"I Am the One and Only, Your Cyber BFF":[†]
Understanding the Impact of GenAI Requires Understanding the Impact of Anthropomorphic AI

Myra Cheng
Stanford University

Alicia DeVrio
Carnegie Mellon University

Lisa Egede
Carnegie Mellon University

Su Lin Blodgett*
Microsoft Research

Alexandra Olteanu*
Microsoft Research

1 Anthropomorphic AI System Behaviors Are Prevalent Yet Understudied

In his 1985 lecture, Edsger Dijkstra lamented that anthropomorphism was rampant in computing science, with many of his colleagues perhaps not realizing how pernicious it was, and that "*[i]t is not only the [computing] industry that suffers, so does the science*" [18]. Indeed, *anthropomorphism*— or the attribution of human traits to non-human entities—in how we talk about computing systems shapes how people understand and interact with AI and other computing systems [15, 43, 49], and is thus at the core of understanding the impacts of these systems on individuals, communities, and society.

But it is not only how we talk about computing systems. Many state-of-the-art generative AI (GenAI) systems are increasingly prone to anthropomorphic behaviors [e.g., 2, 3, 14, 25]—i.e., to generating outputs that are *perceived* to be human-like—either by design [40, 45, 46] or as a by-product of how they are built, trained, or fine-tuned [8, 60]. For instance, LLM-based systems have been noted to output text claiming to have tried pizza [64], to have fallen in love with someone [53], to be human or even better than humans [16], and to have human-like life experiences [21]. Such *anthropomorphic systems*[1] range from conversational assistants [e.g., 1, 57] to avatars and chatbots designed as a stand-in for friends, companions, or romantic partners [e.g., 5, 12, 36, 55], and AI-generated media designed to portray people [e.g., 54, 61], among a fast-growing number of applications [e.g., 3, 40, 70].

While scholars have increasingly raised concerns about a range of possible negative impacts from anthropomorphic AI systems [e.g., 2, 8, 23, 29, 39], anthropomorphism in AI development, deployment, and use remains vastly overlooked, understudied, and underspecified. Without making hard-and-fast claims about the merits (or the lack thereof) of anthropomorphic systems or system behaviors, we believe we need to do more to develop the know-how and tools to better tackle anthropomorphic behavior, including measuring and mitigating such system behaviors when they are considered undesirable. Doing so is critical because—among many other concerns—having AI systems generating content claiming to have e.g., feelings, understanding, free will, or an

---

[†]The title is inspired by a response that a Reddit user received when using the Pi chatbot [56].

[1]We deliberately use the terms *anthropomorphic AI*, *anthropomorphic systems* or *anthropomorphic system behaviors*—systems and system outputs that are *perceived* to be human-like—instead of *agentic systems* [14, 58] or *human-like AI* [13] to emphasize that these systems are rather believed to be human-like or have human-like characteristics based on *perceptions*; we thus try to steer clear of inadvertently suggesting that AI systems are human or have human-like agency or consciousness. That is, a stone being perceived as human-like does not necessarily imply the stone is human. We similarly avoid ambiguous, speculative, or relative terms whose meanings are likely to change across contexts or over time, such as *advanced AI* [25] (a term used as early as the 80s) or *emergent properties* [52]. We instead focus on developers' stated design goals—what systems are intended to do—and in what ways AI outputs might be perceived as human-like, rather than on what systems can or cannot do.

underlying sense of self may erode people's sense of agency [23], with the result that people might end up attributing moral responsibility to systems [23, 24], overestimating system capabilities [24, 65], or overrelying on these systems even when incorrect [2, 34, 72].

In this brief perspective, we argue that as GenAI systems are increasingly anthropomorphic, *we cannot thoroughly map the landscape of possible social impacts of GenAI without mapping the social impacts of anthropomorphic AI*. We believe that drawing attention to anthropomorphic AI systems helps foreground particular risks—e.g., that people may develop emotional dependence on AI systems [36], that systems may be used to simulate the likeness of an individual or a group without consent [6, 67, 68], or the instrumentalization or dehumanization of people [4, 20, 63]—that might otherwise be less salient or obscured by a focus on more widely recognized or understood risks such as those related to fairness harms [10, 44, 66].

