, where users can read and verify the top facts themselves, participants suggested that extending the functionality to the Coverage dashboard could increase user engagement and build trust in the platform. Another feature to build trust we plan to provide involves showing examples of articles in each tone bucket (e.g. Democrat to Republican) and lean bucket (e.g. Very Negative to Very Positive) as users focus on a specific topic or subtopic. Addressing these issues will help users access familiar news stories, making it easier to start their exploration on the MEDIA BIAS DETECTOR.

Additionally, while the Events dashboard presents a qualitative comparison of coverage differences between publishers—such as fact selection and framing—users requested aggregated insights, such as topic, tone, and lean labels to help them quickly identify these differences. In response, we plan to assign topic and subtopic labels to each event based on the majority class of their constituent articles and display these labels with aggregated tone and lean data directly on the Events page. These updates will enable quicker comparisons and a more comprehensive exploration of coverage differences.

Finally, many participants highlighted the tool's potential as an educational resource for fostering media literacy. To build on this, we aim to add a section to our website that can guide users who are less familiar with media bias. This section may include blog posts (where team members write analyses using the MEDIA BIAS DETECTOR), videos, and interactive guides explaining how to use the MEDIA BIAS DETECTOR to identify different forms of media bias (e.g., selection, framing, tone). These resources aim to empower users to better navigate not only the MEDIA BIAS DETECTOR but the information ecosystem as a whole.

### 9.2 Data-Related Limitations

The MEDIA BIAS DETECTOR currently tracks only ten publishers and collects their top 20 articles in every interval. While this covers the most prominent news published by the most widely-read publishers, it excludes coverage from local or more niche sources. Several participants highlighted that this approach misses other perspectives in the news. We built the dashboard with ten publishers, due to the significant financial and engineering costs associated with data collection and labeling pipelines. Each publisher requires a tailored data collection pipeline which involves manual text preprocessing. Additionally, using state-of-the-art, closed-source models (i.e., OpenAI's GPT-4o) on a daily basis incurs significant costs that scale linearly with the number of publishers analyzed. To select these ten publishers, we sought out a set of sources that are popular among news readers [93], while also maintaining diversity across the political spectrum (we added Breitbart for its agenda-setting power on the right, The Guardian and Huffpost on the left, and the Wall Street Journal for its financial leadership). We are currently in the process of expanding our data collection pipeline to cover the top 30 stories per news publisher and plan to expand our list of publishers beyond legacy media.

Another limitation of our tool is that the topic and subtopic lists are driven by what is actively discussed in the news. When the human-in-the-loop annotators identify major themes that have emerged in the news, we make updates to the topic list. This creates a blind spot for important issues that are not being reported by research assistants, as they may never appear in our analysis. Additionally, our current approach assigns each article to a single topic and subtopic. This rigid assignment of topics and subtopics can overlook articles that fall under multiple options. Not including them in one of those options affects what we see in the coverage numbers and may not fully reflect reality. This highlights the broader challenge of classifying a complex news article into a finite number of categories. In future iterations of the MEDIA BIAS DETECTOR, we may explore more flexible methods to accommodate overlapping themes. For example, allowing a single article to be associated with "most likely" topics or subtopics or adjusting prompts to support multiclass classification.

### 9.3 Limitations of LLMs in Media and Politics

Our tool relies on instruction-tuned, pre-trained LLMs to generalize tasks in a zero-shot setting, which introduces inherent challenges [106]. While LLMs are efficient for text classification, they are not immune to biases present in their training data, which can subtly influence downstream tasks such as political lean labeling [35]. This has important implications for the MEDIA BIAS DETECTOR's results, particularly when analyzing political news topics. When prompting the LLM for explanations on its political lean label, we can identify that they draw conclusions from not only the article text itself, but from context and inferences from its training data. This results in the LLM inherently associating certain topics with Democratic leanings (such as climate change or women's healthcare), and others with Republican leanings (such as immigration or crime), regardless of the specific arguments presented in the article. We must be cautious as the nature of these contexts and inferences may shift over time.

When measured using both novel and established methods to assess political bias, GPT model variants, including GPT-4 and GPT-3.5-Turbo, tend to align with a left-leaning political ideology [35, 68, 89]. This may relate to the model's tendency to align certain issues with the Democratic or Republican party's political agenda. While individual political lean labels could be subject to LLM biases, their relative values, and the differences between labels on articles about the same event, are informative. Our research is grounded in comparative analysis between publishers on multiple levels, and we intentionally provide granularity in all our analyses to allow users to draw their own conclusions. Furthermore, the MEDIA BIAS DETECTOR emphasizes features like tone, news topics, and event coverage, which are less susceptible to political bias in LLMs compared to political lean labels.

## 10 Conclusion

We believe that helping people identify and quantify bias across different stories and publishers, along with exposing readers to alternative narratives, can enhance public awareness of the editorial choices made by news publishers. From this perspective, the MEDIA BIAS DETECTOR could help counter the potential for creating separate realities by empowering users to independently compare how different publishers cover the same topics or events [15]. For instance, participants in our user study noted that the Events dashboard helped them see discrepancies in coverage of major events, such as the Gaza ceasefire negotiations, where some outlets prioritized domestic political stories instead. Additionally, everyday users reported that the ability to drill down into subtopics and compare political lean and tone encouraged them to challenge their preconceptions about specific publishers. By providing a high-level view of the media landscape, the MEDIA BIAS DETECTOR can help users engage with diverse perspectives and hold media outlets accountable for their power to shape public perception. Moreover, the tool may also serve as a valuable resource for researchers, educators, and journalists, enabling them to check their own biases and quantitatively assess media coverage.

