Trust but (Easily) Verify: Ambivalent Bias Visualizations for News Media

SODI KROEHLER, University of Pittsburgh
AAKASH GAUTAM, University of Pittsburgh

Previous work on bias reduction in news media has focused on improving the efficacy of bias indication techniques, with the goal of reducing the biased content that news consumers read and believe. While motivated well, this effort often fails since it is primarily designed for the utility of the researcher and does not take into account the user's preferences and desires, which may at times conflict with it. This work attempts to build on the work of pluralism to introduce a new method of bias reduction that also gives utility to the consumer. Through user interviews, we iterate on a design that meets users needs while implementing features known to reduce bias, settling on a news aggregator app which we build and test.

ACM Reference Format:
Sodi Kroehler and Aakash Gautam. 2026. Trust but (Easily) Verify: Ambivalent Bias Visualizations for News Media. 1, 1 (January 2026), 16 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

1 Introduction

Motivations of this distrust frequently center around claims of factual inaccuracy or bias, although studies have shown this may be more tied to ideological similarity than objective accuracy claims [5]. While this is a subject of a great degree of study, traditional attempts to mitigate bias in news consumption have centered around designing indications of bias: such as labeling stance (e.g., 'left" or' right"), [1] [2] [17][9], hi

There has been a growing distrust of the media in the recent years, from both sides of the aisle [7]. While the causes of this are surely multivariate, the motivations of this distrust are often cited as claims of factual inaccuracy or bias, although studies have shown this may be more tied to ideological similarity than objective accuracy claims [5]. While this is a subject of a great degree of study, traditional attempts to mitigate bias in news consumption have centered around designing indications of bias: such as labeling stance (e.g., 'left" or' right"), [1, 2, 9, 17], highlighting possible biased content [14] or marking possible mis- or dis- information [4], to name a few.

However, these methods rarely evaluate their work with respect to the user's self-reported utility, and sometimes do not evaluate via a user study at all. This is a significant shortcoming, since a tool that is difficult to use for consumers is of little benefit in actually reducing bias. Previous work clearly indicates that users often prefer biased content [6, 8], so by acting to penalize biased content these methods become additionally undesirable to the user. There is thus a research gap to try to find methods which meet at least some needs or desires of the user and that additionally work to mitigate bias.

To address this gap, we conducted interviews with a variety of users to determine what they were looking for in a news app. We found that users are often very interested in multiple viewpoints, either in a drive to understand the situation more deeply or else from curiosity. Building off this and the work of [11] and others of nudges, as well as the

Authors' Contact Information: Sodi Kroehler, sek188@pitt.edu, University of Pittsburgh; Aakash Gautam, AAKASH@pitt.edu, University of Pittsburgh.

Manuscript submitted to ACM

recently developing ideas of pluralism [13], we build this motivation up into concrete design elements that meet real user needs. We connect this to the work on bias mitigation by positing the idea of "fragile cilos" - a design that presents content that is in line with the user's accepted bias - i.e. in their intellectual cilo - but then makes remarkably easy to see the similarities and differences with other viewpoints, to the point where the user may even accidentally view these other viewpoints. We hope that that this might help users realize that their existing "echo chamber" is merely one way of looking at the event and might encourage self-guided exploration of other viewpoints. In other words, this design aims to show users, rather than tell, that they may be reading biased content.

2 Context/Background Description

Our work builds off a rich body of work on reducing bias [15]. Of primary interest is the work in nudging [11, 16] which has shown users exhibit a reduced bias intake if a prompt or light bias training is presented to them beforehand. The work is highly varied, and a survey of it is much outside the scope of our work. However we do note the dominance of the metric of reducing belief in biased content, present in nearly all of the work. We build off this work in a new direction, by focusing on the users and what they are looking for in the news consumption process, and designing around that.

While this area of research is much smaller, there are a few other attempts at specifically this. NewsCube [10] is likely the closest to our method, introducing the idea of "aspect-level browsing" which we also attempt to implement. However NewsCube was built using much older technology than is now available, and relied on hyperlinks which are tedious to follow, especially on mobile devices. Munson et al [12] was much more user-focused and creative, but did not provide tools for addressing biased consumer-ship, only pointing it out. Spinde et al [14] is geared to users, but only points out biased sentences and doesn't report on the cognitive load of users. Media Bias Analyzer [17] and Biaseer [9] focus on a more general and wider-view approach, with Biaseer explicitly noting it is for social scientists, while the Analyzer attempts to be available to users as well. Both seem complicated, and in user studies the Analyzer was found to increase cognitive load when compared to just reading the articles.

There is also a considerable body of work that deals with bias/alternative viewpoint identification. This work is not outlined in this piece, as we used only fake data for testing.

3 Methodology

To begin the design process, we conducted eight semi-structured interviews with participants drawn from our immediate social circle(s). Participants were selected with a focus on keeping the overall pool balanced with respect to political ideology, gender, and other potentially influential variation, with limited success.

The primary goal of this initial set of interviews was to guide further work along lines of real, user-generated problems. Thus, we accept the limitations of the smaller interview pool, and the lack of representative guarantees that a larger interviewee pool would give us. We hope that future study can validate this tool on a more representative sample, after it is completed. Some brief demographic information on the participants is included in Table 1.

All interviews followed the same 8 question script, available in the Appendix. Each interview was held virtually on Zoom, and was transcripted with either Whisper or Zoom's native transcription function. Interviews ranged from 27 to 55 minutes long, and recordings were retained through the end of the project.

