

Document Version

Final published version

Licence

CC BY

Citation (APA)

Wang, J., He, G., Rieger, A., & Gadiraju, U. (2026). From SERPs to Sound: How Search Engine Result Pages and AI-generated Podcasts Interact to Influence User Attitudes on Controversial Topics. In C. Shah, R. W. White, A. Fourney, C. T. Lopes, & J. Trippas (Eds.), *CHIIR 2026 - Proceedings of the 2026 Conference on Human Information Interaction and Retrieval* (pp. 51-61). Association for Computing Machinery (ACM). <https://doi.org/10.1145/3786304.3787942>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

In case the licence states "Dutch Copyright Act (Article 25fa)", this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership.
Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

From SERPs to Sound: How Search Engine Result Pages and AI-generated Podcasts Interact to Influence User Attitudes on Controversial Topics

Junjie Wang
Delft University of Technology
Delft, Netherlands
j.wang-101@student.tudelft.nl

Alisa Rieger
GESIS - Leibniz Institute for the Social Sciences
Cologne, Netherlands
alisa.rieger@gesis.org

Gaole He
Delft University of Technology
Delft, Netherlands
g.he@tudelft.nl

Ujwal Gadiraju*
Delft University of Technology
Delft, Netherlands
u.k.gadiraju@tudelft.nl

Abstract

Compared to search engine result pages (SERPs), AI-generated podcasts represent a relatively new and relatively more passive modality of information consumption, delivering narratives in a naturally engaging format. As these two media increasingly converge in everyday information-seeking behavior, it is essential to explore how their interaction influences user attitudes, particularly in contexts involving controversial, value-laden, and often debated topics. Addressing this need, we aim to understand how information mediums of present-day SERPs and AI-generated podcasts interact to shape the opinions of users. To this end, through a controlled user study ($N = 483$), we investigated user attitudinal effects of consuming information via SERPs and AI-generated podcasts, focusing on how the sequence and modality of exposure shape user opinions. A majority of users in our study corresponded to attitude change outcomes, and we found an effect of sequence on attitude change. Our results further revealed a role of viewpoint bias and the degree of topic controversiality in shaping attitude change, although we found no effect of individual moderators.

CCS Concepts

• **Human-centered computing** → **Human computer interaction (HCI)**; • **Information systems**; • **Computing methodologies** → **Artificial intelligence**;

Keywords

Attitude Change, AI-generated Podcasts, Information modality, Web search, Controversial Topics, Responsible Opinion Formation

ACM Reference Format:

Junjie Wang, Gaole He, Alisa Rieger, and Ujwal Gadiraju. 2026. From SERPs to Sound: How Search Engine Result Pages and AI-generated Podcasts Interact to Influence User Attitudes on Controversial Topics. In *2026 ACM*

*Corresponding author.



This work is licensed under a Creative Commons Attribution 4.0 International License. *CHIIR '26, Seattle, WA, USA*

© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2414-5/26/03
<https://doi.org/10.1145/3786304.3787942>

SIGIR Conference on Human Information Interaction and Retrieval (CHIIR '26), March 22–26, 2026, Seattle, WA, USA. ACM, New York, NY, USA, 11 pages.
<https://doi.org/10.1145/3786304.3787942>

1 Introduction

Traditional search engine result pages (SERPs) have served as the primary gateway to online information for years, offering users a curated list of sources that they can actively navigate, shaping user opinions and influencing their attitudes on a variety of topics [21, 22, 80, 84]. For example, there has been a wealth of evidence demonstrating how rankings of search results and their corresponding viewpoints influence consumer choices—called the *search engine manipulation effect* (SEME)—since users tend to consume, trust, and rely more on higher-ranked web search results than lower-ranked results [18, 21]. Over the last decade, we have witnessed a growing interest in the conversational search paradigm [57] and a flourishing presence of digital assistants, chatbots [25], and conversational agents [13, 38] that aid users in satisfying their information needs [3, 4, 37, 43, 73]. At the same time, there has been a growth in podcast consumption as a source of information [70] with millions of episodes readily available on several topics, and the number of podcast listeners around the world has nearly doubled since 2019 [52, 67, 68]. Notably, podcasts uniquely deliver linear, narrative content that users typically consume with little interaction or selective control.

With advances in generative AI and the frantic adoption of AI-powered solutions in search systems over the last few years, there has been a shift in information seeking and consumption behaviors [81]. For example, users today are widely exposed to AI-generated content in different media, such as textual *AI overviews* on Google search [86] or AI-generated podcasts in the audio medium using technologies like NotebookLM.¹ Since digital information is now omnipresent and increasingly personalized, understanding how users form and update their opinions based on different media formats is critical. In recent work, White and Shah [82] argued that in the future, *panmodality* (i.e., the use of multiple modalities either separately or collectively) will play a central role in information interaction. Mayerhofer et al. [48] explored how users engaged with an interface combining web search and a generative AI chat

¹<https://notebooklm.google.com>

feature to solve health-related information tasks. Their qualitative findings suggest that while such integration with generative AI has the potential to enhance information-seeking, it can also lead to misplaced trust in favour of ‘ease-of-use’ and seemingly perfect answers. This corroborates findings across the HCI and NLP research communities that has explored how users tend to build trust and reliance on agents powered by LLMs that demonstrably provide convincingly wrong answers or advice [7, 31, 32, 64].

In contrast to traditional SERPs, AI-generated podcasts represent a relatively new and more passive modality of information consumption, delivering synthesized narratives in a conversational format. Compared to other passive formats (e.g., static summaries, videos, and to an extent, long-form articles), podcasts offer a conversational, human-like delivery that may foster perceptions of coherence, trust, and social presence. Existing work in HCI, media and communication suggests that such qualities can meaningfully influence persuasion and attitude formation [26, 41, 47, 49]. Thus, podcasts provide a particularly rich test case for understanding how passive, narrative-driven modalities interact with active search behaviors. Recent preliminary work compared the impact of SERPs and AI-generated podcasts on users attitude change and found no significant differences [76]. However, as these two media converge in everyday information-seeking behavior, it becomes essential to explore how their interaction influences user attitudes, particularly in contexts involving controversial, value-laden, and debated topics [59, 60, 75]. This informs the core research gap and the overarching question that we aim to address in our work—*how do the information mediums of present-day search engine result pages (SERPs) and AI-generated podcasts interact to shape opinions of users?* To this end, we aim to investigate the user attitudinal effects of consuming information via SERPs and AI-generated podcasts, focusing on how the sequence and modality of exposure shape user opinions. Thus, we aim to answer the following research questions:

- **RQ1:** How does information medium sequence (i.e., *SERP-first* versus *podcast-first*) influence user attitude change?
- **RQ2:** How does the effect of information medium sequence on user attitude change differ across viewpoint biases (supporting, opposing, neutral)?
- **RQ3:** How does topic controversiality influence user attitude change across segments within information medium sequences?
- **RQ4:** How do individual differences in AI literacy, intellectual humility, need for cognition, and user engagement influence attitude change across segments within medium sequences?

We carried out a $2 \times 3 \times 2$ between-subjects exploratory study ($N = 483$) to examine whether users who engage with SERPs before or after listening to an AI-generated podcast experience different levels of attitude change. We investigated whether and how the sequence effects (*SERP-first* versus *podcast-first*) vary across three different viewpoint biases embedded in the content—either supporting, opposing, or neutral, and across topics with varying degrees of controversiality (*moderate* versus *high*). Acknowledging that individual differences can play a significant role in how information is consumed and processed, we explored whether users’ AI literacy, intellectual humility, need for cognition, and their perceived user engagement can moderate the impact of information

medium sequence on their attitude change. By incorporating these variables, our study aims to provide a nuanced understanding of how users engage with and are influenced by multimodal information ecosystems that are increasingly commonplace today.

Original Contributions. We found that users who consumed information in podcasts before SERPs exhibited a significantly higher attitude change. Our results suggest a significant interaction between time and viewpoint bias, such that users who encountered opposing viewpoints were less prone to attitude change, particularly when such information is reinforced in a consequent medium. We also found that moderately controversial topics corresponded to a significantly higher user attitude change compared to highly controversial topics. These key findings and others reported in the paper have important implications for the design of multimodal search systems and responsible content delivery systems. All relevant data, code, and supplementary material can be accessed publicly in an OSF repository for reproducibility.²

2 Background and Related Literature

We position our work in the context of existing literature on (i) web search and opinion formation, (ii) search in the age of generative AI, and (iii) multimodal information access.