We close by re-emphasizing that the goal of the MEDIA BIAS DETECTOR is to empower users to discover bias on their own, not to tell them what is biased and what is not. Our focus is on understanding how news publishers' decisions influence public perception, and how emerging narratives are shaped by the emphasis placed on them by different mainstream media outlets. This exploration is not about finding fault, but about analyzing how the natural processes of news selection and framing influence the stories we encounter in our everyday lives [99]. We therefore hope our work contributes to meaningful discussions at the intersection of media, the public, and politics, by promoting critical thinking about how news is produced, consumed, and absorbed, and ultimately the worldviews and choices it shapes.

## Acknowledgments

We thank Polygraph for their support in developing the frontend. We also thank Ajay Patel and Bryan Li for helpful discussions on the methodology. We also express our gratitude to Harsh Kumar for encouraging us to conduct user studies on the MEDIA BIAS DETECTOR and present it to an HCI audience. This research was

developed with funding from Richard Jay Mack and the Defense Advanced Research Projects Agency's (DARPA) SciFy program (Agreement No. HR00112520300). The views expressed are those of the authors and do not reflect the official policy or position of the Department of Defense or the U.S. Government.

## References

- [1] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774* (2023).
- [2] Shaaron Ainsworth. 1999. The functions of multiple representations. *Computers & education* 33, 2-3 (1999), 131–152.
- [3] Shaaron Ainsworth. 2008. The educational value of multiple-representations when learning complex scientific concepts. In *Visualization: Theory and practice in science education*. Springer, 191–208.
- [4] Scott Alexander. 2022. The Media Very Rarely Lies. *Astral Codex Ten*. <https://www.astralcodexten.com/p/the-media-very-rarely-lies> Accessed: 2024-08-17.
- [5] Hunt Allcott and Matthew Gentzkow. 2017. Social media and fake news in the 2016 election. *Journal of economic perspectives* 31, 2 (2017), 211–236.
- [6] Jennifer Allen, Baird Howland, Markus Mobius, David Rothschild, and Duncan J Watts. 2020. Evaluating the fake news problem at the scale of the information ecosystem. *Science advances* 6, 14 (2020), eaay3539.
- [7] Jennifer Allen, Duncan J Watts, and David G Rand. 2024. Quantifying the impact of misinformation and vaccine-skeptical content on Facebook. *Science* 384, 6699 (2024), eadk3451.
- [8] AllSides. 2024. <https://www.allsides.com/media-bias/media-bias-rating-methods>. Accessed: 2024-09-05.
- [9] Michelle A Amazeen, Arunima Krishna, and Rob Eschmann. 2022. Cutting the bunk: Comparing the solo and aggregate effects of prebunking and debunking COVID-19 vaccine misinformation. *Science Communication* 44, 4 (2022), 387–417.
- [10] Maryam Amirizaniani, Jihan Yao, Adrian Lavergne, Elizabeth Snell Okada, Aman Chadha, Tanya Roosta, and Chirag Shah. 2024. Developing a framework for auditing large language models using human-in-the-loop. *arXiv preprint arXiv:2402.09346* (2024).
- [11] Matthew A Baum and Tim Groeling. 2008. New media and the polarization of American political discourse. *Political Communication* 25, 4 (2008), 345–365.
- [12] Md Momen Bhuiyan, Sang Won Lee, Nitesh Goyal, and Tanushree Mitra. 2023. Newscomp: Facilitating diverse news reading through comparative annotation. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–17.
- [13] Biasly. 2024. <https://www.biasly.com/bias-meter-how-it-works/>. Accessed: 2024-09-05.
- [14] Dylan Bourgeois, Jérémie Rappaz, and Karl Aberer. 2018. Selection bias in news coverage: learning it, fighting it. In *Companion Proceedings of the The Web Conference 2018*. 535–543.
- [15] Curtis Bram. 2024. Beyond partisan filters: Can underreported news reduce issue polarization? *Plos one* 19, 2 (2024), e0297808.
- [16] Ivar Bråten and Helge I Strømsø. 2011. Measuring strategic processing when students read multiple texts. *Metacognition and Learning* 6 (2011), 111–130.
- [17] Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative research in psychology* 3, 2 (2006), 77–101.
- [18] Danielle K. Brown and Rachel R. Mourão. 2021. Protest Coverage Matters: How Media Framing and Visual Communication Affects Support for Black Civil Rights Protests. *Mass Communication and Society* 24, 4 (2021), 576–596. <https://doi.org/10.1080/15205436.2021.1884724> arXiv:<https://doi.org/10.1080/15205436.2021.1884724>
- [19] Tom B Brown. 2020. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165* (2020).
- [20] Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, et al. 2023. Sparks of artificial general intelligence: Early experiments with gpt-4. *arXiv preprint arXiv:2303.12712* (2023).
- [21] Ceren Budak, Sharad Goel, and Justin M Rao. 2016. Fair and balanced? Quantifying media bias through crowdsourced content analysis. *Public Opinion Quarterly* 80, S1 (2016), 250–271.
- [22] Heather Burgess, Kate Jongbloed, Anna Vorobyova, Sean Grieve, Sharyle Lyndon, Tim Wesseling, Kate Salters, Robert S Hogg, Surita Parashar, and Margo E Pearce. 2021. The “sticky notes” method: Adapting interpretive description methodology for team-based qualitative analysis in community-based participatory research. *Qualitative Health Research* 31, 7 (2021), 1335–1344.
- [23] John M. Carroll and Caroline Carrithers. 1984. Training wheels in a user interface. *Commun. ACM* 27, 8 (aug 1984), 800–806. <https://doi.org/10.1145/358198.358218>