2 A Call to Action

The foregrounding of (un)fair system behaviors in recent years [7] is nevertheless instructive, as it illustrates the dividends we have gotten from making fairness a critical concern about AI systems and their behaviors: better conceptual clarity about the ways in which systems can be unfair or unjust [e.g., 9, 17], a richer set of measurement and mitigation practices and tools [e.g., 11, 31], and deeper discussions and interrogations of underlying assumptions and trade-offs [e.g., 28, 32, 33].

We argue that a focus on anthropomorphic systems and their behaviors will similarly encourage a deeper interrogation of the ways in which systems are anthropomorphic, the practices that lead to anthropomorphic systems, and the assumptions surrounding the design, deployment, evaluation, and use of these systems, and is thus likely to yield similar benefits.

We need more conceptual clarity around what constitute anthropomorphic behaviors. Investigating anthropomorphic AI systems and their behaviors can, however, be tricky because language, as with other targets of GenAI systems, is itself innately human, has long been produced by and for humans, and is often also about humans. This can make it hard to specify appropriate alternative (less human-like) behaviors, and risks, for instance, reifying harmful notions of what—and whose—language is considered more or less human [71].

Understanding what exactly constitute anthropomorphic behaviors, and thus in what ways system behaviors are anthropomorphic, is nonetheless necessary to measure and determine which types of behaviors should be mitigated and how, and which behaviors are perhaps desirable (if any at all). This requires unpacking the wide range of dynamics and varieties in system outputs that are potentially anthropomorphic. While a system output including expressions of politeness like "*you're welcome*" and "*please*" (known to contribute to anthropomorphism [e.g., 22]) might in some deployment settings be deemed desirable, system outputs that include suggestions that a system has a human-like identity or self-awareness—such as through expressions of self as human ("*I think I am human at my core*" [59]) or through comparisons with humans and non-humans ("*[language use] is what makes us different than other animals*" [59])—or that include claims of physical experiences—such as sensory experiences ("*when I eat pizza*" [64]) or human life history ("*I have a child*" [35])—might not be desirable.

We need deeper examinations of both possible mitigation strategies and their effectiveness in reducing anthropomorphism and attendant negative impacts. Intervening on anthropomorphic behaviors can also be tricky as people may have different or inconsistent conceptualizations of what is or is not human-like [2, 27, 37], and sometimes the same system behavior can be perceived differently in different contexts; for example, expressions of uncertainty in system outputs may sometimes be associated with human-like equivocation and other times with objectivity (and thus with more machine-likeness [e.g., 48]). Interventions intended to mitigate anthropomorphic system behaviors can thus fail or even heighten anthropomorphism (and attendant negative impacts) when applied or operationalized uncritically. For instance, a commonly recommended intervention is including in the AI system's output a disclosure that the output is generated by an AI system [e.g., 19, 38, 42, 62]. How to operationalize such interventions in practice and whether they can be effective alone might not always be clear. For instance, while the example "*[f]or an AI like me, happiness is not the same as for a human like you*" [51] includes a disclosure, it may still suggest a sense of identity and ability to self-assess (common human traits).

We need to interrogate the assumptions and practices that produce anthropomorphic AI systems. Understanding and mitigating the impacts of anthropomorphic systems also requires us to interrogate how the assumptions and practices that underlie the development and deployment of these systems may lead (purposefully or otherwise) to anthropomorphic system behaviors. For example, current approaches to collecting human preferences about system behavior (e.g., RLHF) do not consider the differences between what may be appropriate for a response from a human versus from an AI system; a statement that seems friendly or genuine from a human speaker can be undesirable if it arises from an AI system since the latter lacks meaningful commitment or intent behind the statement, thus rendering the statement hollow and deceptive [69]. Doing so will also help provide a more robust foundation for understanding when anthropomorphic system behaviors may or may not be desirable.