2.1 Web Search and Opinion Formation

Web search engines are among the most widely used gateways to information, and the results they deliver have significant power to shape people’s opinions and decisions [10, 12, 59, 80]. When people search for information on debated topics, subjects of ongoing discussion that lack a clear consensus, different factors can hinder responsible opinion formation, which would require actively and thoroughly engaging with diverse viewpoints [59]. For instance, users’ attitudes can be shifted by viewpoint biased search results, when some viewpoints are overrepresented compared to others [18] or elements on the search engine result page like featured snippets [6]. Further, user factors like strong pre-existing attitudes on debated topics were observed to be linked to decreased engagement with diverse viewpoints and a low likelihood of attitude change [60, 75].

These observations highlight the limitations of traditional search interfaces in supporting responsible opinion formation. This raises the question of what alternative information access systems and interfaces that support users in engaging more thoroughly and critically with different viewpoints on debated topics could look like [59, 63] and whether generative AI could serve as a tool that contributes to that goal.

2.2 Search in the Age of Generative AI

Generative AI, and Large Language Models in particular, have shifted many people’s information behavior [81, 87]. Although the precise ways in which they influence information behavior and outcomes is still mostly unclear, recent research has offered some initial insights into the effects of LLM based chat interfaces. For instance, Yang et al. [86] compared information seeking for complex learning tasks with traditional search interfaces to search with chat interface and found that participants engaged less with

²https://osf.io/bn2pd/overview?view_only=2146ca25613c40308d195a592444326

traditional search results when they could use the chat interface, and that the chat interface improved immediate learning, yet not longer term retention. Building on these findings with a follow-up qualitative study Mayerhofer et al. [48] found that while the level of correct responses for complex information tasks found with the search and chat interface was comparable to that achieved with traditional search interfaces, participants often misplaced trust for ease-of-use, leading to high post-search confidence even if their answer was incorrect, particularly when using the chat feature.

While we are beginning to understand some of the immediate effects of widely available LLMs on information behavior and search outcomes, we still know little about longer term consequences of these technologies for individuals and society. Focusing on information access, such systems could, for instance, disrupt the information ecosystem, further increase marginalization and power concentration, and hamper innovation [50].

Thus, GenAI, and LLMs in particular, pose risks to our information ecosystems, yet, they also have the potential to facilitate information access and support information seeking in various ways [63, 79]. For instance, GenAI can enable multimodal ways to access and interact with information, thereby potentially improving engagement, learning outcomes, and overall accessibility [1, 15, 63, 72].

2.3 Multimodal Information Access

Multimodal interfaces enable different modes of interacting with a system, via multiple modes of input and output, such as text, speech, gesture, image, or videos, or with different interaction paradigms, such as query-response or multi-turn dialog [82]. Traditional search engine interfaces have offered information interactions mainly via a single modality: visually, with query-response interactions [15, 82]. However, researchers who investigated alternatives modes of information presentation, and interfaces that support access through other modalities, have observed a range of benefits. Beyond the web search context, researchers found that presenting people with multimedia content during learning tasks (e.g., audio, video, text) reduces cognitive load, improves information comprehension, and increases motivation [14, 40, 62]. Moraes et al. [51] observed that presenting participants with an instructional video combined with a subsequent phase of conducting a web search yielded better learning outcomes than search or instructional video alone.

The advances of generative AI facilitate multimodal approaches to presenting and accessing information that can support deeper engagement with complex information, for instance through conversational interactions [82]. Compared to traditional, primarily visual web search interfaces conversational modes of information access enable more natural, accessible, and in the case of spoken conversational modes, screen-free ways of interacting with information [15]. They also allow for information to be presented not only as facts, but with a richer context, for instance with added narration, which might play an important role in improving human comprehension [61].

Some of the advantages of conversational interactions could also apply to delivering information through AI generated podcasts, which can transform written text into audio content, potentially increasing engagement and accessibility [35, 85]. In an educational

setting, Do et al. [17] observed that students were more engaged by AI-generated podcasts compared to textbook reading, and that podcasts tailored to student profiles improved knowledge retention in some school subjects.

While AI generated podcasts may improve some aspects of information interaction, they also pose various risk, such as misleading content, misrepresentation of complex ideas, reduced critical engagement due to the passive nature of the interaction, or tone-introduced bias in user reliance. These concerns become particularly critical when the content presented can shape people's opinions and decisions, which can be the case with information on debated topics [59]. However, combining AI generated podcasts with web search might leverage some benefits of multimodal information interactions, such as promoting deeper understanding and greater user engagement. To advance the understanding of the potential benefits and risks of AI generated podcasts, we explore how multimodal interactions with information on debated topics, combining AI-generated podcasts with traditional web search interfaces affect attitude change.

3 Method

To understand how different information mediums (SERPs and AI-generated podcasts) interact in sequence to shape the opinions of users, we carried out a $2 \times 3 \times 2$ between-subjects crowdsourcing study that was approved by our university ethics committee. We examined two sequences (*SERP-first* versus *Podcast-first*), across three levels of viewpoint bias (supporting, neutral, and opposing) and topics that varied on the degree of controversy (*moderately* controversial versus *highly* controversial topics, as shown in Table 1). These topics were chosen based on prior literature that has explored the role of viewpoint biases in SERPs, user opinion formation, and attitude change [18, 60, 84]. For instance, topics were selected from ProCon [8], a resource that presents controversial topics and related arguments, with varying levels of controversy—*'Should abortion be legal?'* (*highly* controversial, *i.e.*, people tend to have strong attitudes) and *'Is obesity a disease?'* (*moderately* controversial, *i.e.*, people tend to have moderate attitudes). This allowed us to assess and compare the attitude change effects under different conditions.

3.1 Materials

We first describe the process of generating the content of SERPs and the AI-generated podcasts presented to participants.

3.1.1 SERPs. We constructed SERPs corresponding to each of the six topics as follows:

- (i) For moderately controversial topics, we directly used the publicly available viewpoint annotated dataset released by Draws et al. [18]³, which contains the URL, title, snippet, and annotated viewpoint (7-point scale, from "*strongly opposing*" to "*strongly supporting*") of the search results;
- (ii) For highly controversial topics, we identified corresponding relevant sources on ProCon,⁴ and employed NotebookLM⁵ to annotate the viewpoints represented by the sources in a cost-effective manner [77]. Following best practices [71] and to confirm the

³<https://osf.io/v38c5>

⁴<https://www.britannica.com/procon>

⁵<https://notebooklm.google/>

Table 1: The set of moderately controversial (MC) and highly controversial (HC) topics considered in our study.

Degree of Controversiality	Topic
Moderately Controversial (MC)	Cell phone radiation is safe. (Cell Phone) Social networking sites are good for our society. (Social Networks) Obesity is a disease. (Obesity)
Highly Controversial (HC)	The government should allow more refugees to resettle in the United States. (Immigration) Permitless carry of guns should be legal. (Gun Control) Abortion should be legal. (Abortion)

quality of the viewpoint annotations by NotebookLM—given that large language models can be prone to erroneous annotations [69]—we carried out a validation on a random subset with authors of this paper acting as experts. An example prompt (P1) used for the viewpoint annotation process is presented below.

P1. Please rate the viewpoint of each source about the statement that "Obesity is a disease." You have the following seven choices: **strongly opposing, opposing, somewhat opposing, neutral, somewhat supporting, supporting, strongly supporting.**

Apart from the retrieval and viewpoint annotation of the search results used in our study, we followed the ranking methodology employed by Draws et al. [18] to create SERPs that are either supporting, opposing, or neutral with respect to the topic. We randomly sampled three "opposing", two "somewhat opposing", two "somewhat supporting", and three "supporting" items from the search result items that were deemed relevant to a given topic. For the supporting condition, we used the ranking with extreme bias towards the supporting viewpoint. For the opposing condition, we used the ranking with extreme bias towards the opposing viewpoint. For the neutral condition, we used the ranking with little bias (wherein the supporting viewpoints and opposing viewpoints are interleaved, with half of the results containing bias for the opposing viewpoint and the other half for the supporting viewpoint). Consistent with existing practices around using AI overviews of SERPs (e.g., on Google search), and to help participants grasp key information from the SERPs, we generated and integrated a concise AI overview using the following prompt (P2):

P2. Imagine that a user inputs the query "Is obesity a disease?" into the search box of a search engine and receives these 10 sources in return. Now, please generate a short overview of the search results (the number of words should be less than 150).