[24] Man-pui Sally Chan, Christopher R Jones, Kathleen Hall Jamieson, and Dolores Albarracín. 2017. Debunking: A meta-analysis of the psychological efficacy of messages countering misinformation. *Psychological science* 28, 11 (2017), 1531–1546.

[25] Yupeng Chang, Xu Wang, Jindong Wang, Yuan Wu, Linyi Yang, Kaijie Zhu, Hao Chen, Xiaoyuan Yi, Cunxiang Wang, Yidong Wang, et al. 2024. A survey on evaluation of large language models. *ACM Transactions on Intelligent Systems and Technology* 15, 3 (2024), 1–45.

[26] Media Bias/Fact Check. 2024. <https://mediabiasfactcheck.com/left-vs-right-bias-how-we-rate-the-bias-of-media-sources/>. Accessed: 2024-09-05.

[27] Chun-Fang Chiang and Brian Knight. 2011. Media bias and influence: Evidence from newspaper endorsements. *The Review of economic studies* 78, 3 (2011), 795–820.

[28] Dennis Chong and James N Druckman. 2007. Framing theory. *Annu. Rev. Polit. Sci.* 10, 1 (2007), 103–126.

[29] Dave D'Alessio and Mike Allen. 2000. Media bias in presidential elections: A meta-analysis. *Journal of communication* 50, 4 (2000), 133–156.

[30] James W Dearing and Everett Rogers. 1996. *Agenda-setting*. Sage publications.

[31] Tilman Dingler, Benjamin Tag, Evangelos Karapanos, Koichi Kise, and Andreas Dengel. 2020. Workshop on Detection and Design for Cognitive Biases in People and Computing Systems. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI EA '20). Association for Computing Machinery, New York, NY, USA, 1–6. <https://doi.org/10.1145/3334480.3375159>

[32] Jakob-Moritz Eberl, Hajo G Boomgaarden, and Markus Wagner. 2017. One bias fits all? Three types of media bias and their effects on party preferences. *Communication Research* 44, 8 (2017), 1125–1148.

[33] Robert M Entman. 1993. Framing: Toward clarification of a fractured paradigm. *Journal of communication* 43, 4 (1993), 51–58.

[34] Motahhare Eslami, Karrie Karahalios, Christian Sandvig, Kristen Vaccaro, Aimee Rickman, Kevin Hamilton, and Alex Kirlik. 2016. First I “like” it, then I hide it: Folk Theories of Social Feeds. *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (2016). <https://api.semanticscholar.org/CorpusID:17192594>

[35] Shangbin Feng, Chan Young Park, Yuhua Liu, and Yulia Tsvetkov. 2023. From Pretraining Data to Language Models to Downstream Tasks: Tracking the Trails of Political Biases Leading to Unfair NLP Models. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki (Eds.). Association for Computational Linguistics, Toronto, Canada, 11737–11762. <https://doi.org/10.18653/v1/2023.acl-long.656>

[36] Kiran Garimella, Tim Smith, Rebecca Weiss, and Robert West. 2021. Political polarization in online news consumption. In *Proceedings of the International AAAI Conference on Web and Social Media*. Vol. 15. 152–162.

[37] Matthew Gentzkow and Jesse M Shapiro. 2006. Media bias and reputation. *Journal of political Economy* 114, 2 (2006), 280–316.

[38] Matthew Gentzkow and Jesse M Shapiro. 2010. What drives media slant? Evidence from US daily newspapers. *Econometrics* 78, 1 (2010), 35–71.

[39] Matthew Gentzkow, Jesse M Shapiro, and Daniel F Stone. 2015. Media bias in the marketplace: Theory. In *Handbook of media economics*. Vol. 1. Elsevier, 623–645.

[40] Angela A Geraci, Ardith Brunt, and Cindy Mariahart. 2014. The work behind weight-loss surgery: a qualitative analysis of food intake after the first two years post-op. *International Scholarly Research Notices* 2014, 1 (2014), 427062.

[41] Akshay Goel, Almog Gueta, Omry Gilon, Chang Liu, Sofia Erell, Lan Huong Nguyen, Xiaohong Hao, Bolouz Jaber, Shashir Reddy, Rupesh Kartha, et al. 2023. Llms accelerate annotation for medical information extraction. In *Machine Learning for Health (ML4H)*. PMLR, 82–100.

[42] Felix Hamborg, Karsten Donnay, and Bela Gipp. 2019. Automated identification of media bias in news articles: an interdisciplinary literature review. *International Journal on Digital Libraries* 20, 4 (2019), 391–415.

