

Hazem Ibrahim

hazem.ibrahim@nyu.edu | hazemibr.com | [Google Scholar](#)

BIO

My research studies how algorithms, data, and institutions create systematic inequalities in who is seen, heard, and valued. I treat bias in socio-technical systems — social-media platforms, large language models, academic publishing — not as a bug in a single model but as an emergent property of tightly coupled computational and social processes. My work moves from measurement (large-scale audits quantifying inequality) to causation (field experiments isolating how platform decisions produce bias) to mitigation (algorithmic and policy interventions). This spans two empirical domains: social media and LLMs, where I audit recommendation algorithms and measure the political behavior of language models, and academia, where I uncover citation manipulation and demonstrate racial and institutional biases in access to scientific knowledge.

EDUCATION

2023–2027 (expected)	Ph.D. , Computer Science, New York University Thesis: <i>Understanding bias, inequality, and human behavior in socio-technical systems.</i> Advisor: Yasir Zaki
2018–2020	M.Sc. , Computer Science, University of Toronto Advisor: Syed Ishtiaque Ahmed
2014–2018	B.Sc. , Computer Engineering, New York University Abu Dhabi

AWARDS & HONORS

2025	Best Poster Award, AI4GS
2024	Best Parallel Talk Award, IC2S2
2024	Best Poster Award, IC2S2
2023	Innovators Under 35, MIT Technology Review
2018	Founder's Day Award, NYU Abu Dhabi

PUBLICATIONS

* denotes corresponding author; † denotes joint first authorship.

Published

- [1] **Ibrahim, H.**[†], Jang, H. D.[†], Aldahoul, N., Kaufman, A. R., Rahwan, T. & Zaki, Y. Systematic partisan content skews in TikTok during the 2024 U.S. elections. *Nature* (2026).
- [2] Jang, H. D., **Ibrahim, H.**, Asim, R., Varvello, M. & Zaki, Y. A tale of three location trackers: AirTag, SmartTag, and Tile. In *Proc. ACM Internet Measurement Conference* (2026). *Forthcoming*.

- [3] **Ibrahim, H.**, Khan, F., T., Rahwan, T. & Zaki, Y. Analyzing political stances on Twitter/X in the lead-up to the 2024 U.S. election. *PoliticalNLP Workshop at LREC* (2026).
- [4] **Ibrahim, H.[†]**, Liu, F.[†], Zaki, Y. & Rahwan, T. Citation manipulation through citation mills and pre-print servers. *Scientific Reports* 15, 5480 (2025).
- [5] **Ibrahim, H.**, Sabie, D., Roy, P., Bhattacharjee, A., Alam, S. M. R., Mim, N. J. & Ahmed, S. I. Heritage language maintenance: The case of Bangladeshi immigrants in Canada. *Interaction Design and Architecture(s) Journal (IxD&A)*, 170–206 (2025).
- [6] **Ibrahim, H.**, Debicki, M., Rahwan, T. & Zaki, Y. Big tech dominance despite global mistrust. *IEEE Transactions on Computational Social Systems* 11, 3741–3752 (2024).
- [7] **Ibrahim, H.**, AlDahoul, N., Lee, S., Rahwan, T. & Zaki, Y. YouTube’s recommendation algorithm is left-leaning in the United States. *PNAS Nexus* 2, pgad264 (2023).
- [8] **Ibrahim, H.**, Asim, R., Zaffar, F., Rahwan, T. & Zaki, Y. Rethinking homework in the age of artificial intelligence. *IEEE Intelligent Systems* 38, 24–27 (2023).
- [9] **Ibrahim, H.**, Liu, F., Asim, R., Battu, B., Benabderrahmane, S., Alhafni, B., Adnan, W., Alhanai, T., AlShebli, B., Baghdadadi, R. et al. Perception, performance, and detectability of conversational artificial intelligence across 32 university courses. *Scientific Reports* 13, 12187 (2023).
- [10] **Ibrahim, H.**, Asim, R., Varvello, M. & Zaki, Y. I tag, you tag, everybody tags! In *Proc. ACM Internet Measurement Conference* (2023), 561–568.
- [11] **Ibrahim, H.** & Ibrahim, W. Gamification in online educational systems. In *6th International Conference on Higher Education Advances (HEAd’20)* (2020), 1217–1224.
- [12] Ibrahim, W. & **Ibrahim, H.**. Multithreaded and reconvergent aware algorithms for accurate digital circuits reliability estimation. *IEEE Transactions on Reliability* 68, 514–525 (2018).

Preprints & Under Review

- [13] **Ibrahim, H.[†]**, Liu, F.[†], Mengal, K., Kaufman, A., Zaki, Y. & Rahwan, T. Causal evidence of racial and institutional biases in accessing paywalled articles and scientific data (2025). *Revise and resubmit at Science*.
- [14] **Ibrahim, H.**, Rahwan, T., Zaki, Y. & Park, M. On-screen inertia: persistent racial and gender disparities in Hollywood film (1900–2024). *Under review at PNAS*.
- [15] **Ibrahim, H.[†]**, AlDahoul, N.[†], Abbasi, S. M. A., Zaffar, F., Rahwan, T. & Zaki, Y. A longitudinal analysis of racial and gender bias in New York Times and Fox News images and articles (2025). *Under review at Scientific Reports*.
- [16] Kuo, C. W.[†], Chu, K.[†], AlDahoul, N.[†], **Ibrahim, H.[†]**, Rahwan, T. & Zaki, Y. Neutralizing the narrative: AI-powered debiasing of online news articles (2025). *Revise and resubmit at Engineering Applications of Artificial Intelligence*.
- [17] AlDahoul, N.[†], **Ibrahim, H.[†]**, Park, M., Rahwan, T. & Zaki, Y. Inclusive content reduces racial and gender biases, yet non-inclusive content dominates popular media outlets (2025). *Under review at EPJ Data Science*.
- [18] AlDahoul, N., **Ibrahim, H.**, Park, M., Rahwan, T. & Zaki, Y. Large language models are often politically extreme, usually ideologically inconsistent, and persuasive even in informational contexts. arXiv:2505.04171 (2025). *Revise and resubmit at American Political Science Review*.