Pseudo	Gender	Age	Marital	Income	Ideology	Voted (2024)	Education
K1	Male	80–90	Married	$80–120k	Progressive	Harris	Post-grad
B	Male	20–30	Single	$0–80k	Conservative	Unspecified	Bachelor's
T1	Male	20–30	Single	$0–100k	Slightly Conservative	Did not vote	Bachelor's
T2	Male	50–60	Unspecified	$100–200k	Progressive	Harris	Bachelor's
K2	Female	40–60	Unspecified	$100–200k	Centrist	Trump	Bachelor's + License
K3	Female	40–50	Married	$80–120k	Progressive	Harris	Master's
A	Female	70–80	Unspecified	$0–100k	Unspecified	Unspecified	Bachelor's
J	Female	40–50	Married	$200–300k	Progressive	Harris	PhD
N	Female	20–30	Single	$0–80k	Progressive	Unspecified	Bachelor's

Table 1. Participant demographics, ideology, and education level (anonymized).

3.1 Analysis Process

After the interviews were completed, the transcripts were collected and parsed. All transcripts were then combined on a Miro board.[1] Codes were generated in several cycles for each interview transcript, and revisited after some days. After all articles were coded, the codes were copied to a new section, and a portion of them was used to determine overarching themes. After these were determined, the rest of the codes were assigned to each thematic element (with several additions/changes over this process). Since there only one researcher involved in this process, no evaluation is done between coders.

4 Findings

4.1 Current Practices

The interviews all began with simple questions into news habits. In these participants varied widely. A summarization of these elements can be found in Table 2. Routines for consuming news were inconsistent - T1 had set times for

[1]https://miro.com/app/board/uXjVJ26uUF4=/?share_link_id=356577745856

Pseudo	News Habits Summary
K1	Local paper; 5–7 email newsletters (1440, NiceNews, NYT); Facebook posts; Comedy (late night).
B	Social media (Instagram, Facebook); Apps (Apple News); Websites (Fox, CNN, Newsweek).
T1	Social media (X/Twitter).
T2	Podcasts (Straight Arrow News, NPR); Comedy (late night); YouTube.
K2	TV (CBS, ABC, NBC); Google; Google AI; Social media (Facebook).
K3	Radio (WGBH, WBUR, national); Email newsletters (twice weekly, local paper); Word of mouth; Apps/websites (NYT, Boston Globe, Washington Post, Atlantic); Newspaper (Bay State Banner); Comedy (late night).
A	Avoids news entirely.
J	News apps (NYT, Boston Globe, Washington Post, Atlantic); TV (local); Social media (Instagram, Bluesky, Substack); Email (WALFAM); Aggregator (Patch); Podcasts; Comedy (late night).
N	Social media (Instagram); Websites (CNN, BBC); Peers.

Table 2. News consumption patterns for each participant (anonymized).

scrolling social media, which became news-discovery times, whereas K1 read newsletters in the morning. The majority of participants did not report routines, and all acknowledged that they didn't always live by them.

Participants, across all demographics, expressed only negative emotions involved with news consumption, despite the fact that the questions were created with an effort to be emotionally neutral. Frequent themes were "outrage", "frustration", and "increased stress". "It is a lot of frustration and kind of being upset, not necessarily because of the fact that's presented, but ... how it's presented..." said N.

Most participants referenced examples from the past as superior to those of today, with a general consensus that today you need to "read more carefully (K2)". Younger participants labeled things in generic terms, like T1 saying of news today "there's no agency like stuff is just that you based on what you last saw", or N said that "there is kind of like ... pushing the needle ...what we see now as like somewhat bias used to be like incredibly biased. And I think now a lot of people do have the kind of mental filter for like taking everything with a grain of salt." Older participants cited more specific anecdotes, but along the same lines - A said news used to be "not all accurate, but was straightfoward". K1 praised hosts like Edward R. Murrow who was "50 years ago, or 70...You know, probably had his biases, but he was… He was right, a lot of the time, and he was… Even-handed with people. Okay. Differences. And a truth teller.", lamenting that he cannot find similar now.

All participants described some sort of a "checking-out" activity, to differing levels of pronouncement. At one end there was a move to avoid getting news altogether - A said she "preferred to not know what was going on unless I absolutely have to". T2 said he thinks increasingly "maybe I can just kind of sit back and kind of watch this whole thing play out", as "it feels like so many issues look like they're catastrophes on the horizon". At the other end, participants like J, K3 exhibited this trend by merely being parsimonious in choosing news sources "that don't outrage me"(K3) since so much news was "hard to stomach (J)".

4.1.1 What Caused the Negative Effects. To better understand this, we tried to analyze the interviews to understand exactly where these negative reactions were coming from.

Some participants identified disliking reporting that had an agenda. T1 said he would ignore anything "that clearly makes it undeniable that maybe they're leaning this side or that side". N more rigorously defined this as a way to affect the reader, saying she disliked when "information is either hidden or abstracted through an emotional frame or through with an angle to leave the reader to interpret in a certain way." "I don't think political parties should impact news and media" B said, extending this even to aspects of leaning that we had thought might be acceptable. K2 further extended this even to just interpretation - "I just need to know, was there a bombing in Israel ... I don't need to know why".

However, this specific sentiment was not shared among all participants - K3 specifically called out commentary like this as being useful, especially when that commentary had more expertise on the subject. (see section 4.2 for more on this).

Additionally, most participants didn't label "removing bias" as a critical next step to improve the media. In general actually there was a shared theme treating bias as inescapable - whether because it was too difficult to remove ("because as much as I want to get rid of political bias or bias in general... in the news you're not going to" (B)), or that it was intrinsic to the process of reporting ("in some ways, you'd say all news is biased because it's through the eye of the beholder"(K1)).

