Information Research - Vol. 31 No. iConf (2026)

Disability misinformation on Facebook: a comparison of LLM-based fact-checking tools

Ian Prazak, Leah Padovani, Yool Lim, Julia H-P. Hsu, and Myeong Lee

DOI: https://doi.org/10.47989/ir31iConf64259

Abstract

Introduction. Social media has become a prominent space for seeking and sharing information, but it also enables misinformation to spread. When it comes to disability-related information, such as how to apply for a Medicaid waiver, understanding the prevalence of false information on social media becomes further complicated due to varying content types. To provide an initial exploration of this problem space, we investigated misinformation propensity within disability-related Facebook groups, group factors associated with it, and the performance of AI fact-checking tools in detecting this type of information.

Method. We identified target Facebook groups through a large-scale survey. From 20 public Facebook groups mentioned in the survey, we scraped 1,407 informational and fact-checkable posts. GPT-4o, GPTo1, and Originality.ai were used to classify the posts and compared.

Analysis. The results were validated against the ground-truths generated manually, providing benchmarks for assessing AI tools in detecting misinformation on Facebook.

Results. Our findings reveal that groups centered on developmental disabilities tend to be more vulnerable to misinformation. AI factchecking tools are generally effective in classifying accurate information but presented varying performance in detecting misinformation.

Conclusion. This work provides an initial assessment of the prevalence of misinformation about disability services and the performance of LLM-based tools.

Introduction

Social media platforms, such as Facebook, are frequently used as information sources for people's everyday life information seeking (Pentina and Tarafdar, 2014). The abundance of information not only makes it easy for people to access key information, but also increases the likelihood that they will encounter misinformation (Suarez-Lledo and Alvarez-Galvez, 2021). Individuals who rely on such information may unknowingly act on false content, which can have negative consequences on their health and well-being (Simko et al., 2019; Wang et al., 2019). This issue is especially troubling within disability-related communities, where misinformation can negatively affect individuals' understanding of their conditions, the services available to them, and effective treatments. Moreover, people with cognitive difficulties are particularly vulnerable to misinformation (Morrison et al., 2019), underscoring the importance of addressing disability-related misinformation.

However, finding disability-related information on social media and assessing its quality is challenging because disability policies evolve rapidly and social media postings tend to be nuanced and interwoven with personal experiences (Lee et al., 2024). Furthermore, there is limited understanding of the information sources that are popular among people with disabilities and their caregivers, which makes it difficult to assess the quality of disability-information quality. To address this gap, this study aims to identify key information sources on Facebook, the most widely used social media platform among people with disabilities and their caregivers. In addition, we examine the prevalence of misinformation across social media platforms and the key factors associated with susceptibility to misinformation. Specifically, we ask the following questions:

- **RQ1.** What kinds of misinformation exist in Facebook groups focused on people with disabilities?
- **RQ2.** Can AI tools identify misinformation about disabilities from Facebook?
- **RQ3.** How is group formality (or other key group characteristics) associated with misinformation propensity in Facebook groups?

Related work

Misinformation in the health context

Misinformation has been widely studied for its impact on publichealth perception and decisionmaking. Marginalised communities, such as individuals with low socioeconomic status or disabilities, were particularly vulnerable to it (Chi et al., 2020; Gibson and Martin III, 2019; Nicol et al., 2022). During the COVID–19 pandemic, for example, social media served as a primary source of misinformation for underrepresented communities, amplifying conspiracy theories (Sanfilippo et al., 2025) and allowing nonexperts to circulate unverified medical claims (Liu and Regulagedda, 2023). Although experts and authorities have attempted to counter these falsehoods with accurate information (Ghosh and Mitra, 2023; Song et al., 2024), the rapid pace at which misinformation spreads on social media has often outstripped their efforts (Amriza et al., 2025).

Content types and community dynamics both contributed to the spread of misinformation. Prior research showed that emotionally charged posts, especially those with negative sentiment, facilitate its dissemination (Conroy et al., 2015; Giachanou et al., 2021; Karami et al., 2021). In addition, the structural features of online communities and the kinds of contents shared within them also influence how misinformation circulates (Huvila and Gorichanaz, 2025).