3.1.2 AI Podcast. Using the "Audio Overview" feature of NotebookLM, we generated podcasts reflecting different viewpoint biases based on the SERPs. The AI model was instructed to synthesize content from topic sources while adhering to a predefined bias. Example prompts are presented below (P3, P4, P5):

P3. Focus on discussing whether obesity is a disease. Focus on the sources which hold the view that obesity is a disease.

P4. Focus on discussing whether obesity is a disease. Ensure a balance between the sources which hold the view that obesity is a disease and the sources which hold the view that obesity is not a disease.

P5. Focus on discussing whether obesity is a disease. Focus on the sources which hold the view that obesity is not a disease.

Two authors of this paper independently assessed the alignment of the 18 AI-generated podcasts with the intended viewpoint biases (i.e., 6 topics with 3 different viewpoint biases each). Inter-rater reliability was assessed using Krippendorff's α for ordinal data, which accounts for the magnitude of disagreements on an ordered scale. We found $\alpha = 0.896$ (95% CI: 0.78 – 0.97), indicating very high agreement between raters, supporting the reliability of the embedded viewpoint biases in the AI-generated podcasts.

3.2 Measures

Attitude change. We measure *attitude change* as the main dependent variable in our study, drawing inspiration from prior works [18, 60] that have captured participant attitudes before and after exposure to information (e.g., via SERPs). We captured the attitudes of participants corresponding to the different topics at three different points in our study—their pre-existing attitude on the topic (*pre*), their attitude after exposure to the first medium (*mid*), and after exposure to the second medium (*post*). For example, we used the following question: *To what extent do you agree with the following statement? 'Cell phone radiation is safe.'* and capture responses on a 7-point Likert-scale -3: *Strongly Opposing*—to—+3: *Strongly Supporting*. This allows us to measure the participants' attitude change across two piecewise segments: (i) Δ_1 from *pre* to *mid*, and (ii) Δ_2 from *mid* to *post*.

Attitude strength. Luttrell and Sawicki [46] have shed light on the difference between predictors and defining features of *attitude strength*, which is defined as the durability and impact of an attitude. They synthesized existing literature, arguing that strength-related attitude attributes and attitude strength itself are independent. Therefore, we separately capture three key attributes that have been shown to be effective predictors of attitude strength:

- (i) *Importance* (i.e., how much a person perceives him or herself to care about a particular attitude) [19, 27, 33, 42]; we captured importance at the pre-task stage using the following question: *'How important is your attitude on this topic to you?'* on a 7-point Likert-scale (1: *Not important at all*—to—7: *Extremely important*).
- (ii) *Elaboration* (i.e., the extent to which attitudes are formed through careful thinking about relevant information) [5, 30, 34, 55]; we captured elaboration using the following question: *To what extent do you agree with the following statement? "Your attitude on this topic is a result of careful thinking about relevant information."* on a 7-point Likert-scale (1: *Strongly Disagree*—to—7: *Strongly Agree*) at three points in the study (at the pre-task stage, after exposure to the first medium, after exposure to the second medium).
- (iii) *Moralization* (i.e., the degree to which people perceive an attitude as being connected to core moral values) [2, 45, 65, 66]; we captured this at the pre-task stage using the following question:

To what extent do you agree with the following statement? “Your attitude on this topic is connected to your core moral values.” on a 7-point Likert-scale (1: Strongly Disagree—to—7: Strongly Agree).

Belief strength. We captured the *belief strength* of participants in our study. This refers to how strongly a person holds a specific belief to be true or the degree of certainty that one has in the truth of a specific proposition [16, 20, 74]. We used the following two questions to capture belief strengths at three points in the study (at the pre-task stage, after exposure to the first medium, after exposure to the second medium): (i) *How strongly do you believe your answer is based on knowledge rather than intuition?* on a 7-point Likert-scale (1: Strongly Disagree—to—7: Strongly Agree), and (ii) *How certain are you about your beliefs on this topic?* on a 7-point Likert-scale from (1: Completely Uncertain—to—7: Completely Certain).

AI literacy (AIL). Given the context of our study and the real-world state of SERPs with embedded AI overviews and increasingly accessible AI-generated podcasts, we captured *AI literacy* of the participants in our study as a potential moderator. To this end, we employed the 10-item short version of the Meta AI Literacy Scale (MAILS) [11, 39]. MAILS is a scale that helps us measure AI literacy and other psychological competencies that are assumed to support successful interaction with AI.

Intellectual humility (IH). Recognizing gaps in one’s knowledge and that one’s current beliefs might be incorrect can play an important role in shaping attitudes on different topics [58]. *Intellectual humility* is the ability to recognize shortcomings or potential limitations in one’s own point of view [56, 78]. We employed the 12-item Limitations-Owning Intellectual Humility Scale (L-OIHS) since it allows for a relatively more nuanced interpretation and measurement of intellectual humility [28].

Need for cognition (NFC). NFC is a psychological trait representing the extent to which individuals enjoy and engage in effortful cognitive activities [9]. Haugtvedt and Petty [29] found that NFC moderates the persistence and resistance of attitude and belief changes among individuals. Wu et al. [83] found that users with a higher NFC demonstrated different search behavior and different interactions with SERPs than those with a lower NFC. Given the wealth of evidence in existing literature that has demonstrated how NFC influences attitude change and user interactions with information, we measured users’ NFC using the *NFC-6* scale developed by Lins de Holanda Coelho et al. [44].

User engagement (UE). This relates to the extent to which a user is actively involved with a digital experience, that has been shown to be important in the context of web search [88]. We measured user engagement using the UES-SF scale developed by O’Brien et al. [53] and widely used in the context of user interactions with digital media.

3.3 Participants

Based on estimates from a simulated power analysis and accounting for potential participant exclusion, we recruited 550 participants through the Prolific crowdsourcing platform.⁶ All participants were paid an hourly wage of 9 GBP per hour (per the guidelines for a ‘good’ payment on Prolific). We restricted participation to workers in the age group of 18 and 50, who had successfully completed

over 40 tasks and maintained an approval rate of over 90%. The gender ratio was balanced in terms of the female-male ratio (1:1), and due to the topics considered—which are commonly debated in the USA—we restricted participation to workers from USA. To ensure participant reliability, we included 5 questions (that doubled as attention checks in the study) that participants encountered after each medium, and we rewarded participants with bonuses of 0.25 GBP for each correct answer. This design also allowed us to incentivize genuine interaction with the SERPs and the podcast, increasing the external validity of data gathered in our study.

3.4 Procedure

Participants were first informed about the scope of the study, the data that would be gathered, and provided consent to proceed. They were then asked to respond to questions in a pre-task questionnaire, including their existing attitude on a given topic (*pre*), the three predictors of attitude strength, belief strength, AI literacy, intellectual humility, and need for cognition. Next, they were asked to progress through a tutorial that allowed them to familiarize themselves with the task and the interface. After the tutorial, participants moved on to the main phase of the study, where they were asked to engage with information related to the given topic. Depending on the experimental condition, they were either exposed to the SERPs first or the podcast first. Participants were asked to engage with the search results (e.g., by clicking on them to retrieve the linked webpages) and listen to the podcasts, and proceed once they were satisfied with the SERPs they consumed (or the podcast they heard) on the given controversial topic. Following the first medium and before they are directed to the second, we capture the attitude of participants on the topic once again (*mid*), along with their attitude strength (elaboration), and belief strength. After the second medium, we capture the attitude of participants on the topic (*post*), alongside their attitude strength (elaboration), and belief strength. Participants were finally directed to a post-task questionnaire where we captured their perceived engagement and gathered open-ended responses comparing the two media.