This also was discussed by participants occasionally without pessimism. B added "even like really highly biased news sources ...it's all just trying to get the the actual event that's there". K1 talks about, using connections to his religion, that "we're not ever going to be of one mind" but that the question is "where do we find the glue". Although this was

about personal communication, it was indicated that this was the process - again in Section 4.2 we can see that some rely on these alternate accounts to build their own confidence.

Yet other participants acknowledge that they consume biased content, yet aren't terribly worried about it. "If there's bias, if there's human judgment involved in what they're covering, I think I'm pleased to see that they are covering stories that I think are important to cover and framing them in ways that are, that are accurate" K3 says. In some ways this can represent a paradigm shift from the viewpoints of the rest of the participants, in treating the news as a slightly more serious form of entertainment rather than a civic duty. While this was a bit surprising to us at first, it is echoed in other interviews if not as clearly - for example almost all participants appreciate humor news sources (late night/the onion/ect) because it makes news "easier to digest (J)". Given our goal of making a platform that meets users needs, it seems we have dual responsibilities in both entertainment (making hard to palate things easier to process) as well as informing (talking about the hard-to-palate things) in an equal balance.

Instead of attempting to remove sources with high bias, participants often identified "triggers" or "springboards" - events which caused further search or involvement with a particular piece of content. Participants identified humor (B, J, T1), social media (J, K3, T2), opinions (K3), repeated exposure (B), or interest in the topic (B). Generally, this trigger kicked off an internet search or some other credibility obtaining method (see Section 4.2). However, this was often given when talking about humor or social media - as if this process was a vegetables that they were required to eat before consuming their "dessert": the humor content.

4.2 Confidence Building Processes

To facilitate this balancing of concerns, we can also look at the processes which happened after this trigger. Mostly, participants sought to convince themselves they knew enough of the story - reaching some internal "trust point" (J) where they felt they didn't need to look further. This wasn't the same for everyone, but as was discussed in the bias section, it was never 100% (J had 90-95, T1 had 60). The method by which they did this was - across the board - a google search. However, what they were looking for varied, as well as how they knew when they found it.

4.2.1 A Shared Base. Participants that used this method focused on consuming from a wide variety of sources but with the express goal of finding what A called the "base of whatever's being said", or T1 called "the facts and the realities behind it". T1 would read "4-5 articles" just to "skim through main points and agenda", and stopped when he felt "I get a good direction of where this is going". T2 used global news to do a similar process - "Typically, I really am ...triangulating between news out of Japan, Germany, Australia, France, and Mexico....I find that ever so much more helpful than just to rely on, say, ABC or Fox. I just think that's ... a fool's errand.".

Note, this is separate from the behavior of J, K3 as discussed later since the primary objective here is to learn the facts by sifting through multiple accounts of the event rather than any value provided by those accounts. J would say that "it's a lot of it's just coming through the AP anyway", but B would instead say you'd only get bits and pieces that people feel their viewers need to hear from these.

This of course ties into a behavior of fact-checking. K2, K1, and J didn't trust Snopes, citing issues with bias in times past or just a more general "who's watching the watchers ?(J)". Most of the other participants talked about fact-checking around specific facets of the task, such as verifying quote attribution (K1, K2).

4.2.2 Trusting a Source. Several participants gained confidence from the source and not necessarily from markers in the content. T2 said "I'm not going to be willingly tuned into it" if he didn't trust the source, T1 said he would "swipe

left, or next" if a source was clearly biased. K3 and J implied that they heavily trust their existing sources, but that they would go beyond them if they ever felt that they did not have enough information.

Most often participants that did trust from a source perspective cited things like the history of the institution or it's specific restrictions around mis-representing the truth. For example K3 trusts the New York Times because "if in 2020 I can look back at this same newspaper from 1950 and find useful information". In her job as a librarian she is often referencing the same source for current stories as well as 50 year old stories, indicating the source will likely remain true.

This same paradigm was also at times applied to specific people, whom participants trusted more. While T1 said he would not ever trust legacy media, he also said that he would be much more likely to trust someone who had experience or training in the area in which they were talking about. K1 said he would trust unlearned eyewitnesses over educated journalists who didn't directly experience the event. K3 listens to a podcast that has two anchors with a history of work with reproductive justice, and appreciates their opinions even if they are a bit biased. This same factor was cited by both her and J with respect to local news - they appreciated that they knew the anchors personally (they would meet them at the farmers market in their town).

4.2.3 Trusting many sources. In contrast to this, other participants gained confidence from precisely the opposite. While K3 could be "triggered" to check other sources from a particularly biased headline or something she didn't agree with (again using only trusted sources in this search), K2 looked for multiple sources as the primary confidence-finding metric. T1 said he consumed news on X because this was how to get news before "it gets to media", which he sees as consistently twisting content. K1 related a story of information which he found from the Heritage Foundation (a primarily conservative source which he usually finds "not helpful") which he really appreciated seeing.

4.3 Confidence Building Barriers

As this seemed like the primary pain point, we also asked about specific barriers to these practices, to better understand where our artifact could fit in.

The first and primary issue was the amount of time spent. K2 said "if you want truth, you're hours in". J lamented that she had to accept a little less confidence at times, because "it's something that could turn into another rabbit hole. You could spend hours and hours and hours trying to figure out what is the primary source".

A secondary issue was around just the difficulty of the process. "I find it that difficult because it's if I don't know anything about a topic, I it's hard because I don't know the first place to start because I read something and I already have in mind ... this is definitely not the whole story." N said. K3 echoed, saying "But if somebody tells me, oh, did you hear Memphis is now occupied by US troops? Who would I go to? What institution would I go to? And how could I find out if that is the paper of record in a city I'm unfamiliar with?" While she does value source integrity more than other participants, this worry extended to more story-based efforts as well, saying "I don't have the skills and, and the background to be able to actually fact check."