AI-based fact-checkers

AI-driven fact-checking methods have emerged as scalable solutions for identifying and addressing online misinformation (Ghosh and Shah, 2018; Jahanbakhsh et al., 2023; Singh et al., 2021), whereas conventional approaches are often timeconsuming to implement by practitioners

(Burel and Alani, 2023; Gangopadhyay et al., 2024; Hasanain and Elsayed, 2022). Machine learning models, for instance, detected misinformation by leveraging social network topological features (Pierri et al., 2020) and linguistic cues (Ahmad et al., 2020). Although AI-based systems can inherit biases from their training data and may struggle to capture context or subtle nuances (Labib et al., 2022; Li and Sinnamon, 2024), the rapid progress of large language models and the fact-checking tools built on them offers a promising path to content-level, large-scale detection of misinformation.

However, it remains unclear how these models perform in detecting disabilityrelated misinformation, the primary channels through which such misinformation spreads, and the specific types of information involved. To address these gaps, we conduct a study that (1) identifies the main channels through which disabilityrelated (mis)information is disseminated, (2) analYSes its types and prevalence, and (3) evaluates the performance of widely used fact-checking tools in detecting this content. We focus on Facebook Groups and Pages, as they are among the most frequently used information sources for people with disabilities and their family members, according to our survey results.

Methods

Data

We identified the primary sources of disability-related information through a survey of 886 residents in an eastern U.S. state that examined how participants sought disability-related services. The survey was reviewed and approved by the authors' institutional review board (IRB) and was conducted in May 2024. Participants were recruited through a professional survey firm using stratified random sampling to maximise geographic and demographic representativeness across the state. The survey population consisted of people with disabilities and their caregivers. Participants were asked, in an open-ended format, to report the information sources they had previously used to obtain disability-related service information. The survey required participants to provide at least three and up to five information sources.

We then analYSed these responses descriptively by standardising the entries, as participants' answers contained typos and inconsistent terminology referring to similar information sources (Hsu and Lee, 2025). Facebook emerged as the most frequently cited source (39.7% of respondents), surpassing Google, government agencies, and friends. Some respondents mentioned specific Facebook Groups or Pages. Out of the survey responses, we manually identified 20 public Facebook Groups and Pages, as they were specifically mentioned by the respondents. We did not include other groups or pages that were not mentioned in the survey responses. We scraped all their posts, yielding roughly 21,400 entries.

Each Facebook post was evaluated using an OpenAI API call (GPT-4o) to determine whether it contained fact-checkable information or not (https://openai.com/api/). For example, instructions on how to apply for Medicaid services were included as fact-checkable information, while personal feelings about disability services were not included in our analysis. Posts that were not disability-related were also excluded. In total, 1,407 posts met these criteria and were subjected to further analysis.

Fact-checking with AI and ground truth generation

We used three fact-checking tools for the filtered posts, GPT-4o, GPT-o1, and Originality.ai (https://originality.ai/), and generated ground-truths through manual verification. Where applicable, each tool was queried with identical prompts. Posts were classified into three categories: *true*, *false*, or *not fact-checkable* (CBPC).

Data analysis

We conducted statistical analyses to compare the performances of GPT-4o, GPTo1, and Originality.ai against manual fact-checking. First, ANOVA was used to test for differences in classification accuracy across the AI-based tools. Next, Ordinary Least Squares (OLS) regression analysis was applied to examine how group characteristics related to the tendency for misinformation. Finally, Spearman's rank correlation coefficient (Spearman, 1961) was used to assess the consistency of rankings produced by the different methods. Group characteristics considered in the analysis include: whether the entity was a Facebook Group or Page, group size, the primary disability type the group focuses on (e.g., developmental, psychological, sensory, physical, or other), whether the group is managed by an official organisation outside of Facebook, whether the group holds offline or online meetings, and the proportion of comments within the group exhibiting positive or negative sentiment. These variables were generated manually through a qualitative coding of each Group or Page. Collectively, these tests provide an initial assessment of misinformation prevalence in Facebook Groups and Pages and the relative effectiveness of each fact-checking approach for disability-related content.

Results

RQ1: Types of (mis)information across Facebook groups focused on disabilities

Based on a qualitative review of the sample posts, we divided content into three categories: (1) promotions of disability-related products, events, or services (1,023 posts); (2) personal experiences of disability (189 posts); and (3) scientific discussions of disability (195posts). Figure 1 and Table 2 present the distribution of (mis)information across these categories. Overall, promotional content dominated the Facebook Groups and Pages. To determine information accuracy, we manually verified posts against the authoritative websites and articles, such as government websites (e.g., CDC) and academic articles. While most promotional posts were accurate, about 15% of the promotional posts contained misinformation, with the proportion varying across individual Groups and Pages.