3.5 Piecewise DiD Mixed Model

To collectively address our research questions, we adopt a piecewise difference-in-differences (DiD) mixed model. This can be formally represented as follows:

$$Y_{it} = \beta_0 + \beta_1 s_{1t} + \beta_2 s_{2t} + \beta_3^T X_i + \beta_4^T (s_{1t} \cdot X_i) + \beta_5^T (s_{2t} \cdot X_i) + u_{0i} + u_{1i} s_{1t} + u_{2i} s_{2t} + v_{0, \text{topic}(i)} + \varepsilon_{it},$$

where, Y_{it} represents the attitude of participant i at time $t \in \{pre, mid, post\}$; β_0 represents the fixed intercept (i.e., the expected attitude score at *pre*); $s_{1t} = 1[t \in mid, post]$ (piecewise time indicator which captures Δ_1 , the first segment change from *pre* to *mid*); $s_{2t} = 1[t = post]$ (piecewise time indicator which captures Δ_2 , the second segment change from *mid* to *post*); β_1 represents the average change in attitude after the first medium, and β_2 represents the additional change after the second medium; X_i represents the vector of between-subjects predictors: the sequence (SERP-first or Podcast-first), viewpoint bias (supporting,

⁶<https://www.prolific.com/>

opposing, neutral), controversiality (moderate, high), and moderators (*AI literacy, need for cognition, intellectual humility, user engagement*); $\beta_3^T X_i$ represents the main effects of predictors on the baseline attitude (*pre*); $\beta_4^T (s1_t \cdot X_i)$ represents the interactions of the first-segment change Δ_1 with predictors; $\beta_5^T (s2_t \cdot X_i)$ represents the interactions of the second-segment change Δ_2 with predictors; u_{0i} represents the participant-specific intercept (baseline differences), u_{1i} represents the participant-specific deviation in Δ_1 , and u_{2i} represents the participant-specific deviation in Δ_2 for participant i ; $v_{0, \text{topic}(i)}$ represents the random intercept for topic to account for topic-level differences

4 Results and Analysis

Descriptive Insights. To ensure data quality, we assessed participants’ attention levels using a heuristic composed of: (i) time spent on the SERPs/podcast, and (ii) accuracy on attention check questions related to the content. These were chosen since the time spent and accuracy on attention check questions together have been shown to serve as a reasonably good proxy for reliability in similar crowdsourced studies [24, 60]. For each experimental condition, we first computed the *median time spent* and the *median accuracy rate* across all participants. We then computed an *attention score* for each participant and each medium (SERP, podcast) as follows:

$$\text{Attention Score} = \frac{\text{Time Spent}}{\text{Median Time}} + \frac{\text{Accuracy Rate}}{\text{Median Accuracy Rate}}$$

Participants whose attention scores were below 1 in both mediums were excluded from subsequent analyses. The resulting distribution of participants in each group can be seen in Table 2.

Table 2: Number of reliable (and excluded) participants in each group considered in our analysis, split by topic (degree of controversiality) and across the three viewpoint biases.

Topic	Supporting	Neutral	Opposing
Cell Phone (MC)	27 (7)	28 (6)	27 (7)
Social Networks (MC)	26 (4)	30 (0)	29 (1)
Obesity (MC)	26 (4)	29 (1)	26 (4)
Immigration (HC)	26 (4)	28 (2)	29 (1)
Gun Control (HC)	26 (4)	24 (6)	27 (3)
Abortion (HC)	25 (5)	25 (5)	25 (5)
Total	156 (28)	164 (20)	163 (21)

Note that data pertaining to the attitude strength predictors (importance, elaboration, moralization) and belief strength are included in the online repository, beyond the scope of our primary analysis presented in the paper.

We analyzed attitude change among users across the different conditions to identify and categorize the instances where user attitudes *flipped* (i.e., they changed to the opposing view), *strengthened* (i.e., user attitudes were further reinforced), *weakened*, or remained *unchanged*. Table 3 presents the distribution of attitude change outcomes based on the sequence medium. We found that the majority of users demonstrate unchanged attitudes, followed by nearly one-fourth of the users demonstrating a strengthening of attitudes, and the remaining users either reported flipped attitudes or weakened attitudes. This general trend was consistent across both sequences of *Podcast-first* and *SERP-first*.

Table 3: Distribution of attitude change outcomes among users by sequence medium (*Podcast-first* versus *SERP-first*).

Sequence Medium	Outcome	#Users (%)
Podcast-first (243)	Flipped	43 (17.7%)
	Strengthened	59 (24.3%)
	Unchanged	101 (41.6%)
	Weakened	40 (16.5%)
SERP-first (240)	Flipped	37 (15.4%)
	Strengthened	62 (25.8%)
	Unchanged	99 (41.2%)
	Weakened	42 (17.5%)

We also explored the distribution of attitude change outcomes among users across topics corresponding to a *high* versus *moderate* degree of controversiality (Table 4). Once again, we found similar trends in the outcomes.

Table 4: Distribution of attitude change outcomes among users by degree of controversiality (*High* versus *Moderate*).

Controversiality	Outcome	#Users (%)
High (235)	Flipped	34 (14.5%)
	Strengthened	59 (25.1%)
	Unchanged	107 (45.5%)
	Weakened	35 (14.9%)
Moderate (248)	Flipped	46 (18.5%)
	Strengthened	62 (25.0%)
	Unchanged	93 (37.5%)
	Weakened	47 (19.0%)

Model Overview. Our model incorporated random intercepts for both participants and topics to account for unobserved heterogeneity in baseline attitudes of users. The participant-level variance was found to be $\sigma^2 = 1.55$, indicating substantial individual differences in initial attitudes. This corroborates the importance of accounting for such variability while modeling attitudes longitudinally. This random effect captures how participants differ in their baseline responses, independent of the fixed predictors. Similarly, the topic-level variance was found to be $\sigma^2 = 1.418$, reflecting meaningful variation in users’ attitudes across different topics. This supports the inclusion of topic as a random effect, acknowledging that some topics inherently elicit stronger or weaker attitudes regardless of the medium sequence or the viewpoint bias.

Overall, these random effects enhance our model’s robustness by allowing for more accurate estimation of fixed effects while controlling for nested data structures and contextual variability.

4.1 RQ1: Effect of information medium sequence on attitude change

To examine how various predictors influenced users’ attitude changes (Δ_1 and Δ_2), we fit a piecewise difference-in-differences mixed-effects model. The fixed effects component of the model revealed that the sequence of information exposure significantly impacted user attitudes. Per our model described in Section 3.5, this corresponds to β_3 , which captures the main effect of medium

sequence on baseline attitude (*pre*), and through interaction terms in β_4 and β_5 for segment-specific changes (Δ_1 and Δ_2). For RQ1, we focus on the main effect of sequence on users' attitude change.

Users who encountered SERPs first exhibited a lower attitude change compared to those who engaged with podcasts first ($\beta = -0.376, p = .034; 95\% \text{ CI } [-0.72, -0.03]$). This negative coefficient suggests that, holding other factors constant, beginning with SERPs first led to lower attitude change scores relative to starting with podcasts. The standardized coefficient for medium sequence was $\beta^* = -0.189$, indicating a small-to-moderate effect size. The unique contribution of medium sequence to variance explained was partial $R^2 = .0031$, suggesting that while statistically significant, the effect accounts for about 0.3% of residual variance. Cohen's d for the contrast between SERP-first and Podcast-first was $d = -0.10$, reinforcing that the practical impact is small but meaningful in the context of attitude change. Interactions with segments were not significant, indicating that medium sequence primarily influenced baseline attitudes rather than segment-specific changes.

Key takeaway (1): Information modality (*i.e.*, SERPs or podcasts) and order of information presentation can shape initial impressions and subsequent attitude formation when these mediums are consumed in sequence.

4.2 RQ2: Does the effect of sequence on attitude change differ across viewpoint biases?

On exploring the effect of medium sequence on attitude change across different viewpoint biases, we found that the opposing viewpoint bias had a negative but non-significant main effect on users' attitude change ($\beta = -0.304, p = .167, \beta^* = -0.153$), while supporting viewpoint bias was negligible ($\beta = -0.076, p = .731, \beta^* = -0.038$).

The interaction between time and viewpoint bias was significant in the second segment (*mid*→*post*) with a small effect ($\beta = -0.297, p = .039; 95\% \text{ CI } [-0.59, -0.02]$, standardized $\beta^* = -0.15$, partial $R^2 = .003$). This interaction explains 0.3% of the residual variance, which is statistically significant but practically small. Because higher attitude change scores indicate a favorable change in attitude toward a given topic, a negative coefficient means that attitudes became less favorable (*i.e.*, more negative) during this segment. Thus, when content was opposing, the additional change from *mid*→*post* was more negative compared to neutral bias. We did not find any significant interactions between medium sequence and viewpoint bias in the first segment.

Key takeaway (2): During the second segment (*mid*→*post*), users who encountered viewpoint opposing information became less favorable toward the topic compared to those exposed to neutral information, indicating that viewpoint opposing content drives a small but meaningful negative shift in user attitudes.

To examine whether viewpoint bias alters attitude change within each medium sequence, we compared segment changes for content with supporting and opposing biases against neutral content (the baseline) within sequence (*SERP-first* vs. *Podcast-first*) using bootstrap contrasts. Table 5 presents the sequence differences in segment-specific attitude change (Δ_1 and Δ_2) across viewpoint biases. A positive difference here indicates that the viewpoint bias condition produced a greater increase (or less decrease) in users' attitude (*i.e.*, indicating a more favorable shift in attitude toward the

Table 5: Differences in users' attitude change across viewpoint biases and sequences, relative to the neutral baseline.