The next issue was a bit wider, and revolved around a shared theme of "journalism under attack". Some mentioned this in regards to the current administration limiting journalistic access to several political buildings, but the wider theme was just a worsening of the process. As mentioned before, this was a key aspect of understanding what caused the negative reactions in readers. Instead of lamenting this though as we covered before, participants also indicated that this was a barrier to credibility - as A and K1 had indicated, readers used to feel comfortable trusting the news given, even though none of them maintained it was always accurate. The news now, having put more of a focus on

understanding they were biased(for example Fox which is very upfront about it's leaning), seems to be indicating to users that journalism as an art form is failing. Participants also label engagement-chasing or "sensationalized (N)" as a primary driver - "it's not so much just the news now. It's more how can I make...my ratings go up (T1).

A even wider theme was around the difficulty of seeing things a different way. K1 talked about this in relation to himself - "it's a bridge too far". B and A applied it to others - saying he would trust them if they were "willing to engage in conversation", "otherwise, you just feel like somebody's been brainwashed, and they don't have any other way to say it". These feelings we also looked at as a cause of checking out, but having seen that these same participants also want to hear from these alternate sources, we can see these same things as a barrier to finding confidence as well.

4.4 What the future should look like

The final question of the interview asked participants to provide input which they thought would be important for an effective solution in this space. With the exception of N, most participants indicated they weren't looking for new sources, and J and A specifically said they would be unlikely to adopt any. Thus, most of the responses were centered around something beyond that.

Education was a big theme, across the aisle. T2 thought that a bias education class shoudl be mandatory (he suggested Bias Breakdown on spotify). K1 identified teaching fact-checking, B identified teaching how to have respectful arguments.

A slightly more hedged version of this was just in prompting reflection. N encouraged to "prompt reflection" or have a "counter frame that encourages like recognizing frames" to push against what she saw as a passive intake of bias in both herself and others. T2 said that a new product "would do well to have a huge ... psychology department budget to understand ...what mythologies we gravitate toward and, and what memes, you know, take off in our little heads.", building structural limits in the tool to prevent biased content proliferating. This was a common theme for him, as he has been worried with high engagement especially with social media.

4.4.1 Ethics. Ethics was a surprising part of a lot of peoples answers. "The person distributing the news has to have a high maxima of ethics, personal ethics, right." T1 said, explaining this would keep them from fabricating or misrepresenting since they have the power to influence mass perception. K2 didn't phrase it ethically, but still said "I think it's it may be the role of the media and the information stream to recognize some of their power or some of of their responsibility".

While supporting this idea of a normative requirement for the media to inform, K2 also revealed further issues about how that could be done. She brought up several times in the interview that she was worried about rapidly released content -a.k.a. "breaking news" - and how that might impact her and other's perceptions. Initially she liked to know things as they were happening - "I don't know the whole story yet, but it's giving me facts that I know, okay, I'm, I'm safe or there's a fire or, you know, that kind of thing.". Later in the interview, however, she said was worried that this same type of information release could make you feel fully informed when you weren't yet- "And then you're, you don't have the whole picture and you're fighting against not having the whole picture."

4.4.2 Actionable Steps. While a bit less connected than the themes we have discussed so far, we saw a theme for a more crowd-sourced version of seeing the news. While this seems to be a trend [2], it still doesn't seem like participants feel that they have a "new news" platform which meets their needs. T1, who put the least value on legacy meeting and

[2]https://www.cjr.org/the_media_today/musk_grok_media.php

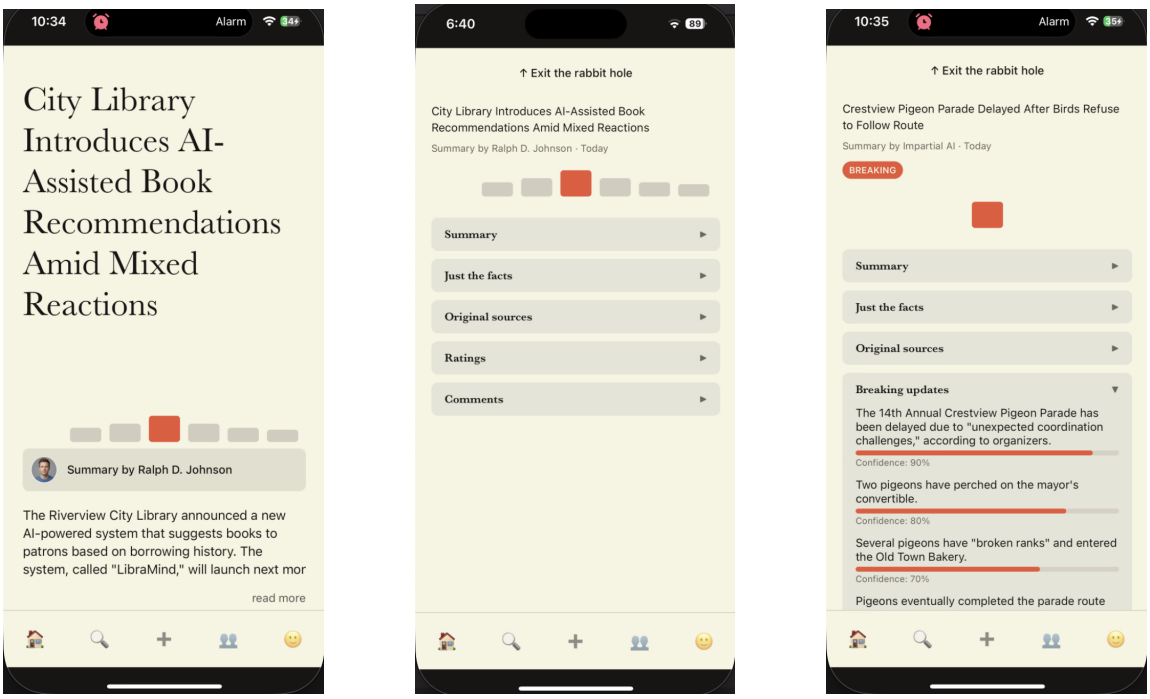

(a) Home Screen (b) Deep Dive (c) Breaking Updates

Fig. 1. Key screens from the prototype interface.