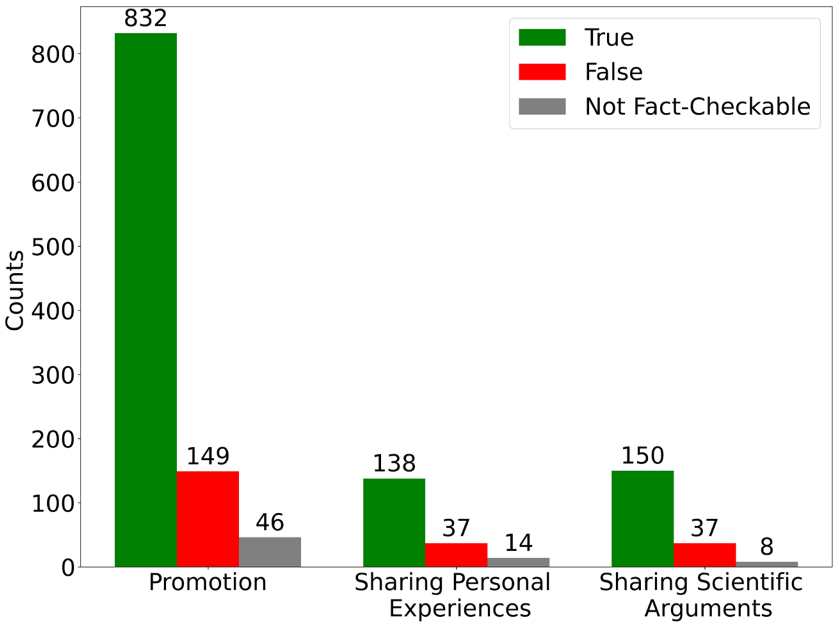

Figure 1. Distribution of (mis)information across post types.

Classification	True	False	CBFC
Promotion	0.808	0.147	0.045
Sharing Personal Experiences	0.735	0.191	0.074
Sharing Scientific Arguments	0.769	0.189	0.041

Table 1. Proportions of information types and their accuracy.

RQ2: AI performance in identifying misinformation

To evaluate the reliability of different fact-checking tools (Originality.ai, GPT-4o, GPT-o1) in comparison to manual fact-checking, a boxplot was generated to visualise the proportions across true, false, and not-fact-checkable (CBFC) categories. Each point in the scatter plot represents the proportion of a given category within a Facebook Group. The boxplot further summarises the distribution of these proportions by highlighting variations and potential outliers, as well as displaying the median proportion for each category. To statistically assess the differences between categories, pairwise t-tests were conducted, where the t-statistics and corresponding p-values were calculated to determine whether the differences between category proportions were statistically significant. A non-significant result indicates that the classifications produced by the two methods are similar.

Figure 2a shows that GPT-o1 performed comparably to manual fact-checking (i.e., approximate ground truth) for both true and false categories, with no statistically significant differences observed (true: $p = .4148$, $t = 0.82$; false: $p = .1304$, $t = -1.81$). GPT-4o likewise performed comparably to manual fact-checking in the false category ($p = .7371$, $t = -0.34$). Figure 2b presents comparisons among fact-checking tools within the advertisement category. In the *true* category, GPT-o1 did not differ significantly from manual fact-checking ($p = .2912$, $t = 1.07$). In the false category, both GPT-4o and GPT-o1 also showed no significant differences relative to manual fact-checking (GPT-4o: $p = .4755$, $t = 0.72$; GPT-o1: $p = .2062$, $t = -1.31$). Overall, for Facebook posts classified as advertisements, GPT-o1 most closely aligned with manual fact-checking in the true category, while both GPT-4o and GPT-o1 showed comparable performance in the false category. In contrast, Originality.ai demonstrated substantially poorer performance across categories.

Finally, when comparing two fact-checking methods with information classified as true, all methods exhibited Spearman's rank correlation coefficients greater than 0.971, while all other pair combinations (false, not-fact-checkable) were less than 0.8502. This suggests that fact-checking tools are more reliable when evaluating true information compared to false or ambiguous information.