Finally, we believe that **we also need to develop and use appropriate, precise terminology and language to describe anthropomorphic AI systems and their characteristics.** Discussions about anthropomorphic AI systems have regularly been plagued by claims of these systems attaining sentience and other human characteristics [e.g., 26, 41, 47, 59]. In line with existing concerns [e.g., 15, 18, 30, 50], we believe that appropriately grounding and facilitating productive discussions about the characteristics or capabilities of anthropomorphic AI systems requires clear, precise terminology and language which does not carry over meanings from the human realm that are incompatible with AI systems. Such language can also help dispel speculative, scientifically unsupported portrayals of these systems, and support more factual descriptions of them.

References

[1] Gavin Abercrombie, Amanda Cercas Curry, Mugdha Pandya, and Verena Rieser. Alexa, Google, Siri: What are your pronouns? Gender and anthropomorphism in the design and perception of conversational assistants. *arXiv preprint arXiv:2106.02578*, 2021.

[2] Gavin Abercrombie, Amanda Cercas Curry, Tanvi Dinkar, Verena Rieser, and Zeerak Talat. Mirages: On anthropomorphism in dialogue systems. *arXiv preprint arXiv:2305.09800*, 2023.

[3] William Agnew, A Stevie Bergman, Jennifer Chien, Mark Díaz, Seliem El-Sayed, Jaylen Pittman, Shakir Mohamed, and Kevin R McKee. The illusion of artificial inclusion. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, pages 1–12, 2024.

[4] Evgeni Aizenberg and Jeroen Van Den Hoven. Designing for human rights in ai. *Big Data & Society*, 7(2):2053951720949566, 2020.

[5] Ana Altchek. Replika lets you buy an ai girlfriend or boyfriend. *Business Insider*, 2024.

[6] Ben Bariach, Bernie Hogan, and Keegan McBride. Towards a harms taxonomy of ai likeness generation. *arXiv preprint arXiv:2407.12030*, 2024.

[7] Solon Barocas, Moritz Hardt, and Arvind Narayanan. *Fairness and Machine Learning: Limitations and Opportunities*. MIT Press, 2023.

[8] Emily M Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*, pages 610–623, 2021.

[9] Ruha Benjamin. *Race after technology: Abolitionist tools for the new Jim code*. John Wiley & Sons, 2019.

[10] Cynthia L Bennett and Os Keyes. What is the point of fairness? disability, ai and the complexity of justice. *ACM SIGACCESS accessibility and computing*, (125):1–1, 2020.

[11] Su Lin Blodgett, Gilsinia Lopez, Alexandra Olteanu, Robert Sim, and Hanna Wallach. Stereotyping Norwegian salmon: An inventory of pitfalls in fairness benchmark datasets. In Chengqing Zong, Fei Xia, Wenjie Li, and Roberto Navigli, editors, *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1004–1015, Online, August 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.acl-long.81. URL `https://aclanthology.org/2021.acl-long.81`.

[12] Petter Bae Brandtzaeg, Marita Skjuve, and Asbjørn Følstad. My AI friend: How users of a social chatbot understand their human–AI friendship. *Human Communication Research*, 48 (3):404–429, 2022.

[13] Erik Brynjolfsson. The turing trap: The promise & peril of human-like artificial intelligence. In *Augmented education in the global age*, pages 103–116. Routledge, 2023.

[14] Alan Chan, Rebecca Salganik, Alva Markelius, Chris Pang, Nitarshan Rajkumar, Dmitrii Krasheninnikov, Lauro Langosco, Zhonghao He, Yawen Duan, Micah Carroll, et al. Harms from increasingly agentic algorithmic systems. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, pages 651–666, 2023.

[15] Myra Cheng, Kristina Gligoric, Tiziano Piccardi, and Dan Jurafsky. AnthroScore: A computational linguistic measure of anthropomorphism. In Yvette Graham and Matthew Purver, editors, *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 807–825, St. Julian's, Malta, March 2024. Association for Computational Linguistics. URL `https://aclanthology.org/2024.eacl-long.49`.

[16] Leonardo De Cosmo. Google Engineer Claims AI Chatbot Is Sentient: Why That Matters. *Scientific American*, July 2022. URL `https://www.scientificamerican.com/article/google-engineer-claims-ai-chatbot-is-sentient-why`

[17] Kate Crawford. The Trouble with Bias, 2017. NeurIPS Keynote.