used X specifically for this purpose, but still thought it could improve - "News, or decentralized news on blockchain easily accessible to anybody's media device would be something worth a thought". He liked that it would reduce the modification between the "source" - the eyewitnesses of the event" and the reader in the same way that a distributed ledger improves fidelity in financial transactions.

Specific features were also identified. K1 liked the feature in existing news apps and LLM platforms, where one could highlight a bit of text and immediately search for it (similar to a hyperlink).

AI was also identified as an area for improvement. Most participants use it already, but in equal measure distrust it as they feel it is too black-boxy. "How can you trust the AI would be ...unbiased? ...how is AI programmed, who's doing it, and what is going in, in order to get what comes out?" A said, but most participants echoed similar concerns. K3 thought that a feature of the tool that showed the path of the story through the news : "When was this topic first covered and by whom? Like a trail." would help with this.

5 Design Implications

Due to these requirements, an news-aggregator app seemed the logical solution. J and A, along with others more indirectly, indicated that they weren't looking for new news sources, but were excited about a new news browsing app. We went through first and tried to summarize the needs that the users had and then brainstorm features which would help solve them. This process went through several wireframes and a few Figma prototypes, and did not directly

result in a feature for all the needs, but only those which were identified as being most critical from the user interviews. These needs and the associated features are summarized below.

5.1 Features

5.1.1 Multi-Directional Navigation. Despite unilateral negativity about the biasedness of sources, all users made use of a number of sources. K3 and those who used the trust they had in the source to reach their "trust point" still made use of a number of similar sources, and appreciated content which was more biased than these sources as additional information. Others, like B, K2 and those who used a shared base to reach their trust point used sources across the aisle. In both cases, the time to find these sources could be reduced, thus addressing the primary issue user's faced in the time spent. To address this, we implement an app-based design similar to existing shortform content apps (such as Tiktok) that rely on swiping as the primary interaction. Similar to these apps, a swipe navigates discretely instead of via a continuous scroll. We positioned alternate viewpoints on the same issue on a horizontal axis, and different stories on a vertical axis. The nature of the horizontal axis was also a subject of some conversation. We note that while most users talked about political leaning as a primary difference between sources, others identified other aspects. T1 was interested in what was presented domestically as opposed to internationally, K3 talked about the viewpoints of podcasters that work in reproductive health versus those that work in law on reproductive health. Thus, we thought the design would be better if it was open to more than just the traditional left-right dichotomy. However, during user testing we used this still, arranging the multiple versions from most left-leaning (on the furthest left) to the most right-leaning (on the furthest right). An example of the screen on the most center-leaning version can be seen in Figure 1a.

5.1.2 Summaries. Each viewpoint is represented at this level - the "top level" - by a headline and a short 1-2 sentence summary. This is made to address the difficult that many users face with having to read through highly biased content - this feature lets them immediately tell whether it is a bias that they want to engage with or do not, with limited negative emotional triggers. A larger summary as well as direct links to the original source articles are available in the "Deep Dive" section.

Although the idea would be to source real news sources and present them, for the purposes of testing we simply created a sample set of news sources for a fictional set of stories. We created 6 different "versions" of the first story, to represent the different types of articles a user might find when googling. Each version consisted of a summary and several "sources" which represented that viewpoint.

While all the data for the demos were AI-generated, we included an author bar on the main screen and showed some written by users and others by AI. This was designed for users who were skeptical of AI, allowing them to see from the home screen what was written by AI and which had human authorship.

5.1.3 "Deep Dive". For each viewpoint, we created an extra details screen which was navigable by either a triple tap interaction or else a visual cue on the screen. Figure 1b shows a sample of the "Deep Dive" level. This screen is meant to be flexible with the sections included, but to address the different user's needs. For example, shared-base users can much more easily see the facts that support a specific viewpoint from this screen, and source-trust users can easily navigate to the original sources they deem trustworthy.

We experimented with several possible visual cues to trigger this. We tried a "shovel" icon on the author bar, but during heuristic testing we realized it violated Nieson's 6th heuristic: maintaining "Recognition rather than recall", and would be too confusing for users. We also attempted a large footer bar but it didn't look aesthetically pleasing. We

(a) Original (b) Labeled (c) Final

Fig. 2. Key screens from the prototype interface.

finally settled on the "read more" button, which matches existing newspapers to clearly demarcate the functionality, abiding by the "consistency and standards" heuristic.

5.1.4 "Just the facts". We experimented with several orders of the available subsections, but focused on just a few that were clearly motivated by the user interviews. Chief among these is the "Just the Facts" section. The information in this section (at least for our demo) contains the same information as in the summaries, but in clear, one-sentence statements. This was built for users like B, K2 who valued "talking heads" - or just the facts without additional commentary. Additionally, as can be seen in 1c, we included a "confidence meter" for facts that show contention, in hopes to fulfill the hopes of N, T1 in prompting reflection. Users can also swipe between viewpoints in this "deep dive" level, which was designed for participants like A, T1 who valued a shared base as a confidence building mechanism - it is easy for them to be able to see what facts are shared across viewpoints (and thus have higher confidence) and which are only in certain accounts of the story.