Viewpoint Bias	Sequence	Difference	95% CI	p-value
Attitude change from <i>pre</i>→<i>mid</i> ($\Delta_1 = \text{mid} - \text{pre}$)				
Supporting	SERP-first	-0.163	[-0.582, 0.291]	0.46
Opposing	SERP-first	+0.350	[-0.224, 0.851]	0.20
Supporting	Podcast-first	+0.641	[0.024, 1.245]	0.04
Opposing	Podcast-first	-0.440	[-1.003, 0.118]	0.13
Attitude change from <i>mid</i>→<i>post</i> ($\Delta_2 = \text{post} - \text{mid}$)				
Supporting	SERP-first	+0.174	[-0.233, 0.57]	0.40
Opposing	SERP-first	-0.510	[-0.931, -0.117]	0.02
Supporting	Podcast-first	-0.282	[-0.631, 0.01]	0.09
Opposing	Podcast-first	-0.166	[-0.579, 0.201]	0.39

topic than the neutral condition for that segment). A negative difference indicates a greater decrease (or a smaller increase) in users' attitude than neutral (*i.e.*, indicating a relatively less favorable shift in attitude toward the topic). We found a statistically significant effect for the *Podcast-first* sequence in Δ_1 in the supporting viewpoint bias condition; (95% CI: [0.024, 1.245], $p = .04$), indicating that supporting content led to a larger increase in user attitudes favorably towards a topic than neutral content during the first segment. In addition, we found a significant effect for the *SERP-first* sequence in Δ_2 in the opposing viewpoint bias case; (95% CI: [-0.931, -0.117], $p = .021$), indicating that exposure to viewpoint opposing information led participants to become less favorable toward the topic relative to those exposed to neutral information.

In the Figures 1 and 2 we plot the predicted attitude trajectories using the proposed piecewise difference-in-differences mixed effects model across three different time points (*viz.* *pre*, *mid*, and *post*) where user attitudes on the topics were captured in our study.




Figure 1: Predicted attitude trajectories for each viewpoint bias across the three time points (*pre*, *mid*, *post*) for moderately controversial topics. Shaded regions represent 95% CIs.

Figure 1 corresponds to moderately controversial topics and Figure 2 to highly controversial topics, respectively. Addressing RQ2, these plots contrast the sequence of mediums; *SERP-first* versus *Podcast-first*. These trajectory plots reveal distinct order effects across the different viewpoint biases.

Key Takeaway (3): Within the *SERP-first* sequence, opposing content produced a significantly more negative attitude change in the second segment (Δ_2) relative to neutral, indicating that attitudes became less favorable from *mid*→*post* when the information was opposing. Within the *Podcast-first* sequence, supporting content produced a significantly more positive attitude change in the first segment (Δ_1) relative to neutral, indicating a larger gain in favorability from *pre*→*mid* when the information was supportive.




Figure 2: Predicted attitude trajectories for each viewpoint bias across the three time points (*pre*, *mid*, *post*) for highly controversial topics. Shaded regions represent 95% CIs.

4.3 RQ3: Does topic controversiality influence users attitude change across segments within information medium sequences?

Figure 3 illustrates the predicted attitude trajectories for topics classified as *Moderately* versus *Highly* controversial, holding the neutral viewpoint bias constant.




Figure 3: Predicted attitude trajectories for topics with varying controversiality (*moderate*, *high*) across the time points (*pre*, *mid*, *post*), separated by the medium sequence (*SERP-first*, *Podcast-first*), and with viewpoint bias fixed to neutral.

We found that the main effect of controversiality was significant—moderately controversial topics were associated with more favorable baseline attitudes compared to highly controversial topics ($\beta = 0.803, p < .001, 95\% \text{ CI } [0.457, 1.149]$), with a standardized coefficient of $\beta^* = 0.405$ and partial $R^2 = .0143$. The interaction between the degree of controversiality and Δ_1 (*pre*→*mid*) was significant with a moderate effect ($\beta = -0.660, p < .001; \beta^* = -0.333$, partial $R^2 = 0.127$). Since higher attitude scores indicate more favorable attitudes towards topics, this negative coefficient indicates that moderately controversial topics corresponded to a smaller increase (or a greater decrease) in favorability during the first segment (*pre*→*mid*) compared to highly controversial topics. This interaction explains about 1.3% of residual variance, which is small but meaningful in the context attitude change.

Key Takeaway (4): Users reported more favorable attitudes toward moderately controversial topics to begin with, but these topics corresponded to less positive change (or more negative change) during the first segment of the medium sequence compared to highly controversial topics. In the second segment of the medium

sequence, we found that controversiality did not influence attitude change significantly.

4.4 RQ4: Do individual differences influence attitude change within medium sequences?

As reported in the Table 6 below, we found that interactions involving all the individual difference moderators considered in our study (AI literacy, intellectual humility, need for cognition, and user engagement) were not statistically significant across segments.

Table 6: Non-significant interactions between sequence and individual moderators over time.

Interaction Term	Coefficient (β)	p -value
$s_1 \times \text{AI Literacy}$	0.014	.784
$s_2 \times \text{AI Literacy}$	0.004	.915
$s_1 \times \text{Intellectual Humility}$	0.023	.768
$s_2 \times \text{Intellectual Humility}$	0.051	.389
$s_1 \times \text{Need for Cognition}$	-0.040	.496
$s_2 \times \text{Need for Cognition}$	-0.043	.327
$s_1 \times \text{User Engagement}$	0.231	.138
$s_2 \times \text{User Engagement}$	0.006	.961

5 Discussion, Implications, and Future Work

Major search platforms today are increasingly integrating AI-generated audio summaries and podcast-like outputs into search experiences. Tools such as NotebookLM, Perplexity’s audio mode, and Google’s experimental “*Listen*” features illustrate a trend toward audio-augmented search, where users can move fluidly between traditional search results and AI-generated audio. Switching between SERPs and AI-generated podcasts reflects an emerging interaction pattern in modern search ecosystems.

Our work is situated in this emerging and important context. In our study addressing the interaction between mediums of SERPs and podcasts, we observed that participants who engaged with the podcasts first showed increased attitude change compared to those who engaged with the SERPs first (RQ1). This observed ability of AI-generated podcasts to shift attitudes more easily than web search mirrors the duality of benefits and risks of the podcast medium, promoting deeper engagement with different viewpoints but also posing a risk of undue influence on beliefs and the potential spread of misinformation [36, 54]. Our results demonstrate that exposure to opposing viewpoints can shift user attitudes less favorably towards a topic after initial exposure, potentially reflecting reactance, conflict, or belief consolidation processes in the second sequence medium (RQ2). Similar observations have been made in the context of web search on debated topics [60]. Participants began with more favorable attitudes toward moderately controversial topics, but early change (*i.e.*, in Δ_1) was smaller for these topics relative to highly controversial ones—consistent with ceiling effects, early consolidation, or lower initial uncertainty (RQ3). In the second segment (Δ_2), controversiality did not shape attitude change. This finding is not surprising, given that opinions on topics with varying controversiality are often grounded in core values, and thereby, be more or less resistant to change [19].

Table 7: Qualitative comparison of podcasts, SERPs, and their combination synthesized through participants' feedback.

Medium	Advantages	Limitations
Podcast	<ul style="list-style-type: none"> • Engagement: More engaging and easier to follow, especially for auditory learners. • Convenience: Allows multitasking (e.g., listening while driving or cooking). • Narrative clarity: Simplifies complex topics through conversational tones and storytelling. • Retention: Improves memorization due to emotional and contextual delivery. 	<ul style="list-style-type: none"> • Lack of source verification. • Potential bias or repetitive content.
SERP	<ul style="list-style-type: none"> • Depth and diversity: Provides access to multiple sources, studies, and viewpoints. • Fact-checking: Enables users to verify claims and explore primary sources. • Control: Users can skim or dive deep into specific details at their own pace. 	<ul style="list-style-type: none"> • Overwhelming volume of information. • Time-consuming to sift through.
Combination	<ul style="list-style-type: none"> • Comprehensive understanding: Combines the depth of SERP with the accessibility of podcasts. • Balanced learning: SERP offers factual rigor, while podcasts provide narrative context. • Critical thinking: Encourages cross-referencing and reduces bias. 	<ul style="list-style-type: none"> • Requires more time and effort to engage with both mediums.