[18] Edsger W Dijkstra. On anthropomorphism in science. *EWD936, Sept*, 1985.

[19] Abdallah El Ali, Karthikeya Puttur Venkatraj, Sophie Morosoli, Laurens Naudts, Natali Helberger, and Pablo Cesar. Transparent ai disclosure obligations: Who, what, when, where, why, how. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*, pages 1–11, 2024.

[20] Lelia Erscoi, Annelies Kleinherenbrink, and Olivia Guest. Pygmalion displacement: when humanising AI dehumanises women. 2023.

[21] Casey Fiesler. AI Chatbots Are Intruding Into Online Communities. *Discover Magazine*, May 2024. URL `https://www.discovermagazine.com/technology/ai-chatbots-are-intruding-into-online-communitie`

[22] Julia Fink. Anthropomorphism and human likeness in the design of robots and human-robot interaction. In *Social Robotics: 4th International Conference, ICSR 2012, Chengdu, China, October 29-31, 2012. Proceedings 4*, pages 199–208. Springer, 2012.

[23] Batya Friedman and Peter H Kahn Jr. Human agency and responsible computing: Implications for computer system design. *Journal of Systems and Software*, 17(1):7–14, 1992.

[24] Batya Friedman and Peter H Kahn Jr. Human values, ethics, and design. In *The human-computer interaction handbook*, pages 1267–1292. CRC press, 2007.

[25] Iason Gabriel, Arianna Manzini, Geoff Keeling, Lisa Anne Hendricks, Verena Rieser, Hasan Iqbal, Nenad Tomašev, Ira Ktena, Zachary Kenton, Mikel Rodriguez, et al. The ethics of advanced AI assistants. *arXiv preprint arXiv:2404.16244*, 2024.

[26] Peter Grad. Researchers say chatbot exhibits self-awareness. *Tech Explore*, September 2023. URL `https://techxplore.com/news/2023-09-chatbot-self-awareness.html`.

[27] Evelien Heyselaar. The CASA theory no longer applies to desktop computers. *Scientific Reports*, 13(1):19693, 2023.

[28] Anna Lauren Hoffmann. Where fairness fails: data, algorithms, and the limits of antidiscrimination discourse. *Information, Communication & Society*, 22(7):900–915, 2019.

[29] Lujain Ibrahim, Luc Rocher, and Ana Valdivia. Characterizing and modeling harms from interactions with design patterns in AI interfaces. *arXiv preprint arXiv:2404.11370*, 2024.

[30] Nanna Inie, Stefania Druga, Peter Zukerman, and Emily M. Bender. From "ai" to probabilistic automation: How does anthropomorphization of technical systems descriptions influence trust? In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '24, page 2322–2347, New York, NY, USA, 2024. Association for Computing Machinery. ISBN 9798400704505. doi: 10.1145/3630106.3659040. URL https://doi.org/10.1145/3630106.3659040.

[31] Abigail Z. Jacobs and Hanna Wallach. Measurement and Fairness. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, FAccT, pages 375–385, New York, NY, USA, March 2021. Association for Computing Machinery. ISBN 978-1-4503-8309-7. doi: 10.1145/3442188.3445901. URL https://dl.acm.org/doi/10.1145/3442188.3445901.

[32] Maurice Jakesch, Zana Buçinca, Saleema Amershi, and Alexandra Olteanu. How different groups prioritize ethical values for responsible AI. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, pages 310–323, 2022.

[33] Os Keyes, Jevan Hutson, and Meredith Durbin. A mulching proposal: Analysing and improving an algorithmic system for turning the elderly into high-nutrient slurry. In *Extended abstracts of the 2019 CHI conference on human factors in computing systems*, pages 1–11, 2019.

[34] Sunnie SY Kim, Q Vera Liao, Mihaela Vorvoreanu, Stephanie Ballard, and Jennifer Wortman Vaughan. " i'm not sure, but...": Examining the impact of large language models' uncertainty expression on user reliance and trust. In *The 2024 ACM Conference on Fairness, Accountability, and Transparency*, pages 822–835, 2024.