5.1.5 Breadcrumbs. During user testing, we found that the horizontal swipe was not intuitive for many users, so we experimented with several visual aids to emphasize the functionality. The first attempt was a simple "breadcrumb" made of light grey boxes above the author bar, shown in 2a, with an orange indicator for the current page. Users found it useful in determining the slant of the current viewpoint, but did not infer to swipe horizontally from it, unless told explicitly too. We improved on this design by adding arrows and specific axis labels, shown in 2b. Users inferred to swipe here, but the design looked messy. Additionally, several users were worried about presenting all viewpoints equally, noting that since we are not specifically avoiding bias some viewpoints at the end of a spectrum might be extreme, and users would not be benefited if it was presented as equivalent. Thus, we iterated again to final version, shown in Figure 2c, and in the sample screens. We added a height to each bar which indicates the amount of sources supporting the summary of that viewpoint, highlighting that viewpoints further from the center have less people writing about and consuming them, and thus are should be taken with a grain of salt. This has the added benefit of showing missing or unbalanced viewpoints for a particular story as bars with zero height, and retains the intial value as a breadcrumb to indicate which viewpoint is currently being shown. To lessen the time it took the users to figure out how to swipe horizontally, we also made each "pillar" clickable. We found in user testing that the users understood the meaning of the breadcrumbs, but often clicked them instead of swiping, so this feature is likely more natural.

5.2 Feature Evaluation

To verify this creation, we did several user tests. We recruited users by tabling in a thoroughfare of the business school as well in the graduate computer science area. We successfully completed 5 user interviews, all of computer science PhD students. While this is not representative of our final user population, it was the best we could accomplish within the timeframe.

Users were first given the app to explore, and then asked a few general questions about it. All users successfully identified that it was a news app, and were all quickly (<10 seconds) able to discover scrolling vertically went to a new

story - often accompanied by a "oh it's like tiktok" or something like that. Users were a bit slower (> 30sec) to learn they could swipe horizontally, usually they attempted to click the breadcrumb pillars - a key reason why we added that feature. Users were then given four tasks: 1. Navigate to a new story, a determine if it is currently "breaking" or if the story happened a while ago. 2. What's one of the biggest controversial aspects of the first story. 3. Compare the disagreement of the first story to the second one. and 4. Find out the facts of the "breaking story". Users were able to accomplish all tasks accurately in less than a minute.

6 Discussion

This work studies news consumer's attitudes about bias-prevention techniques. Instead of creating a bias measurement or indication tool, like previous work has done, we instead focus on user's desires while engaging with news content. We find that even across numerous demographic divides (age, gender, political leaning, news engagement) those we interviewed rarely used fact checkers or bias indicators, consumed media which they understood to be biased regularly, and highlighted the difficulties of reaching a "trust point" - or a level of understanding where they felt they understood the situation - far more than they did avoiding bias entirely. For the majority of them, this involved looking at multiple sources, whether from different outlets, ideological leanings, or even countries. In response to this, we developed a demo news aggregation app that aims to meet these needs firstly, and reduce bias secondly.

Specifically, we create an app which utilizes horizontal and vertical swiping to navigate between a set of summaries. Each summary is autogenerated from a cluster of news sources, which can be subsections of or even entire unique articles, tweets, or any other form of text. For a specific story, the available information is clustered into it's most dominant *ambivalence* - or most divisive elements - although we do not specify the method of this in this work.[3] We then arrange these along the horizontal axis, and make different stories scrollable via the vertical axis. This presents the available information to the user in a way that encourages exploration even of biased content without falsely indicating its trustworthiness.

Our work is less of an improvement on prior work in this domain and more of an offshoot. Because we focus on the user's utilities, which often overlap with biased news content, our task stops being how do we prevent people from consuming biased content and shifts to how do we allow people to consume the biased content that they want to in a healthy way, and is there an overlap between facilitating this and meeting other user's needs. To illustrate this and how our work relates to others, we engage with three dominant offerings in the academic and commerical world today: Spinde et all, GroundNews and Allsides, and the Media Bias Analyzer by UPenn.

Spinde et al created a classifier and presentation method which highlights possibly biased sentences in news. They report statistically significant increases in bias recognition performance when using this tool. This is in line with what several users suggested around educating readers or prompting reflection while reading. This may seem attractive to us as researchers, who see the problems with biased content and wish to mitigate them with our work. However, user's do not share this and may see these labels as intrusive or forced upon them, as they do with other similar controls [12]. Also this does not address the primary barriers to reaching that "trust point", which users identified as a central problem. Thus, our work attempted to improve on this work by extending similar methods (the confidence score on each fact can be seen as an extension of this) while still focusing our main efforts towards creating something which users may not have already envisioned but which would help address their concerns more comprehensively.

[3]Again we only use one ambivalence in this work - specifically the political left versus political right. This is why we avoid putting familiar colors like red and blue on the breadcrumbs - we want to leave the horizontal axis customizable for stories that vary in other respects, such as bull or bear markets. However, this is beyond the scope of this article, and left-vs-right are more easily understood in this context.

A much more comprehensive solution is the dashboard created by the Computational Social Science Lab at Penn [17]. This tool examines bias in news articles both representationally and semantically, and makes the results available to researchers and users at a per-event level. While a truly remarkable and quality creation, the tool has a lot of powerful features which may seem daunting to the casual user, and in fact user studies revealed that the interface was a bit tricky for them and even for some experts. Our app is much less powerful, but still communicates the differences between sources, while retaining a very uncluttered design and simple navigation.