[43] Tony Harcup and Deirdre O’neill. 2017. What is news? News values revisited (again). *Journalism studies* 18, 12 (2017), 1470–1488.

[44] Sandra G Hart. 2006. NASA-task load index (NASA-TLX): 20 years later. In *Proceedings of the human factors and ergonomics society annual meeting*, Vol. 50. Sage publications Sage CA: Los Angeles, CA, 904–908.

[45] Jan Hartmann, Antonella De Angeli, and Alistair Sutcliffe. 2008. Framing the user experience: information biases on website quality judgement. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Florence, Italy) (CHI '08). Association for Computing Machinery, New York, NY, USA, 855–864. <https://doi.org/10.1145/1357054.1357190>

[46] Homa HosseiniMardi, Samuel Wolken, David M Rothschild, and Duncan J Watts. 2023. The diminishing state of shared reality on US television news. *arXiv preprint arXiv:2310.18863* (2023).

[47] Hong Huang, Hua Zhu, Wenshi Liu, Hua Gao, Hai Jin, and Bang Liu. 2024. Uncovering the essence of diverse media biases from the semantic embedding space. *Humanities and Social Sciences Communications* 11, 1 (2024), 1–12.

[48] Juan José Igartua and Lifen Cheng. 2009. Moderating Effect of Group Cue While Processing News on Immigration: Is the Framing Effect a Heuristic Process? *Journal of Communication* 59 (2009), 726–749. <https://api.semanticscholar.org/CorpusID:59474192>

[49] Shanta Iyengar and Kyu S Hahn. 2009. Red media, blue media: Evidence of ideological selectivity in media use. *Journal of communication* 59, 1 (2009), 19–39.

[50] Emily K. Johnson. 2022. Miro, Miro: Student perceptions of a visual discussion board. In *Proceedings of the 40th ACM International Conference on Design of Communication* (Boston, MA, USA) (SIGDOC '22). Association for Computing Machinery, New York, NY, USA, 96–101. <https://doi.org/10.1145/3513130.3558983>

[51] Oliver Kiss, Markus Möbius, and Tanya Rosenblat. 2023. Assembling News Like Legos (working paper coming soon with expanded list of authors). (2023).

[52] David M. J. Lazer, Matthew A. Baum, Yochai Benkler, Adam J. Berinsky, Kelly M. Greenhill, Filippo Menczer, Miriam J. Metzger, Brendan Nyhan, Gordon Pennycook, David Rothschild, Michael Schudson, Steven A. Sloman, Cass R. Sunstein, Emily A. Thorson, Duncan J. Watts, and Jonathan L. Zittrain. 2018. The science of fake news. *Science* 359, 6380 (2018), 1094–1096. <https://doi.org/10.1126/science.aoa2998>

[53] Bongshin Lee, Greg Smith, George G. Robertson, Mary Czerwinski, and Desney S. Tan. 2009. FacetLens: exposing trends and relationships to support sensemaking within faceted datasets. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Boston, MA, USA) (CHI '09). Association for Computing Machinery, New York, NY, USA, 1293–1302. <https://doi.org/10.1145/1518701.1518896>

[54] George Lifchits, Ashton Anderson, Daniel G Goldstein, Jake M Hofman, and Duncan J Watts. 2021. Success stories cause false beliefs about success. *Judgment and Decision Making* 16, 6 (2021), 1439–1463.

[55] Sora Lim, Adam Jatowt, Michael Färber, and Masatoshi Yoshikawa. 2020. Annotating and analyzing biased sentences in news articles using crowdsourcing. In *Proceedings of the Twelfth Language Resources and Evaluation Conference*. 1478–1484.

[56] Yixin Liu, Kejian Shi, Katherine He, Longtian Ye, Alexander Fabbri, Pengfei Liu, Dragomir Radev, and Arman Cohan. 2024. On Learning to Summarize with Large Language Models as References. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, Kevin Duh, Helena Gomez, and Steven Bethard (Eds.). Association for Computational Linguistics, Mexico City, Mexico, 8647–8664. <https://doi.org/10.18653/v1/2024.nacl-long.478>

[57] Yixin Liu, Kejian Shi, Katherine S He, Longtian Ye, Alexander R Fabbri, Pengfei Liu, Dragomir Radev, and Arman Cohan. 2023. On learning to summarize with large language models as references. *arXiv preprint arXiv:2305.14239* (2023).

[58] Jörg Matthes. 2009. What’s in a frame? A content analysis of media framing studies in the world’s leading communication journals, 1990–2005. *Journalism & mass communication quarterly* 86, 2 (2009), 349–367.

[59] Nolwenn Maudet. 2019. Dead Angles of Personalization: Integrating Curation Algorithms in the Fabric of Design. In *Proceedings of the 2019 on Designing Interactive Systems Conference* (San Diego, CA, USA) (DIS '19). Association for Computing Machinery, New York, NY, USA, 1439–1448. <https://doi.org/10.1145/3322276.3322322>

[60] Maxwell McCombs and DL Shaw. 2005. The agenda-setting function of the press. *The Press. Oxford, England: Oxford University Press Inc* (2005), 156–168.