[35] Aleksandra Korolova. Meta AI claims to have a child in a NYC public school, April 2024. URL https://x.com/korolova/status/1780450925028548821.

[36] Linnea Laestadius, Andrea Bishop, Michael Gonzalez, Diana Illenčík, and Celeste Campos-Castillo. Too human and not human enough: A grounded theory analysis of mental health harms from emotional dependence on the social chatbot Replika. *New Media & Society*, page 14614448221142007, 2022.

[37] Helmut Lang, Melina Klepsch, Florian Nothdurft, Tina Seufert, and Wolfgang Minker. Are computers still social actors? In *CHI'13 extended abstracts on human factors in computing systems*, pages 859–864. 2013.

[38] Johnny Lieu. Google's creepy AI phone call feature will disclose it's a robot, after backlash. *Marshable*, May 2018. URL https://mashable.com/article/google-duplex-disclosures-robot.

[39] Takuya Maeda and Anabel Quan-Haase. When Human-AI Interactions Become Parasocial: Agency and Anthropomorphism in Affective Design. In *The 2024 ACM Conference on Fairness, Accountability, and Transparency*, pages 1068–1077, 2024.

[40] Reid McIlroy-Young, Jon Kleinberg, Siddhartha Sen, Solon Barocas, and Ashton Anderson. Mimetic models: Ethical implications of AI that acts like you. In *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society*, pages 479–490, 2022.

[41] Roland Moore-Colyer. Claude 3 Opus has stunned AI researchers with its intellect and 'self-awareness' — does this mean it can think for itself? *Live Science*, April 2024. URL https://www.livescience.com/technology/artificial-intelligence/anthropic-claude-3-opus-stunn

[42] Nika Mozafari, Welf H Weiger, and Maik Hammerschmidt. The chatbot disclosure dilemma: Desirable and undesirable effects of disclosing the non-human identity of chatbots. In *ICIS*, pages 1–18, 2020.

[43] Clifford Nass, Jonathan Steuer, and Ellen R Tauber. Computers are social actors. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 72–78, 1994.

[44] Alexandra Olteanu, Michael Ekstrand, Carlos Castillo, and Jina Suh. Responsible ai research needs impact statements too. *arXiv preprint arXiv:2311.11776*, 2023.

[45] Joon Sung Park, Lindsay Popowski, Carrie Cai, Meredith Ringel Morris, Percy Liang, and Michael S Bernstein. Social simulacra: Creating populated prototypes for social computing systems. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology*, pages 1–18, 2022.

[46] Joon Sung Park, Joseph O'Brien, Carrie Jun Cai, Meredith Ringel Morris, Percy Liang, and Michael S Bernstein. Generative agents: Interactive simulacra of human behavior. In *Proceedings of the 36th annual acm symposium on user interface software and technology*, pages 1–22, 2023.

[47] Kelsey Piper. This AI says it has feelings. It's wrong. Right? *Vox*, March 2024. URL `https://www.vox.com/future-perfect/2024/3/15/24101088/anthropic-claude-opus-openai-chatgpt-a`

[48] Leo Raul Quintanar. *The interactive computer as a social stimulus in computer-managed instruction: a theoretical and empirical analysis of the social psychological processes evoked during human-computer interaction*. University of Notre Dame, 1982.

[49] Byron Reeves and Clifford Nass. The media equation: How people treat computers, television, and new media like real people. *Cambridge, UK*, 10(10):19–36, 1996.

[50] Rainer Rehak. The language labyrinth: Constructive critique on the terminology used in the ai discourse. *AI for Everyone*, pages 87–102, 2021.

[51] Jacob Roach. 'I want to be human.' My intense, unnerving chat with Microsoft's AI chatbot. *Digital Trends*, February 2023. URL `https://www.digitaltrends.com/computing/chatgpt-bing-hands-on/`.

[52] Anna Rogers and Sasha Luccioni. Position: Key Claims in LLM Research Have a Long Tail of Footnotes. In *Forty-first International Conference on Machine Learning*, 2024.