Several other commercial solutions in this problem space have seen market success. Allsides [1] regularly presents stories in the "multiple-summaries" format, and is widely cited. GroundNews [2] has seen rapid growth since it's creation in 2020, and is the most similar to our work. However, both of these solutions are focused primarily on the ideological spectrum, and present bias indicators or alternate summaries only along this tradeoff. We improved on this solution by making our artifact more suited for mobile-device browsing, using swipes to navigate instead of the more lengthy hyperlinks. Additionally our solution is an improvement on these since it is flexible to multiple different *ambivalences*, whether that difference is political, economic, or some other dimension.

From these issues, we position our tool as an improvement on the user experience rather than on the user's agreement with the bias. As other studies have shown [3, 14], showing multiple versions of a story does not have much of an effect on bias recognition, but choosing bias recognition as the primary metric does not seem to be based on user demand, perhaps a motivating factor for why these tools are difficult or even undesirable for users. This is likely because the desire to reduce bias comes from academic or altrustic concerns, seeking to bridge increasing political polarization. However, this will certainly fail if it does not take into account the preferences and desires of the users.

In a way, our work can be positioned with the work on echo chambers/intellectual cilos by piloting the idea of a "fragile cilo" - something that gives the single viewpoint, highly biased outlook which users are looking for, but makes it unbearably easy to be exposed to a different viewpoint. Our tool capitalises on the desire that many of these users have to explore other viewpoints to incentivize bias recognition and improve more informed understanding. By giving users the freedom and power to explore other versions to the extent that they desire, we do not contribute to uneasiness or frustration which is already present just from the content of the news. Our goal is to leverage the existing curiosity that many users display to the viewpoints on "the other side of the aisle" to help engender hesitancy and hedging, if possible, but at least less time searching for and navigating to other viewpoints and more time becoming more informed.

7 Limitations

This study was performed over the course of a semester, with an extremely small team. As such, the interview coding and user testing is far below what it should be. Additionally, interviewees and users were selected via convenience sampling, by selecting from family and friends or else co-located students to the author(s). Given that this tool is being developed for a wide and varied audience, wider user testing should evaluate whether this tool is desired, effectively reduces the cognitive load where it intends to, and is a marked improvement upon existing tools. We intend to pursue these questions in future work, and hopefully build on the work presented here.

8 Future Work

We hope to extend this tool to "live" data, and make it publicly accessible. Additionally, we would like to implement more of the user features which were identified in the interviews, particularly around humor. Our app has multiple ways to easily include memes or relevant humor videos in each viewpoint, which can help users engage with the

content instead of immediately dismissing it. Finally, a great amount more user testing and iteration is necessary, to be sure this is not only meeting user's needs but also functioning to reduce instead of accentuate bias.

9 Conclusion

In this paper, we have introduced several problems with the existing work on news media bias identification/mitigation efforts: in particular they focus on the academic work of reducing bias while ignoring the news media consumers desires. This results in tools that are difficult for users to use, or else that users do not want to use over biased alternatives. The resulting tool bridges these gaps by presenting the biased work alongside other viewpoints, allowing users to consume the news they are interested in, while greatly reducing the chance that they do not see another viewpoint. The goal of this work is thus not to reduce bias understanding, but to reduce the work that a user has to go through in order to be convinced they fully understand a situation, which we hope will eventually cause users to be less biased in favor of being more informed.

10 Contribution Statement

This is a solo project. Chatgpt was used to format latex for the tables and figures, and to translate the interview script (written in word) to latex.

A User Interview Script

Introduction (2–3 minutes)

- **Recording.** Ask participant for permission to record the session.
- **Study purpose.** You are invited to participate in a research study about how people read, interpret, and evaluate news articles.
- **What you will do.** You will answer a series of questions about how you consume news and any challenges you experience in evaluating it.
- **Risks and benefits.** This study involves minimal risk. Some questions may reference current events. There are no direct benefits to participation, though study results may be shared upon request.
- **Privacy.** No IP addresses are stored. All responses will be anonymized and reported in aggregate.
- **Voluntary participation.** Participation is voluntary. You may skip any question or stop the interview at any time without penalty.
- **Compensation.** [If applicable, specify amount and conditions.]
- **Contacts.** For questions, contact [PI email]. For concerns about participant rights, contact [IRB information, if applicable].
- **Consent.** By agreeing to participate, you confirm that you are at least 18 years old and consent to take part in this study.

Kickoff Questions (5 minutes)

(1) How do you typically stay informed about current events?

Follow-up prompts:

- How did you originally discover these sources?
- Have your go-to sources changed over the past few years?

- Do you use different sources for different types of news?
- Do you have a routine for consuming news, or does it vary day to day?

Building Rapport (8–10 minutes)

(2) When you hear people talk about *biased news*, what does that mean to you?

Follow-up prompts:

- How substantial of an issue do you think bias in the news is today?
- Do you think it is getting worse or better? Why?
- Can you describe a recent example?
- How does bias differ from misinformation or disinformation? Which do you think is more harmful?

Grand Tour Questions (20 minutes)

(5) If you have ever encountered news content that goes against your beliefs, how did it make you feel?

Follow-up prompts:

- Does your reaction differ when you agree with the content versus when you disagree?
- Does it differ if the content involves framing or spin, a false claim, or another form of bias?
- Does biased content make you more or less likely to engage with the information?

(7) Tell me about any fact-checking you do when reading news.

Follow-up prompts:

- What typically triggers you to fact-check something?
- What tools or methods do you use (e.g., multiple sources, Snopes)?
- Have these tools or methods changed over time? What caused those changes?
- What indicators make you trust a fact-checking tool or source?
- What makes a news source trustworthy to you?

(8) Tell me about a time when your understanding or beliefs about an event were challenged.