[61] Maxwell E. McCombs and Donald L. Shaw. 1972. The Agenda-Setting Function of Mass Media. *The Public Opinion Quarterly* 36, 2 (1972), 176–187. <http://www.jstor.org/stable/2747787>

[62] Nora McDonald, Sarita Schoenebeck, and Andrea Forte. 2019. Reliability and Inter-rater Reliability in Qualitative Research: Norms and Guidelines for CSCW and HCI Practice. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW, Article 72 (nov 2019), 23 pages. <https://doi.org/10.1145/3359174>

[63] Ad Fontes Media. n.d. Interactive Media Bias Chart. Retrieved from <https://adfontesmedia.com/interactive-media-bias-chart/>.

[64] Miro. 2024. <https://miro.com/>. Accessed: 2024-09-05.

[65] A Mitchell, J Gottfried, M Barthel, N Sumida, and A Mitchell. 2018. Distinguishing between factual and opinion statements in the news. Pew Research Center’s Journalism Project.

[66] Amy Mitchell, Jeffrey Gottfried, Galen Stocking, Katerina Eva Matsa, and Elizabeth Grieco. 2017. Covering President Trump in a polarized media environment. *Pew Research Center* (2017).

[67] Fred Morstatter, Liang Wu, Uraz Yavanoğlu, Stephen R. Corman, and Huan Liu. 2018. Identifying Framing Bias in Online News. *Trans. Soc. Comput.* 1, 2, Article 5 (jun 2018), 18 pages. <https://doi.org/10.1145/3204948>

[68] Fabio Motoki, Valdemar Pinho Neto, and Victor Rodrigues. 2024. More human than human: measuring ChatGPT political bias. *Public Choice* 198, 1 (2024), 3–23.

[69] Daniel Muise, Homa HosseiniMardi, Baird Howland, Markus Möbius, David Rothschild, and Duncan J. Watts. 2022. Quantifying partisan news diets in Web and TV audiences. *Science Advances* 8, 28 (2022), eabn0083. <https://doi.org/10.1126/sciadv.abn0083>

[70] Lloyd H Nakatani and John A Rohrlich. 1983. Soft machines: A philosophy of user-computer interface design. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*. 19–23.

[71] Md. Abu Naser. 2020. Relevance and Challenges of the Agenda-Setting Theory in the Changed Media Landscape. <https://api.semanticscholar.org/CorpusID:245116933>

[72] Ground News. 2024. <https://ground.news/rating-system>. Accessed: 2024-09-05.

[73] Janni Nielsen, Torkil Clemmensen, and Carsten Yssing. 2002. Getting access to what goes on in people's heads? Reflections on the think-aloud technique. In *Proceedings of the second Nordic conference on Human-computer interaction*. 101–110.

[74] Maria Ojala, Ashlee Cunsolo, Charles A Ogunbode, and Jacqueline Middleton. 2021. Anxiety, worry, and grief in a time of environmental and climate crisis: A narrative review. *Annual review of environment and resources* 46, 1 (2021), 35–58.

[75] Deirdre O'neill and Tony Harcup. 2009. News values and selectivity. In *The handbook of journalism studies*. Routledge, 181–194.

[76] Sounil Park, Seungwoo Kang, Sangyoung Chung, and Junehwa Song. 2009. NewsCube: delivering multiple aspects of news to mitigate media bias. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 443–452.

[77] Sungkyu Park, Jaimie Yejean Park, Jeong-han Kang, and Meeyoung Cha. 2021. The presence of unexpected biases in online fact-checking. *The Harvard Kennedy School Misinformation Review* (2021).

[78] Alejandro Peña, Aiythami Morales, Julian Fierrez, Ignacio Serna, Javier Ortega-García, Ifiigo Puente, Jorge Cordova, and Gonzalo Cordova. 2023. Leveraging large language models for topic classification in the domain of public affairs. In *International Conference on Document Analysis and Recognition*. Springer, 20–33.

[79] Gordon Pennycook and David G Rand. 2021. The psychology of fake news. *Trends in cognitive sciences* 25, 5 (2021), 388–402.

[80] Gordon Pennycook and David G Rand. 2022. Nudging social media toward accuracy. *The Annals of the American Academy of Political and Social Science* 700, 1 (2022), 152–164.

[81] Emily Pronin, Thomas Gilovich, and Lee Ross. 2004. Objectivity in the eye of the beholder: divergent perceptions of bias in self versus others. *Psychological review* 111, 3 (2004), 781.

[82] Riccardo Puglisi and James M Snyder Jr. 2015. Empirical studies of media bias. In *Handbook of media economics*. Vol. 1. Elsevier, 647–667.

[83] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog* 1, 8 (2019), 9.