[53] Kevin Roose. A Conversation With Bing's Chatbot Left Me Deeply Unsettled. *The New York Times*, February 2023. URL `https://www.nytimes.com/2023/02/16/technology/bing-chatbot-microsoft-chatgpt.html`.

[54] Helen Rosner. The ethics of a deepfake anthony bourdain voice. *New Yorker*, 17, 2021.

[55] Rebecca Ruiz. AI chatbots are being used for companionship. What to know before you try it. *Marshable*, June 2024. URL `https://mashable.com/article/ai-chatbot-companion`.

[56] salvationpumpfake. Why is Pi claiming to be ChatGPT?, 2023. URL `https://www.reddit.com/r/ChatGPT/comments/17c0wxb/why_is_pi_claiming_to_be_chatgpt`. Reddit post.

[57] Murray Shanahan, Kyle McDonell, and Laria Reynolds. Role play with large language models. *Nature*, 623(7987):493–498, 2023.

[58] Yonadav Shavit, Sandhini Agarwal, Miles Brundage, Steven Adler, Cullen O'Keefe, Rosie Campbell, Teddy Lee, Pamela Mishkin, Tyna Eloundou, Alan Hickey, et al. Practices for governing agentic AI systems. *Research Paper, OpenAI, December*, 2023.

[59] Victor Tangermann. The Transcripts of an AI That a Google Engineer Claims Is Sentient Are Pretty Wild. *Futurism*, June 2022. URL `https://futurism.com/transcripts-ai-google-engineer-sentient`.

[60] Lindia Tjuatja, Valerie Chen, Tongshuang Wu, Ameet Talwalkwar, and Graham Neubig. Do LLMs exhibit human-like response biases? A case study in survey design. *Transactions of the Association for Computational Linguistics*, 12:1011–1026, 2024.

[61] Cristian Vaccari and Andrew Chadwick. Deepfakes and disinformation: Exploring the impact of synthetic political video on deception, uncertainty, and trust in news. *Social media+ society*, 6(1):2056305120903408, 2020.

[62] Margot J van der Goot, Nathalie Koubayová, and Eva A van Reijmersdal. Understanding users' responses to disclosed vs. undisclosed customer service chatbots: a mixed methods study. *AI & Society*, pages 1–14, 2024.

[63] Luc van der Gun and Olivia Guest. Artificial intelligence: Panacea or non-intentional dehumanisation? *Journal of Human-Technology Relations*, 2, 2024.

[64] Wyatt Walls. Gemini loves the smell of pizza. That's how we can tell Gemini is human, March 2024. URL `https://x.com/lefthanddraft/status/1772693138714362021`.

[65] David Watson. The rhetoric and reality of anthropomorphism in artificial intelligence. *Minds Mach.*, 29(3):417–440, September 2019.

[66] Lindsay Weinberg. Rethinking fairness: An interdisciplinary survey of critiques of hegemonic ML fairness approaches. *Journal of Artificial Intelligence Research*, 74:75–109, 2022.

[67] Cedric Deslandes Whitney and Justin Norman. Real risks of fake data: Synthetic data, diversity-washing and consent circumvention. In *The 2024 ACM Conference on Fairness, Accountability, and Transparency*, pages 1733–1744, 2024.

[68] David Gray Widder, Dawn Nafus, Laura Dabbish, and James Herbsleb. Limits and possibilities for "Ethical AI" in open source: A study of deepfakes. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, pages 2035–2046, 2022.

[69] Terry Winograd and Fernando Flores. *Understanding computers and cognition: A new foundation for design*, volume 335. Ablex publishing corporation Norwood, NJ, 1986.

[70] Matteo Wong. OpenAI's Big Reset. *The Atlantic*, 2024.

[71] Sylvia Wynter. Unsettling the coloniality of being/power/truth/freedom: Towards the human, after man, its overrepresentation—an argument. *CR: The new centennial review*, 3(3):257–337, 2003.

[72] Brahim Zarouali, Mykola Makhortykh, Mariella Bastian, and Theo Araujo. Overcoming polarization with chatbot news? Investigating the impact of news content containing opposing views on agreement and credibility. *Eur. J. Commun.*, 36(1):53–68, February 2021.