RQ3: Exploring factors affecting misinformation propensity

Ordinary Least Squares regression analysis was conducted to examine the relationship between misinformation propensity and various group characteristics. The variables used in the regression model are listed in the *Data Analysis* section. The lack of statistical significance across all regression models may be attributable to the small sample size ($N_{groups} = 20$). In contrast, the correlation analyses revealed a meaningful association between disability type, specifically *developmental disability*, and misinformation propensity (Figure 3). This suggests that misinformation may be more prevalent in information concerning developmental disabilities.

Figure 4 shows that the non-parametric regression analysis shows a statistically significant correlation between the focus of the group on developmental disabili-

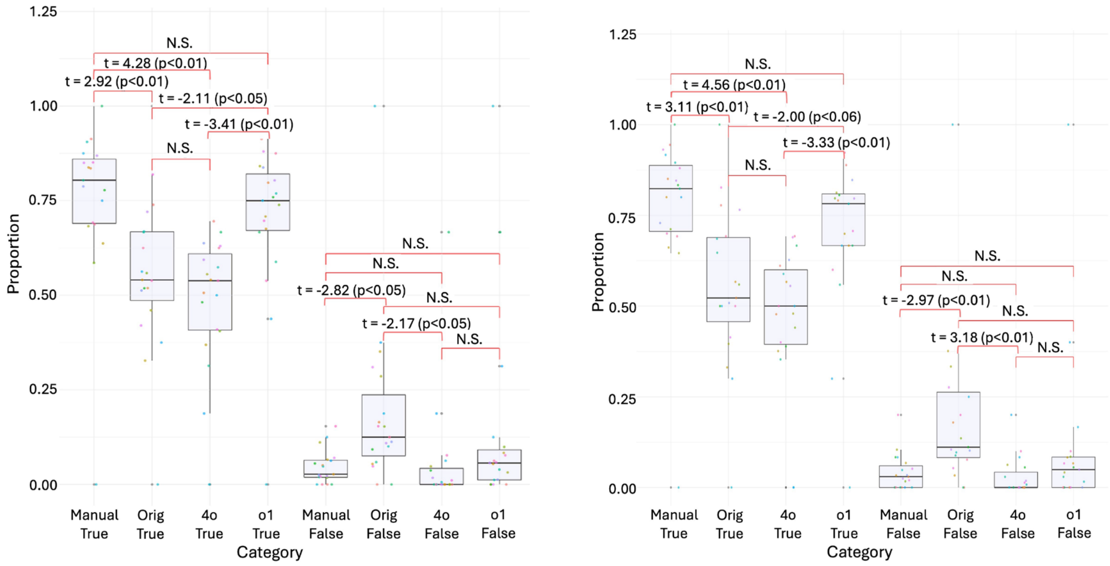

(a) Comparing performances of fact- checking tools for all posts using pairwise t-tests.

(b) Comparing performances of fact checking tools using pairwise t-tests for the advertisement-type postings.

Figure 2. Pair-wise t-tests for the fact-checking results across the tools.

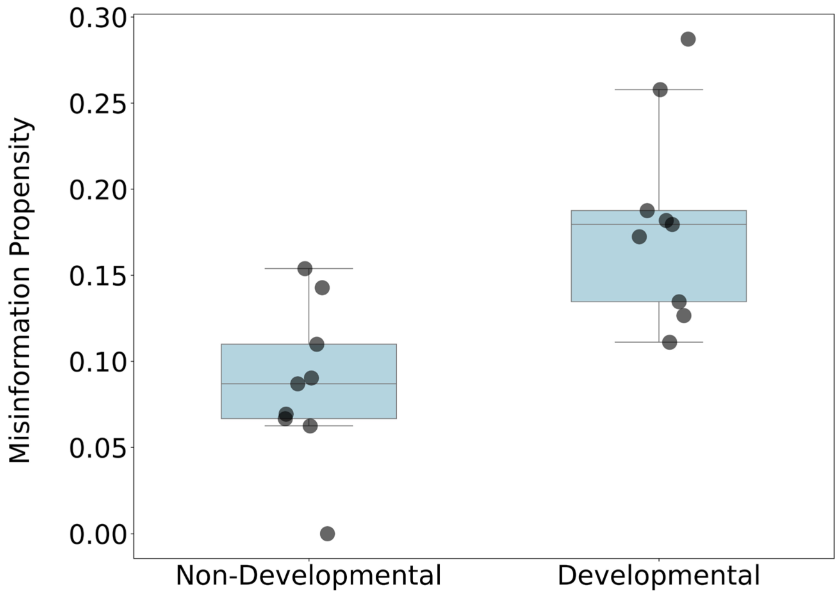

Figure 3. Misinformation propensity vs. whether a group focuses on developmental conditions.