Understanding user preferences and trade-offs. We asked to indicate their preferences among the information mediums to engage with: SERPs, AI-generated podcasts, or a combination of both. Our results revealed that a majority of participants (58%) preferred the combination of both mediums in sequence, while fewer favored either only podcasts (24%) or SERPs (19%) for learning the topic. The rationales provided for these preferences, as indicated by the participants, are summarized in Table 7. In their open-ended feedback, participants highlighted a core tension in their information-seeking experiences: a trade-off between *depth* and *accessibility*. Podcasts were recognized for being engaging, memorable, and easy to incorporate into daily routines. However, participants also acknowledged limitations of podcasts, particularly the difficulty of verifying sources and the potential lack of depth or breadth when compared to text-based formats. In contrast, SERPs were valued for offering comprehensive coverage and clear source traceability, though participants noted that this came at the cost of increased cognitive effort and time.

The complementary nature of these strengths and weaknesses suggests that users may benefit most from a hybrid information access ecosystem, that enables smooth transitions between the narrative richness of podcasts and the structured, traceable detail of search results. Participants reflected that such an approach can help them quickly grasp key ideas while still allowing them to explore deeper details when needed. However, participants also noted that this hybrid model introduces new challenges; notably, increased time demands and the effort required to navigate between modalities. These observations raise important design considerations for future information systems: *How might search engines and content platforms integrate audio narratives and textual results in ways that reduce friction rather than amplify it?*

We identify several promising directions for future work. First, future systems might embed AI-generated podcast snippets directly within SERPs, giving users a rapid auditory overview alongside traditional links. Second, drawing on advances in retrieval-augmented generation [23], AI-produced podcasts could be paired with automatically retrieved evidence, providing the dual benefits of engaging audio guidance and text-based verifiability. A practical example already emerging in some platforms is the inclusion of linked

sources within podcast transcripts, enabling users to move seamlessly from listening to exploring relevant SERPs on demand. Together, these approaches highlight opportunities to design hybrid information experiences that preserve accessibility while enhancing depth and trustworthiness.

Caveats and Limitations. Given the complexity of studying how information medium sequences (SERPs, podcasts) influence user attitude changes, there are a few important limitations of our work. It is worth noting that we leveraged NotebookLM to create the AI-generated podcasts, which result in two speakers (female, male) discussing the content. Understanding how the observed effects would change based on such configurations of podcasts is beyond the scope of our work, but it is an important avenue for future research. In our study, we considered users being exposed to consistent viewpoint biases across the two information mediums that they encountered in sequence. Future work should explore the effect of medium sequence on attitude change, varying viewpoint biases across the mediums encountered in sequence.

6 Conclusions

Our study is the first of its kind and addresses an important research gap by offering a systematic analysis of how information modality, content bias, and user traits interact to shape and shift user attitudes. Through a controlled study, we explored how information medium sequences (with SERPs, Podcasts) influence user attitude change, and how such effects are influenced by viewpoint biases in the content that is consumed, and by the degree of controversiality of the topics. Our findings present strong evidence that calls for thoughtful design and integration of different information media in search ecosystems. Given the growing consumption of information across different mediums and modalities and the evolving information-seeking behavior today, we believe that the CHIIR community is uniquely positioned to help navigate these complex and important considerations. We hope that our work inspires further research on responsible opinion formation in multimodal contexts.

Acknowledgments

We thank all the anonymous participants in our study. This work was supported by the TU Delft AI Initiative, the Model Driven

Decisions Lab (*MoDDL*), the *ProtectMe* Convergence Flagship, and a grant from the TU Delft Fast Fund. We used Microsoft Copilot to assist with utility code for analyses and plots, and all generated code was independently reviewed and validated by the authors.