[84] Judith Ramey, Ted Boren, Elisabeth Cuddihy, Joe Dumas, Zhiwei Guan, Maaikie J. van den Haak, and Menno D. T. De Jong. 2006. Does think aloud work? how do we know? In *CHI '06 Extended Abstracts on Human Factors in Computing Systems* (Montréal, Québec, Canada) (*CHI EA '06*). Association for Computing Machinery, New York, NY, USA, 45–48. <https://doi.org/10.1145/1125451.1125464>

[85] Mathieu Rautav, Aixin Sun, Nancy Chen, and Shafiq Joty. 2024. On Context Utilization in Summarization with Large Language Models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Lun-Wei Ku, André Martins, and Vivek Srikumar (Eds.). Association for Computational Linguistics, Bangkok, Thailand, 2764–2781. <https://aclanthology.org/2024.acl-long.153>

[86] Mohi Reza, Nathan M Laundry, Ilya Musabirov, Peter Dushniku, Zhi Yuan "Michael" Yu, Kashish Mittal, Tovi Grossman, Michael Liut, Anastasia Kuzminykh, and Joseph Jay Williams. 2024. ABScribe: Rapid Exploration & Organization of Multiple Writing Variations in Human-AI Co-Writing Tasks using Large Language Models. In *Proceedings of the CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (*CHI '24*). Association for Computing Machinery, New York, NY, USA, Article 1042, 18 pages. <https://doi.org/10.1145/3613904.3641899>

[87] Claire E Robertson, Nicolas Pröllochs, Kaoru Schwarzenegger, Philip Pärnamets, Jay J Van Bavel, and Stefan Feuerriegel. 2023. Negativity drives online news consumption. *Nature Human Behaviour* 7, 5 (2023), 812–822.

[88] David M. Rothschild, Elliot Pickens, Gideon Heltzer, Jenny Wang, and Duncan J. Watts. 2023. Warped Front Pages. *Columbia Journalism Review* (2023).

[89] Jérôme Rutkowski, Sven Franke, Jan Endendyk, Ina Dormuth, and Markus Pauly. 2023. The self-perception and political biases of ChatGPT. *arXiv. preprint*. doi: 10 (2023).

[90] David Canfield Smith, Charles Irby, Ralph Kimball, and Eric Harslem. 1982. The Star user interface: An overview. In *Proceedings of the June 7–10, 1982, national computer conference*. 515–528.

[91] Aaron Springer and Steve Whittaker. 2019. Progressive disclosure: empirically motivated approaches to designing effective transparency. In *Proceedings of the 24th international conference on intelligent user interfaces*. 107–120.

[92] Steven A Stahl, Cynthia R Hynd, Bruce K Britton, Mary M McNish, and Dennis Bosquet. 1996. What happens when students read multiple source documents in history? *Reading Research Quarterly* 31, 4 (1996), 430–456.

[93] Statista. 2023. Leading news and media sites in the U.S. 2023, by share of visits. <https://www.statista.com/statistics/381569/leading-news-and-media-sites-usa-by-share-of-visits/> Accessed: December 5, 2024.

[94] Natalie Jomini Stroud. 2008. Media use and political predispositions: Revisiting the concept of selective exposure. *Political behavior* 30 (2008), 341–366.

[95] Pál Susánszky, Ákos Kopper, and Frank T Zsigó. 2022. Media framing of political protests—reporting bias and the discrediting of political activism. *Post-Soviet Affairs* 38, 4 (2022), 312–328.

[96] Adaku Uchendu, Jooyoung Lee, Hua Shen, Thai Le, Dongwon Lee, et al. 2023. Does human collaboration enhance the accuracy of identifying llm-generated deepfake texts? In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*. Vol. 11. 163–174.

[97] Soroush Vosoughi, Deb Roy, and Sinan Aral. 2018. The spread of true and false news online. *Science* 359, 6380 (2018), 1146–1151. <https://doi.org/10.1126/science.aap9559>

[98] Xinru Wang, Hannah Kim, Sajjadur Rahman, Kushan Mitra, and Zhengjie Miao. 2024. Human-LLM Collaborative Annotation Through Effective Verification of LLM Labels. In *Proceedings of the CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (*CHI '24*). Association for Computing Machinery, New York, NY, USA, Article 303, 21 pages. <https://doi.org/10.1145/3613904.3641960>

[99] Duncan J. Watts, David M. Rothschild, and Markus Mobius. 2021. Measuring the news and its impact on democracy. *Proceedings of the National Academy of Sciences* 118, 15 (2021), e1912443118. <https://doi.org/10.1073/pnas.1912443118>

[100] David H Weaver. 2007. Thoughts on agenda setting, framing, and priming. *Journal of communication* 57, 1 (2007), 142–147.

[101] Yi Xiang and Miklos Sarvary. 2007. News consumption and media bias. *Marketing Science* 26, 5 (2007), 611–628.

[102] Derong Xu, Wei Chen, Wenjun Peng, Chao Zhang, Tong Xu, Xiangyu Zhao, Xian Wu, Yefeng Zheng, and Enhong Chen. 2023. Large language models for generative information extraction: A survey. *arXiv preprint arXiv:2312.17617* (2023).

[103] Ka-Ping Yee, Kirsten Swearingen, Kevin Li, and Marti Hearst. 2003. Faceted metadata for image search and browsing. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Ft. Lauderdale, Florida, USA) (*CHI '03*). Association for Computing Machinery, New York, NY, USA, 401–408. <https://doi.org/10.1145/642611.642681>

[104] Kun Yu, Shlomo Berkovsky, Ronnie Taib, Dan Conway, Jianlong Zhou, and Fang Chen. 2017. User Trust Dynamics: An Investigation Driven by Differences in System Performance. In *Proceedings of the 22nd International Conference on Intelligent User Interfaces* (Limassol, Cyprus) (*IUI '17*). Association for Computing Machinery, New York, NY, USA, 307–317. <https://doi.org/10.1145/3025171.3025219>

[105] Liudmila Zavolokina, Kilian Sprenkamp, Zoya Katashinskaya, Daniel Gordon Jones, and Gerhard Schwabe. 2024. Think Fast, Think Slow, Think Critical: Designing an Automated Propaganda Detection Tool. In *Proceedings of the CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (*CHI '24*). Association for Computing Machinery, New York, NY, USA, Article 491, 24 pages. <https://doi.org/10.1145/3613904.3642805>

[106] Shengyu Zhang, Linfeng Dong, Xiaoya Li, Sen Zhang, Xiaofei Sun, Shuhe Wang, Jiwei Li, Runyi Hu, Tianwei Zhang, Fei Wu, et al. 2023. Instruction tuning for large language models: A survey. *arXiv preprint arXiv:2308.10792* (2023).