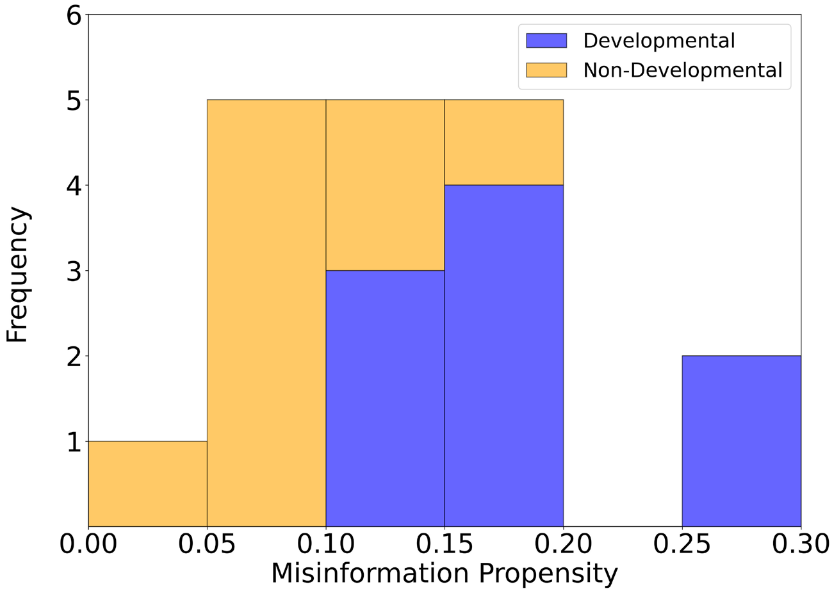

Figure 4. Histogram displaying the correlation of misinformation propensity and whether a group focuses on developmental conditions.

Linear	Kendall's Tau	Spearman's Correlation
$R^2 = 0.028$	$p < 0.005$	$p < 0.005$

Table 2. Statistical analysis results for developmental conditions vs. misinformation propensity

ties and the propensity for misinformation. Additionally, Table 2 highlights significant differences in their distributions, suggesting that groups centered on developmental conditions tend to exhibit higher levels of misinformation propensity.

Figure 5 presents a scatterplot that compares the overall positive sentiment of a group, defined as the proportion of posts classified as having a positive sentiment, with its propensity to misinformation. While prior research often cites sentiment as a factor influencing the presence of misinformation, our data in the context of disability service information on Facebook does not reproduce the findings.

Discussion

Group contexts and misinformation prevalence

Despite prior research identifying sentiment as one of the strongest indicators of misinformation, our analysis found no statistically significant association between sentiment and misinformation propensity. Instead, our results suggest that Facebook groups focused on developmental disabilities disseminate misinformation more frequently. Although this finding should be interpreted cautiously due to the small number of groups examined, our study indicates that both the type and prevalence of misinformation may vary substantially by disability category, not only as a function of population size, but also because of differences in care practices and information needs.

Developmental disabilities, particularly cognitive and learning disorders, are most prevalent among children (Young, 2022). As a result, parents and caregivers may be especially attentive to information related to these conditions, potentially increasing both the volume and circulation of misinformation within these communities. Further research is needed to clarify why and how certain disability types are associated with higher levels of misinformation on social media, particularly through the inclusion of a larger number of groups to improve statistical power.