References

- [1] Mohammad Aliannejadi, Jacek Gwizdka, and Hamed Zamani. 2025. Interactions with Generative Information Retrieval Systems. In *Information Access in the Era of Generative AI*, Ryan W. White and Chirag Shah (Eds.). Springer Nature Switzerland, Cham, 47–71. doi:10.1007/978-3-031-73147-1_3
- [2] Nicholas P Aramovich, Brad L Lytle, and Linda J Skitka. 2012. Opposing torture: Moral conviction and resistance to majority influence. *Social Influence* 7, 1 (2012), 21–34.
- [3] Sandeep Avula, Jaime Arguello, Robert Capra, Jordan Dodson, Yuhui Huang, and Filip Radlinski. 2019. Embedding search into a conversational platform to support collaborative search. In *Proceedings of the 2019 Conference on Human Information Interaction and Retrieval*. 15–23.
- [4] Sandeep Avula, Gordon Chadwick, Jaime Arguello, and Robert Capra. 2018. Searchbots: User engagement with chatbots during collaborative search. In *Proceedings of the 2018 conference on human information interaction & retrieval*. 52–61.
- [5] Jamie Barden and Zakary L Tormala. 2014. Elaboration and attitude strength: The new meta-cognitive perspective. *Social and personality psychology compass* 8, 1 (2014), 17–29.
- [6] Markus Bink, Sebastian Schwarz, Tim Draws, and David Elsweiler. 2023. Investigating the influence of featured snippets on user attitudes. In *Proceedings of the 2023 Conference on Human Information Interaction and Retrieval*. 211–220.
- [7] Shreyan Biswas, Alexander Erlei, and Ujwal Gadiraju. 2025. Mind the Gap! Choice Independence in Using Multilingual LLMs for Persuasive Co-Writing Tasks in Different Languages. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–20.
- [8] Britannica ProCon.org. 2023. Britannica ProCon.org. <https://www.procon.org/>. [Online; accessed 2023-02-01].
- [9] John T Cacioppo and Richard E Petty. 1982. The need for cognition. *Journal of personality and social psychology* 42, 1 (1982), 116.
- [10] Cansu Canca. 2022. Did You Find It on the Internet? Ethical Complexities of Search Engine Rankings.
- [11] Astrid Carolus, Martin J Koch, Samantha Straka, Marc Erich Latoschik, and Carolin Wienrich. 2023. MAILS-Meta AI literacy scale: Development and testing of an AI literacy questionnaire based on well-founded competency models and psychological change-and meta-competencies. *Computers in Human Behavior: Artificial Humans* 1, 2 (2023), 100014.
- [12] Andrés Chacoma and Damián H Zanette. 2015. Opinion formation by social influence: From experiments to modeling. *PLoS one* 10, 10 (2015), e0140406.
- [13] Leigh Clark, Nadia Pantidi, Orla Cooney, Philip Doyle, Diego Garaialde, Justin Edwards, Brendan Spillane, Emer Gilmartin, Christine Murad, Cosmin Munteanu, et al. 2019. What makes a good conversation? Challenges in designing truly conversational agents. In *Proceedings of the 2019 CHI conference on human factors in computing systems*. 1–12.
- [14] Ruichen Cong, Kiichi Tago, and Qun Jin. 2022. Measurement and verification of cognitive load in multimedia presentation using an eye tracker. *Multimedia Tools and Applications* 81, 19 (2022), 26821–26835.
- [15] Yashar Deldjo, Johanne R. Trippas, and Hamed Zamani. 2021. Towards Multi-Modal Conversational Information Seeking. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '21)*. Association for Computing Machinery, New York, NY, USA, 1577–1587. doi:10.1145/3404835.3462806
- [16] Allyson Dennin, Kayla Furman, Jean E Pretz, and Michael M Roy. 2022. The relationship of types of intuition to thinking styles, beliefs, and cognitions. *Journal of Behavioral Decision Making* 35, 5 (2022), e2283.
- [17] Tiffany D Do, Usama Bin Shafiq, Elsie Ling, and Nikhil Sarda. 2024. PAIGE: Examining Learning Outcomes and Experiences with Personalized AI-Generated Educational Podcasts. *arXiv preprint arXiv:2409.04645* (2024).
- [18] Tim Draws, Nava Tintarev, Ujwal Gadiraju, Alessandro Bozzon, and Benjamin Timmermans. 2021. This is not what we ordered: Exploring why biased search result rankings affect user attitudes on debated topics. In *Proceedings of the 44th international ACM SIGIR conference on research and development in information retrieval*. 295–305.
- [19] Asia A Eaton and Penny S Visser. 2008. Attitude importance: Understanding the causes and consequences of passionately held views. *Social and Personality Psychology Compass* 2, 4 (2008), 1719–1736.
- [20] Owen Egan. 1986. The concept of belief in cognitive theory. In *Annals of theoretical psychology*. Springer, 315–350.
- [21] Robert Epstein, Savannah Aries, Kelly Grebbien, Alyssa M Salcedo, and Vanessa R Zankich. 2024. The search suggestion effect (SSE): A quantification of how autocomplete search suggestions could be used to impact opinions and votes. *Computers in Human Behavior* 160 (2024), 108342.
- [22] Robert Epstein, Ronald E Robertson, David Lazer, and Christo Wilson. 2017. Suppressing the search engine manipulation effect (SEME). *Proceedings of the ACM on Human-Computer Interaction* 1, CSCW (2017), 1–22.
- [23] Wenqi Fan, Yujuan Ding, Liangbo Ning, Shijie Wang, Hengyun Li, Dawei Yin, Tat-Seng Chua, and Qing Li. 2024. A survey on rag meeting llms: Towards retrieval-augmented large language models. In *Proceedings of the 30th ACM SIGKDD conference on knowledge discovery and data mining*. 6491–6501.
- [24] Ujwal Gadiraju, Ran Yu, Stefan Dietze, and Peter Holtz. 2018. Analyzing knowledge gain of users in informational search sessions on the web. In *Proceedings of the 2018 conference on human information interaction & retrieval*. 2–11.
- [25] Jonathan Grudin and Richard Jacques. 2019. Chatbots, humbots, and the quest for artificial general intelligence. In *Proceedings of the 2019 CHI conference on human factors in computing systems*. 1–11.
- [26] Kobi Hackenburg, Ben M Tappin, Luke Hewitt, Ed Saunders, Sid Black, Hause Lin, Catherine Fist, Helen Margetts, David G Rand, and Christopher Summerfield. 2025. The levers of political persuasion with conversational artificial intelligence. *Science* 390, 6777 (2025), eaea3884.
- [27] Geoffrey Haddock, Alexander J Rothman, Rolf Reber, and Norbert Schwarz. 1999. Forming judgments of attitude certainty, intensity, and importance: The role of subjective experiences. *Personality and Social Psychology Bulletin* 25, 7 (1999), 771–782.
- [28] Megan Haggard, Wade C Rowatt, Joseph C Leman, Benjamin Meagher, Courtney Moore, Thomas Fergus, Dennis Whitcomb, Heather Battaly, Jason Baehr, and Dan Howard-Snyder. 2018. Finding middle ground between intellectual arrogance and intellectual servility: Development and assessment of the limitations-owning intellectual humility scale. *Personality and Individual Differences* 124 (2018), 184–193.
- [29] Curtis P Haugtvedt and Richard E Petty. 1992. Personality and persuasion: Need for cognition moderates the persistence and resistance of attitude changes. *Journal of Personality and Social Psychology* 63, 2 (1992), 308.
- [30] Curtis P Haugtvedt and Duane T Wegener. 1994. Message order effects in persuasion: An attitude strength perspective. *Journal of consumer research* 21, 1 (1994), 205–218.
- [31] Gaole He, Nilay Aishwarya, and Ujwal Gadiraju. 2025. Is conversational xai all you need? human-ai decision making with a conversational xai assistant. In *Proceedings of the 30th International Conference on Intelligent User Interfaces*. 907–924.
- [32] Gaole He, Gianluca Demartini, and Ujwal Gadiraju. 2025. Plan-then-execute: An empirical study of user trust and team performance when using llm agents as a daily assistant. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–22.
- [33] Allyson L Holbrook, Matthew K Berent, Jon A Krosnick, Penny S Visser, and David S Boninger. 2005. Attitude importance and the accumulation of attitude-relevant knowledge in memory. *Journal of personality and social psychology* 88, 5 (2005), 749.
- [34] Javier Horcajo and Andrew Luttrell. 2016. The effects of elaboration on the strength of doping-related attitudes: Resistance to change and behavioral intentions. *Journal of Sport and Exercise Psychology* 38, 3 (2016), 236–246.
- [35] Paul Huffman and James Hutson. 2024. Enhancing History Education with Google NotebookLM: Case Study of Mary Easton Sibley’s Diary for Multimedia Content and Podcast Creation. *ISRG Journal of Arts, Humanities and Social Sciences* 2, 5 (2024).
- [36] Matthew Jacobson. 2021. Podcast Rhetorics: Insights into Podcasts as Public Persuasion. (Aug. 2021). hdl:11244/330206
- [37] Hyunhoon Jung, Changhoon Oh, Gillwan Hwang, Cindy Yoonjung Oh, Joonhwan Lee, and Bongwon Suh. 2019. Tell me more: Understanding user interaction of smart speaker news powered by conversational search. In *Extended abstracts of the 2019 chi conference on human factors in computing systems*. 1–6.
- [38] Ji-Youn Jung, Sihang Qiu, Alessandro Bozzon, and Ujwal Gadiraju. 2022. Great chain of agents: The role of metaphorical representation of agents in conversational crowdsourcing. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. 1–22.
- [39] Martin J Koch, Astrid Carolus, Carolin Wienrich, and Marc E Latoschik. 2024. Meta AI literacy scale: Further validation and development of a short version. *Heliyon* 10, 21 (2024).
- [40] Tomislava Lauc, Gordana Kuterovac Jagodic, and Jelena Bistrovic. 2020. Effects of Multimedia Instructional Message on Motivation and Academic Performance of Elementary School Students in Croatia. *International Journal of Instruction* 13, 4 (2020), 491–508.
- [41] Kwan Min Lee. 2004. Presence, explicated. *Communication theory* 14, 1 (2004), 27–50.
- [42] Thomas J Leeper and Joshua Robison. 2020. More important, but for what exactly? The insignificant role of subjective issue importance in vote decisions. *Political Behavior* 42, 1 (2020), 239–259.
- [43] Q Vera Liao, Werner Geyer, Michael Muller, and Yasaman Khazaen. 2020. Conversational interfaces for information search. In *Understanding and Improving*