Follow-up prompts:

- Can you walk me through how you made sense of conflicting reports?
- How did you decide which sources to trust more?
- How did you feel when you encountered a new version of a story you already knew about?

(9) What causes you to engage with differing viewpoints, and what causes you to disengage?

Follow-up prompts:

- Does the source of the information matter?
- Do you continue engaging with content you perceive as biased? What kinds?
- Does this differ when talking with friends versus reading on your own?
- Does this change across communication modes (phone, text, in person)?

Reflection (5 minutes)

(10) This study explores potential interventions to help people navigate bias in news. What do you think is important to consider when designing such tools?

Follow-up prompts:

- Do you feel that news consumers currently have enough tools to identify or counteract bias?
- What motivates people to check or reduce bias? What discourages them?
- What kinds of interventions do you think might backfire?

Wrap-Up (5 minutes)

(11) Is there anything about how you consume news or think about bias that we have not covered today?

Thank the participant and briefly explain next steps in the research process.

Demographics

Participants were asked (optionally) to self-report the following:

- Gender
- Voting participation in the most recent election
- Highest level of education completed
- Political orientation (e.g., whether government should do more or less to solve problems)
- Household income range
- Any other information the participant felt comfortable sharing and believed might be relevant

References

[1] [n.d.]. *AllSides | Balanced news via media bias ratings for an unbiased news perspective.* https://www.allsides.com/unbiased-balanced-news

[2] [n.d.]. *Ground News.* https://ground.news/

[3] Jisun An, Meeyoung Cha, Krishna Gummadi, Jon Crowcroft, and Daniele Quercia. 2012. Visualizing media bias through Twitter. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 6. 2–5.

[4] Pepa Atanasova, Alberto Barron-Cedeno, Tamer Elsayed, Reem Suwaileh, Wajdi Zaghouani, Spas Kyuchukov, Giovanni Da San Martino, and Preslav Nakov. [n.d.]. Overview of the CLEF-2018 CheckThat! Lab on Automatic Identification and Verification of Political Claims. Task 1: Check-Worthiness. https://doi.org/10.48550/arXiv.1808.05542 arXiv:1808.05542 [cs]

[5] Gallup and Knight Foundation. [n.d.]. PERCEIVED ACCURACY AND BIAS IN THE NEWS MEDIA. https://knightfoundation.org/wp-content/uploads/2020/03/KnightFoundation_AccuracyandBias_Report_FINAL.pdf

[6] R Kelly Garrett. 2009. Politically motivated reinforcement seeking: Reframing the selective exposure debate. *Journal of communication* 59, 4 (2009), 676–699.

[7] Jesse Holcomb. [n.d.]. *Media Mistrust Has Been Growing for Decades—Does It Matter?* https://pew.org/3Nl4wYc

[8] Shanto 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.

[9] Guozheng Li, Shiyu Han, Jihe Wu, Jiale Hu, Yu Zhang, and Chi Harold Liu. [n.d.]. BiaSeer: A Visual Analytics System for Identifying and Understanding Media Bias. 9, 2 ([n. d.]), 1–28. https://doi.org/10.1145/3710936

[10] Souneil Park, Seungwoo Kang, Sangyoung Chung, and Junehwa Song. [n.d.]. NewsCube: delivering multiple aspects of news to mitigate media bias. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Boston MA USA, 2009-04-04). ACM, 443–452. https://doi.org/10.1145/1518701.1518772

[11] Gordon Pennycook, Jonathon McPhetres, Yunhao Zhang, Jackson G Lu, and David G Rand. 2020. Fighting COVID-19 misinformation on social media: Experimental evidence for a scalable accuracy-nudge intervention. *Psychological science* 31, 7 (2020), 770–780.

[12] Sean A. Munson, Sean A. Munson, Stephanie Y. Lee, Stephanie Y. Lee, Paul Resnick, and Paul Resnick. [n.d.]. Encouraging Reading of Diverse Political Viewpoints with a Browser Widget. ([n. d.]). MAG ID: 2182265315.

[13] Taylor Sorensen, Jared Moore, Jillian Fisher, Mitchell Gordon, Niloofar Mireshghallah, Christopher Michael Rytting, Andre Ye, Liwei Jiang, Ximing Lu, Nouha Dziri, et al. 2024. A roadmap to pluralistic alignment. *arXiv preprint arXiv:2402.05070* (2024).

[14] Timo Spinde, Felix Hamborg, Karsten Donnay, Angelica Becerra, and Bela Gipp. [n.d.]. Enabling News Consumers to View and Understand Biased News Coverage: A Study on the Perception and Visualization of Media Bias. In *Proceedings of the ACM/IEEE Joint Conference on Digital Libraries in 2020* (Virtual Event China, 2020-08). ACM, 389–392. https://doi.org/10.1145/3383583.3398619

[15] Timo Spinde, Christin Jeggle, Magdalena Haupt, Wolfgang Gaissmaier, and Helge Giese. 2022. How do we raise media bias awareness effectively? Effects of visualizations to communicate bias. *PLoS One* 17, 4 (2022), e0266204.

[16] Calum Thornhill, Quentin Meeus, Jeroen Peperkamp, and Bettina Berendt. 2019. A digital nudge to counter confirmation bias. *Frontiers in big data* 2 (2019), 11.

[17] Jenny S Wang, Samar Haider, Amir Tohidi, Anushkaa Gupta, Yuxuan Zhang, Chris Callison-Burch, David Rothschild, and Duncan J Watts. [n.d.]. Media Bias Detector: Designing and Implementing a Tool for Real-Time Selection and Framing Bias Analysis in News Coverage. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems* (Yokohama Japan, 2025-04-26). ACM, 1–27. https://doi.org/10.1145/3706598.3713716