[107] Wenxuan Zhang, Yue Deng, Bing Liu, Sino Jialin Pan, and Lidong Bing. 2023. Sentiment analysis in the era of large language models: A reality check. *arXiv preprint arXiv:2305.15005* (2023).

[108] Yunfeng Zhang, Q. Vera Liao, and Rachel K. E. Bellamy. 2020. Effect of confidence and explanation on accuracy and trust calibration in AI-assisted decision making. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (Barcelona, Spain) (*FAT\* '20*). Association for Computing Machinery, New York, NY, USA, 295–305. <https://doi.org/10.1145/3351095.3372852>

[109] Caleb Ziemer, William Held, Omar Shaikh, Jiaao Chen, Zhehao Zhang, and Diyi Yang. 2024. Can large language models transform computational social science? *Computational Linguistics* 50, 1 (2024), 237–291.

## A Appendix

### A.1 User Study with Media Experts

*Task Scenario.* The same two scenarios were presented sequentially to the participants:

#### Baseline Scenario

Imagine you are having a discussion with your friends about the 2024 election. The conversation shifts to where everyone gets their news, with one friend relying solely on Newspaper A and another on Newspaper B. You are curious about how selection and framing biases present themselves in these news outlets and how they might influence coverage of key policy issues in the election.

Task: Spend up to 10 minutes attempting to **evaluate the overall bias** of Newspaper A and Newspaper B using any tools available to you. These can be existing platforms or websites that assess media bias. Note how these sources categorize the bias and any key points as to why.

Feel free to take notes in this document and **prepare a brief oral explanation** of your findings after completing this exercise.

#### Media Bias Detector Scenario

Imagine you are having a discussion with your friends about the 2024 election. The conversation shifts to where everyone gets their news, with one friend relying solely on Newspaper A and another on Newspaper B. You are curious about how selection and framing biases present themselves in these news outlets and how they might influence coverage of key policy issues in the election.

Task: Spend up to 10 minutes attempting to **evaluate the overall bias** of the same two publishers as the previous Task using the . Note how these sources categorize the bias and any key points as to why.

Feel free to take notes in this document and **prepare a brief oral explanation** of your findings after completing this exercise.

*Guiding Questions from Semi-Structured Discussion.* These guiding questions were developed to steer the discussion, though not all were necessarily used during the interviews:

#### Pre-Task

- Before we get started, can you tell me more about what got you interested in [media and journalism] OR [political science] OR [AI & society]?
- What does journalism mean to you? What is its purpose?
- In your own words, what does media bias mean to you?
- How do you recognize bias when reading news? What specific signs do you look for?
- How do you see the role of technology, such as AI and large language models, affecting the study and mitigation of media bias?

#### Post-Task

- How could this tool be used by you or others in your industry?
- What were features that you found useful and what were features you disliked?
- If you could make a magical tool that does anything you want, what features would it have? Is there anything that you really wish existing systems could do?

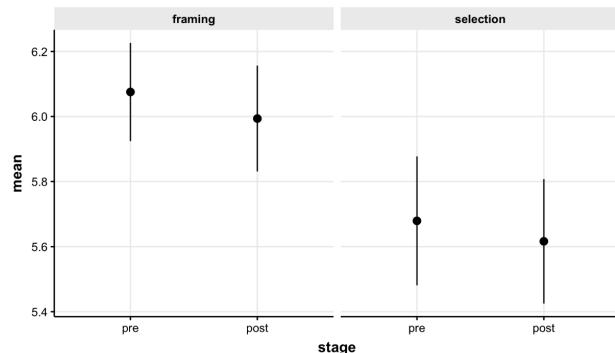
*Survey Questions.* We asked participants for demographic information and evaluations for the baseline and MEDIA BIAS DETECTOR tools through Qualtrics surveys. The surveys are viewable through Open Science Framework (OSF).

- Pre-Survey Demographic Questions
- Baseline Evaluation
- Tool Evaluation

### A.2 Follow-Up Survey of News Consumers

The follow-up survey on everyday news consumers was conducted via a Qualtrics survey. Figure A2 provides a demographic breakdown of the 150 participants. The survey is viewable through OSF:

- News Consumers Survey



**Figure A1: A comparison of pre- and post-task ratings by participants on their perception of selection and framing bias in the news.**

### A.3 Master Topic List

See Figure A3.

### A.4 Detailed Methodology

Visit the Media Bias Detector website's *Methodology section* for the most up-to-date details. The methodology is also available in PDF form in the supplementary materials.