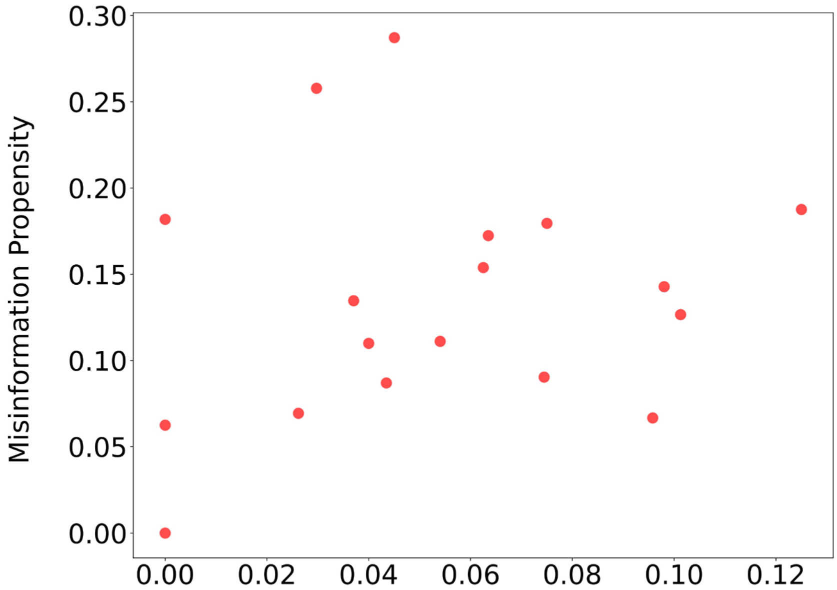

Figure 5. Negative sentiment distribution across groups.

Although our statistical power is constrained by the relatively small number of Groups and Pages analYSed, our findings offer unique insights and implications compared to those that rely on general samples of posts from platforms such as X or Facebook. The Groups and Pages included in this study were explicitly identified by people with disabilities and their caregivers through a large-scale, stratified sampling of participants, strengthening the quality and relevance of our sample. We limited our analysis to publicly accessible Facebook Groups and Pages to address ethical concerns related to privacy and consent.

Implications for AI-based fact-checking tools

Our results indicate that state-of-the-art large language models (LLMs) can reliably flag disability-related misinformation in social media posts. In many cases, the models' chain-of-thought reasoning appears to track the factual logic humans apply when verifying claims, allowing them to separate accurate advice or resource descriptions from misleading content (OpenAI, 2024). Yet, the gap between the ground-truth labels and the performance of GPT-o1 on false statements shows that current systems still fall short of practical, production-level accuracy, especially when it comes to catching subtle or highly specific misinformation. Misclassifications at this scale could mislead end users who rely on the tool's judgments for critical health or service decisions. Because popular search engines such as Google increasingly leverage LLMs to generate search results, search engine engineers and information retriever experts may need to account for the risks associated with LLMs when fine-tuning these models for search applications.

To achieve reliable performance, fact-checking tools may need to incorporate a domain-specific knowledge base that captures authoritative disability information and policy details. Constructing such a database, then plugging it into a retrievalaugmented generation (RAG) pipeline, would give LLMs direct access to vetted sources during inference, improving their reasoning and reducing error rates. In short, pairing advanced models with a curated disability-information repository might be essential in future studies for delivering trustworthy, scalable fact-checking.

Future work

While the study identifies meaningful patterns of misinformation within Facebook Groups and Pages related to disability services, its limited statistical power should be addressed in future research. Expanding the geographical scope of future surveys beyond a single U.S. state would allow for more accurate identification of key information sources. In addition, future work should examine a broader range of information sources referenced in the survey beyond Facebook to support a more comprehensive understanding of misinformation types and their diffusion patterns. Finally, analYSing these patterns at the individual level, by incorporating demographic information from both users and information providers, would enhance the study's nuance and statistical robustness.

Acknowledgment

This project was conducted in part through the Aspiring Scientists Summer Internship Program (ASSIP) at George Mason University, which aims to promote research opportunities for high school and undergraduate students. ASSIP interns Mayuko Karakawa (Stanford Online High School '25; currently an undergraduate student at Stanford University) and Angela Sidhu (Computer Science, George Mason University; currently an undergraduate student at Georgetown University) contributed to this study. The survey component of this research was supported by the Virginia Board for People with Disabilities (VBPD) (Grant #GMU-3-2023). We also thank Teri Morgan and Linh Nguyen from VPBD for their valuable feedback and guidance on the survey. Funding for this project supported, in part, by the Virginia Board for People with Disabilities, under grant number 2301VASCDD, from the U.S. Administration for Community Living (ACL), Department of Health and Human Services, Washington, D.C. 20201. Points of view or opinions do not, therefore, necessarily represent official ACL policy.

Generative AI use

We employed GPT-4o API, OpenAI o1 API, and Originality.ai for testing their performance, as detailed in the Methods section. The authors assume all responsibility for the content of this submission.

About the authors

Ian Prazak is an undergraduate student at Duke University, majoring in Computer Science and Economics. He was an ASSIP summer intern in 2024 at George Mason University while a high school senior at Keystone School.