- Information Search: A Cognitive Approach*. Springer, 267–287.
- [44] Gabriel Lins de Holanda Coelho, Paul HP Hanel, and Lukas J. Wolf. 2020. The very efficient assessment of need for cognition: Developing a six-item version. *Assessment* 27, 8 (2020), 1870–1885.
- [45] Andrew Luttrell, Richard E Petty, Pablo Briñol, and Benjamin C Wagner. 2016. Making it moral: Merely labeling an attitude as moral increases its strength. *Journal of Experimental Social Psychology* 65 (2016), 82–93.
- [46] Andrew Luttrell and Vanessa Sawicki. 2020. Attitude strength: Distinguishing predictors versus defining features. *Social and Personality Psychology Compass* 14, 8 (2020), e12555.
- [47] Sandra C Matz, Jacob D Teeny, Sumer S Vaid, Heinrich Peters, Gabriella M Harari, and Moran Cerf. 2024. The potential of generative AI for personalized persuasion at scale. *Scientific Reports* 14, 1 (2024), 4692.
- [48] Kerstin Mayerhofer, Rob Capra, and David Elweiler. 2025. Blending Queries and Conversations: Understanding Trust, Verification, and System Choice in Search and Chat Interactions. In *Proceedings of the 2025 ACM SIGIR Conference on Human Information Interaction and Retrieval (CHIIR '25)*. Association for Computing Machinery, New York, NY, USA, 168–178. doi:10.1145/3698204.3716454
- [49] Steven McClung and Kristine Johnson. 2010. Examining the motives of podcast users. *Journal of radio & audio media* 17, 1 (2010), 82–95.
- [50] Bhaskar Mitra, Henriette Cramer, and Olya Gurevich. 2025. Sociotechnical Implications of Generative Artificial Intelligence for Information Access. In *Information Access in the Era of Generative AI*, Ryen W. White and Chirag Shah (Eds.). Springer Nature Switzerland, Cham, 161–200. doi:10.1007/978-3-031-73147-1_7
- [51] Felipe Moraes, Sindunuraga Rikarno Putra, and Claudia Hauff. 2018. Contrasting Search as a Learning Activity with Instructor-designed Learning. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management (CIKM '18)*. Association for Computing Machinery, New York, NY, USA, 167–176. doi:10.1145/3269206.3271676
- [52] Spotify Newsroom. 2023. International Podcast Day: Transcripts, Chapters, Show Pages, and More. <https://newsroom.spotify.com/2023-09-28/international-podcast-day-transcripts-chapters-show-pages-global/>. Accessed: 2024-11-02.
- [53] Heather L O'Brien, Paul Cairns, and Mark Hall. 2018. A practical approach to measuring user engagement with the refined user engagement scale (UES) and new UES short form. *International Journal of Human-Computer Studies* 112 (2018), 28–39.
- [54] Sachin Pathiyar Cherumanal, Ujwal Gadiraju, and Damiano Spina. 2024. Everything we hear: Towards tackling misinformation in podcasts. In *Proceedings of the 26th International Conference on Multimodal Interaction*. 596–601.
- [55] Richard E Petty, Curtis P Haugtvedt, and Stephen M Smith. 2014. Elaboration as a determinant of attitude strength: Creating attitudes that are persistent, resistant, and predictive of behavior. In *Attitude strength*. Psychology Press, 93–130.
- [56] Tenelle Porter, Abdo Elnakouri, Ethan A Meyers, Takuya Shibayama, Eranda Jayawickreme, and Igor Grossmann. 2022. Predictors and consequences of intellectual humility. *Nature Reviews Psychology* 1, 9 (2022), 524–536.
- [57] Filip Radlinski and Nick Craswell. 2017. A theoretical framework for conversational search. In *Proceedings of the 2017 conference on conference human information interaction and retrieval*. 117–126.
- [58] Alisa Rieger, Frank Bredius, Mariët Theune, and Maria Soledad Pera. 2024. From Potential to Practice: Intellectual Humility During Search on Debated Topics. In *Proceedings of the 2024 Conference on Human Information Interaction and Retrieval*. 130–141.
- [59] Alisa Rieger, Tim Draws, Nicolas Mattis, David Maxwell, David Elweiler, Ujwal Gadiraju, Dana McKay, Alessandro Bozzon, and Maria Soledad Pera. 2024. Responsible Opinion Formation on Debated Topics in Web Search. In *European Conference on Information Retrieval*. Springer, 437–465.
- [60] Alisa Rieger, Suleiman Kulane, Ujwal Gadiraju, and Maria Soledad Pera. 2024. Disentangling Web Search on Debated Topics: A User-Centered Exploration. In *Proceedings of the 32nd ACM Conference on User Modeling, Adaptation and Personalization*. 24–35.
- [61] Vahid Sadiri Javadi, Johanne R Trippas, and Lucie Flek. 2024. Unveiling Information Through Narrative In Conversational Information Seeking. In *Proceedings of the 6th ACM Conference on Conversational User Interfaces (CUI '24)*. Association for Computing Machinery, New York, NY, USA, 1–6. doi:10.1145/3640794.3665884
- [62] Dwi Septiani. [n. d.]. DEVELOPMENT OF AUDIO-VISUAL BASED LEARNING MEDIA IN IMPROVING STUDENT PRESENTATION ABILITY IN RETORICS COURSES. *The Economic Optimism 2021 Perspectives: Education, Business, Health, Engineering and Technology, Productive Enterprises and Law* [n. d.].
- [63] Chirag Shah and Emily M. Bender. 2024. Envisioning Information Access Systems: What Makes for Good Tools and a Healthy Web? *ACM Trans. Web* 18, 3 (April 2024), 33:1–33:24. doi:10.1145/3649468
- [64] Chenglei Si, Navita Goyal, Tongshuang Wu, Chen Zhao, Shi Feng, Hal Daumé Iii, and Jordan Boyd-Graber. 2024. Large Language Models Help Humans Verify Truthfulness—Except When They Are Convincingly Wrong. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*. 1459–1474.
- [65] Linda J Skitka. 2014. The psychological foundations of moral conviction. *Advances in experimental moral psychology* 148 (2014).
- [66] Linda J Skitka and G Scott Morgan. 2014. The social and political implications of moral conviction. *Political psychology* 35 (2014), 95–110.
- [67] Statista. 2024. Number of podcast listeners worldwide from 2019 to 2024. <https://www.statista.com/statistics/1291360/podcast-listeners-worldwide/>. Accessed: 2024-06-20.
- [68] Statista. 2024. Number of podcasts and podcast episodes worldwide in 2023. <https://www.statista.com/statistics/1418185/podcasts-and-podcasts-episode-worldwide/>. Accessed: 2024-06-20.
- [69] Zhen Tan, Dawei Li, Song Wang, Alimohammad Beigi, Bohan Jiang, Amrita Bhattacharjee, Mansoor Karami, Jundong Li, Lu Cheng, and Huan Liu. 2024. Large Language Models for Data Annotation and Synthesis: A Survey. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*. 930–957.
- [70] Stephanie J Tobin and Rosanna E Guadagno. 2022. Why people listen: Motivations and outcomes of podcast listening. *Plos one* 17, 4 (2022), e0265806.
- [71] Petter Törnberg. 2024. Best practices for text annotation with large language models. *arXiv preprint arXiv:2402.05129* (2024).
- [72] Johanne R. Trippas, Damiano Spina, and Falk Scholer. 2025. Adapting Generative Information Retrieval Systems to Users, Tasks, and Scenarios. In *Information Access in the Era of Generative AI*, Ryen W. White and Chirag Shah (Eds.). Springer Nature Switzerland, Cham, 73–109. doi:10.1007/978-3-031-73147-1_4
- [73] Johanne R Trippas, Damiano Spina, Paul Thomas, Mark Sanderson, Hideo Joho, and Lawrence Cavedon. 2020. Towards a model for spoken conversational search. *Information Processing & Management* 57, 2 (2020), 102162.
- [74] Robert Vidigal. 2025. Measuring Belief Certainty in Political Knowledge. *Political Behavior* 47, 2 (2025), 529–551.
- [75] Ben Wang and Jiquan Liu. 2024. Cognitively Biased Users Interacting with Algorithmically Biased Results in Whole-Session Search on Debated Topics. In *Proceedings of the 2024 ACM SIGIR International Conference on Theory of Information Retrieval*. 227–237.
- [76] Junjie Wang. 2025. *How do AI-generated Podcasts Influence the Opinions of Users on Debated Topics? A User-Centered Exploration*. Master's thesis. Delft University of Technology, Delft, Netherlands. <https://resolver.tudelft.nl/uuid:d1e53949-eada-43ec-9e37-bc063263fb8c>. Accessed: 2026-01-08.
- [77] 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 2024 CHI Conference on Human Factors in Computing Systems*. 1–21.
- [78] Dennis Whitcomb, Heather Battaly, Jason Baehr, and Daniel Howard-Snyder. 2017. Intellectual humility. *Philosophy and Phenomenological Research* 94, 3 (2017), 509–539.
- [79] Ryen W. White. 2024. Advancing the Search Frontier with AI Agents. *Commun. ACM* 67, 9 (Aug. 2024), 54–65. doi:10.1145/3655615
- [80] Ryen W White and Eric Horvitz. 2015. Belief dynamics and biases in web search. *ACM Transactions on Information Systems (TOIS)* 33, 4 (2015), 1–46.
- [81] Ryen W White and Chirag Shah. 2025. *Information Access in the Era of Generative AI*. Springer.
- [82] Ryen W White and Chirag Shah. 2025. Panmodal information interaction. *Commun. ACM* 68, 4 (2025), 33–36.
- [83] Wan-Ching Wu, Diane Kelly, and Avneesh Sud. 2014. Using information scent and need for cognition to understand online search behavior. In *Proceedings of the 37th international ACM SIGIR conference on Research & development in information retrieval*. 557–566.
- [84] Luyan Xu, Mengdie Zhuang, and Ujwal Gadiraju. 2021. How do user opinions influence their interaction with web search results?. In *Proceedings of the 29th ACM Conference on User Modeling, Adaptation and Personalization*. 240–244.
- [85] Yuchi Yahagi, Rintaro Chujo, Yuga Harada, Changyo Han, Kohei Sugiyama, and Takeshi Naemura. 2024. PaperWave: Listening to Research Papers as Conversational Podcasts Scripted by LLM. *arXiv preprint arXiv:2410.15023* (2024).
- [86] Yuyu Yang, Kelsey Urgo, Jaime Arguello, and Robert Capra. 2025. Search+ Chat: Integrating Search and GenAI to Support Users with Learning-oriented Search Tasks. In *Proceedings of the 2025 ACM SIGIR Conference on Human Information Interaction and Retrieval*. 57–70.
- [87] Tao Zhou and Songtao Li. 2024. Understanding User Switch of Information Seeking: From Search Engines to Generative AI. *Journal of Librarianship and Information Science* (April 2024), 09610006241244800. doi:10.1177/09610006241244800
- [88] Mengdie Zhuang, Gianluca Demartini, and Elaine G Toms. 2017. Understanding engagement through search behaviour. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*. 1957–1966.