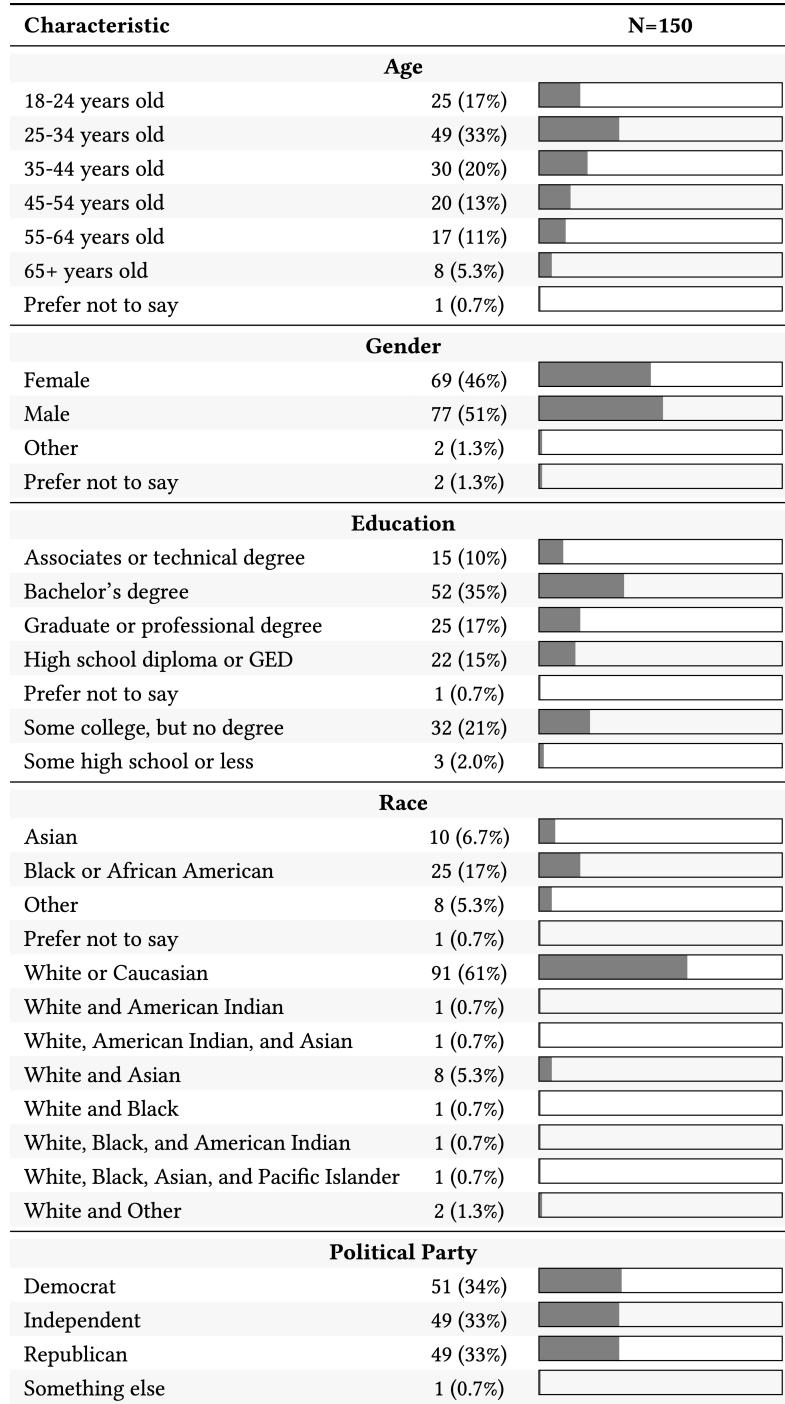


Figure A2: Characteristics of the Follow-Up Survey Population.

Category	Topic	Subtopic
Politics	Israel-Gaza War	Hostage Situation
	Democracy	Israeli Military Operations
	Crime	Voting Access/Rights
	Guns	Violent Crime
	Immigration	Gun Legislation and Regulation
	Russian-Ukraine War	Border Security
	Environment	Russian Military Operations
	Taxes	Climate Change
	Education	Tax Policy and Legislation
	Trump Legal Troubles	Diversity and Inclusion
Health	2024 Election	Jan. 6 Riot
	Foreign Policy	Supreme Court and Presidential Immunity
	Politician	US House
	Reproductive Rights	Senate Horse Race
	Healthcare	China
	LGBTQ	Middle East
	Transportation	US House Representative
	Animals	US Senator
	Arts and Entertainment	State Governor
	Celebrity	European Union
Culture and Lifestyle	Home and Lifestyle	Supreme Court and Presidential Immunity
	Travel and Tourism	Senate Horse Race
	Religion	China
	Holidays	Middle East
	Business - Corporate Strategy	US House Representative
	Business - Industry Developments	US Senator
	Health	State Governor
	Health and Wellness	European Union
	Economy	Supreme Court and Presidential Immunity
	Non-US Economy	Senate Horse Race
Sports	Sports - Baseball	Jan. 6 Riot
	Sports - Basketball	Supreme Court and Presidential Immunity
	Sports - Football	Senate Horse Race
	Sports - Other	China
	Economy	Middle East

Figure A3: Current News Category, Topic, and Subtopic Hierarchy.