Leah Padovani is an undergraduate student in the iSchool at the University of Maryland, College Park. She was an ASSIP summer intern in 2024 at George Mason University while a sophomore at the UMD iSchool.

Yool Lim is an undergraduate student at Rice University, majoring in Statistics. She was an ASSIP summer intern in 2024 at George Mason University while a high school senior at Hendrickson High School.

Julia (Hsin-Ping) Hsu is a PhD candidate in Information Technology with a concentration in Information Sciences and Technology at the School of Computing, George Mason University. She served as a co-mentor for ASSIP 2024. Her research interests include AI for social good, community informatics, and civic technology. She can be contacted at hhsu2@gmu.edu

Myeong Lee is an Assistant Professor in the Department of Information Sciences and Technology at George Mason University. He served as the mentor for the ASSIP 2024 student team. He leads the Community Informatics Lab (CIL) and serves as an AI-in-Gov Fellow at GMU. He is the corresponding author of this paper and can be contacted at mlee89@gmu.edu

References

Ahmad, I., Yousaf, M., Yousaf, S., & Ahmad, M. O. (2020). Fake news detection using machine learning ensemble methods. Complexity, 2020(1), 8885861.

Amriza, R. N. S., Chou, T.-C., & Ratnasari, W. (2025). Understanding the shifting nature of fake news research: Consumption, dissemination, and detection. Journal of the Association for Information Science and Technology, 76(6), 896– 916.

Bannon, L. (2011). Reimagining hci: Toward a more human-centered perspective.

Interactions, 18(4), 50–57. https://doi.org/10.1145/1978822.1978833

Burel, G., & Alani, H. (2023). The fact-checking observatory: Reporting the cospread of misinformation and fact-checks on social media. Proceedings of the 34th ACM Conference on Hypertext and Social Media. https://doi.org/ 10.1145/3603163.3609042

Chi, Y., He, D., & Jeng, W. (2020). Laypeople's source selection in online health information-seeking process. Journal of the Association for Information Science and Technology, 71(12), 1484–1499. https://doi.org/https://doi.org/ 10.1002/asi.24343

Conroy, N. K., Rubin, V. L., & Chen, Y. (2015). Automatic deception detection: Methods for finding fake news. Proceedings of the Association for Information Science and Technology, 52(1), 1–4. https://doi.org/https://doi.org/10.1002/ pra2.2015.145052010082

Gangopadhyay, S., Schellhammer, S., Hafid, S., Dessi, D., Koß, C., Todorov, K., Dietze, S., & Jabeen, H. (2024). Investigating characteristics, biases and evolution of fact-checked claims on the web. Proceedings of the 35th ACM Conference on Hypertext and Social Media, 246–258. https://doi.org/10.1145/ 3648188.3675135

Ghosh, S., & Mitra, P. (2023). Catching lies in the act: A framework for early misinformation detection on social media. Proceedings of the 34th ACM Conference on Hypertext and Social Media. https://doi.org/10.1145/3603163.3609057

Ghosh, S., & Shah, C. (2018). Towards automatic fake news classification. Proceedings of the Association for Information Science and Technology, 55(1), 805– 807. https://doi.org/https://doi.org/10.1002/pra2.2018.14505501125

Giachanou, A., Rosso, P., & Crestani, F. (2021). The impact of emotional signals on credibility assessment. Journal of the Association for Information Science and Technology, 72(9), 1117–1132. https://doi.org/https://doi.org/10.1002/asi. 24480

Gibson, A. N., & Martin III, J. D. (2019). Re-situating information poverty: Information marginalISATion and parents of individuals with disabilities. Journal of the Association for Information Science and Technology, 70(5), 476–487. https: //doi.org/https://doi.org/10.1002/asi.24128

Hasanain, M., & Elsayed, T. (2022). Studying effectiveness of web search for fact checking. Journal of the Association for Information Science and Technology,

73(5), 738–751. https://doi.org/https://doi.org/10.1002/asi.24577

Hsu, J. H.-P., & Lee, M. (2025). From open-ended text to taxonomy: An llm-based framework for information sources for disability services. Proceedings of the Association for Information Science and Technology, 62(1), 915–919.

Huvila, I., & Gorichanaz, T. (2025). Trends in information behavior research, 2016–2022: An annual review of information science and technology (arist) paper. Journal of the Association for

Information Science and Technology, 76(1), 216–237. https://doi.org/https://doi.org/10.1002/asi.24943

Jahanbakhsh, F., Katsis, Y., Wang, D., Popa, L., & Muller, M. (2023). Exploring the use of personalised ai for identifying misinformation on social media. Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems. https://doi.org/10.1145/3544548.3581219

Karami, M., Nazer, T. H., & Liu, H. (2021). Profiling fake news spreaders on social media through psychological and motivational factors. Proceedings of the 32nd ACM Conference on Hypertext and Social Media, 225–230. https://doi. org/10.1145/3465336.3475097

Labib, A., Chakhar, S., Hope, L., Shimell, J., & Malinowski, M. (2022). Analysis of noise and bias errors in intelligence information systems. Journal of the Association for Information Science and Technology, 73(12), 1755–1775. https://doi.

org/https://doi.org/10.1002/asi.24707

Lee, M., Abubakr, L., Shrivastava, T., Hsu, J. H.-p., Whitman, S. A., & Kim, P. (2024). 2024 assessment of virginia's information ecology of the disability services system. Virginia Board for People with Disabilities.

Li, A., & Sinnamon, L. (2024). Generative ai search engines as arbiters of public knowledge: An audit of bias and authority. Proceedings of the Association for Information Science and Technology, 61(1), 205–217. https://doi.org/https: //doi.org/10.1002/pra2.1021

Liu, J., & Regulagedda, R. M. (2023). Social network analysis of misinformation spreading and science communication during covid-19. Proceedings of the Association for Information Science and Technology, 60(1), 1059–1061. https://doi. org/https://doi.org/10.1002/pra2.944

Morrison, J., Forrester-Jones, R., Bradshaw, J., & Murphy, G. (2019). Communication and cross-examination in court for children and adults with intellectual disabilities: A systematic review. The International Journal of Evidence & Proof, 23(4), 366–398.

Nicol, E., Willson, R., Ruthven, I., Elsweiler, D., & Buchanan, G. (2022). Information intermediaries and information resilience: Working to support marginalised groups. Proceedings of the Association for Information Science and Technology, 59(1), 469–473. https://doi.org/https://doi.org/10.1002/pra2.654

OpenAI. (2024). Learning to reason with LLMs. Retrieved September 12, 2024, from https://openai.com/index/learning-to-reason-with-llms/

Pentina, I., & Tarafdar, M. (2014). From 'information' to 'knowing': Exploring the role of social media in contemporary news consumption. Computers in human behavior, 35, 211–223.

Pierri, F., Piccardi, C., & Ceri, S. (2020). Topology comparison of twitter diffusion networks effectively reveals misleading information. Scientific reports, 10(1), 1372.

Sanfilippo, M. R., Zhu, X. A., & Yang, S. (2025). Sociotechnical governance of misinformation: An annual review of information science and technology (arist) paper. Journal of the Association for Information Science and Technology, 76(1), 289–325. https://doi.org/https://doi.org/10.1002/asi.24953

Simko, J., Hanakova, M., Racsko, P., Tomlein, M., Moro, R., & Bielikova, M. (2019). Fake news reading on social media: An eye-tracking study. Proceedings of the 30th ACM Conference on Hypertext and Social Media, 221–230. https:// doi.org/10.1145/3342220.3343642

Singh, V. K., Ghosh, I., & Sonagara, D. (2021). Detecting fake news stories via multimodal analysis. Journal of the Association for Information Science and Technology, 72(1), 3–17. https://doi.org/https://doi.org/10.1002/asi.24359

Song, S., Ying, J., Zhao, Y. (, & Li, J. (2024). Harnessing the power of scientists and livestreaming: Health information outreach in a medical library. Proceedings of the Association for Information Science and Technology, 61(1), 645–

649. https://doi.org/https://doi.org/10.1002/pra2.1073

Spearman, C. (1961). The proof and measurement of association between two things.

Suarez-Lledo, V., & Alvarez-Galvez, J. (2021). Prevalence of health misinformation on social media: Systematic review. Journal of medical Internet research, 23(1), e17187.

Wang, Y., McKee, M., Torbica, A., & Stuckler, D. (2019). Systematic literature review on the spread of health-related misinformation on social media. Social Science & Medicine, 240, 112552. https://doi.org/https://doi.org/10.1016/j. socscimed.2019.112552

Young, N. A. (2022). Childhood disability in the united states: 2019. ACSBR-006.