- [19] Khan, F. K., **Ibrahim, H.**, Rahwan, T. & Zaki, Y. Who gets seen in the age of AI? Adoption patterns of large language models in scholarly writing and citation outcomes (2025). *Revise and resubmit at Journal of Informetrics*.
- [20] **Ibrahim, H.**, et al. Schadenfreude in the digital public sphere: a cross-national and decade-long analysis. arXiv:2601.20413 (2026). *Under review at Journal of Big Data*.
- [21] **Ibrahim, H.*** & Girmay, T. Auditing differential visibility of political content on TikTok. arXiv:2607.17356 (2026). *Under review at ICWSM 2027*.
- [22] **Ibrahim, H.***, Rahwan, T. & Zaki, Y. Bias at the borderline: who gets the benefit of the doubt in peer review? arXiv:2607.26280 (2026). *Preprint*.

In Preparation

- [23] **Ibrahim, H.**, Liu, F., Zaki, Y. & Rahwan, T. Two-thirds of citations to review papers belong to original research. *In preparation*.
- [24] **Ibrahim, H.**, Holovatska, Y., Aldahoul, N., Zaki, Y. & Rahwan, T. Manufactured participation in the Russia–Ukraine information war is industrial in supply but weak in demand. *In preparation*.
- [25] **Ibrahim, H.***, et al. Compression vs. displacement of early-career scientists under generative AI. *In preparation*.
- [26] **Ibrahim, H.***, et al. Productivity gains from AI adoption concentrate among already advantaged and internationally mobile researchers. *In preparation*.
- [27] **Ibrahim, H.***, et al. Fabricated citations spread faster than real ones. *In preparation*.
- [28] **Ibrahim, H.***, et al. Multilingual under-enforcement in deployed content moderation. *In preparation*.
- [29] **Ibrahim, H.***, et al. A forensic analysis of the first deployed algorithmic bias-disclosure mandate. *In preparation*.
- [30] **Ibrahim, H.***, et al. Stale & skewed: auditing LLM-as-annotator validity for multilingual computational social science. *In preparation*.
- [31] **Ibrahim, H.***, et al. Auditing TikTok’s delivery of gambling ads to youth under the Digital Services Act. *In preparation*.
- [32] **Ibrahim, H.***, et al. Chance, persistent advantage, and the generative-AI era in open-source package careers. *In preparation*.
- [33] **Ibrahim, H.***, et al. Is a country’s science more than the sum of its scientists? A coupling-fraction test across sanctions and funding collapses. *In preparation*.
- [34] **Ibrahim, H.***, et al. Loss-chasing on a public ledger: behavioral markers of gambling harm on a prediction market. *In preparation*.
- [35] **Ibrahim, H.***, et al. Epistemic leverage in permissionless and regulated prediction markets. *In preparation*.
- [36] **Ibrahim, H.***, et al. Artificial intelligence and the global diffusion of research technologies. *In preparation*.
- [37] Omari, A, **Ibrahim, H.**, Aldahoul, N., Zaki, Y., Rahwan, T. & Kaufman, A. Measuring the Political Ideology of LLMs Across 90 Countries. *In preparation*.

TEACHING

New York University Abu Dhabi

Instructor of Record

2021–2022	Machine Learning (CS-UH 2220) — Spring 2022
2021–2022	Computational Social Science (CS-UH 2219E) — Spring 2021, Fall 2021
2021–2022	Database Systems (CS-UH 2214) — Fall 2021, Spring 2022
2021–2022	Data Structures (CS-UH 1050) — Spring 2021, Fall 2021
2020–2022	Discrete Mathematics (CS-UH 1002) — Fall 2021, Spring 2022, Summer 2022
2020–2022	Introduction to Computer Science (CS-UH 1001) — Fall 2020, Summer 2021, Fall 2021

University of Toronto

Teaching Assistant

2018–2019	Introduction to Computer Science (CSC148H) — Fall 2018, Spring 2018, Fall 2019
-----------	--

ACADEMIC ADVISING

Undergraduate Students Supervised

2026	Yana Holovatska, NYU Abu Dhabi
2026	Tewoflos Girmay, NYU Abu Dhabi
2025	Farhan Kamrul Khan, NYU Abu Dhabi
2025	Chen Wei Kuo, NYU Abu Dhabi
2025	Kevin Chu, NYU Abu Dhabi

INDUSTRY EXPERIENCE

2017	Software Engineering Intern, Meta, London, United Kingdom
------	---

SELECTED MEDIA COVERAGE

2026	The Guardian: TikTok’s algorithm favored Republican content in 2024 US elections [1]
2026	Nature: TikTok’s algorithm systematically skewed to the right during the 2024 US elections [1]
2026	El País: TikTok helped Trump win the 2024 election [1]
2026	Phys.org: TikTok algorithm shows pro-Republican bias [1]
2025	PsyPost: TikTok’s algorithm exhibited pro-Republican bias during 2024 presidential race, study finds [1]
2025	United Nations: Human Development Report 2025 – A matter of choice: People and possibilities in the age of AI [8]
2024	Nature: The citation black market: schemes selling fake references alarm scientists [3]
2024	Science: Vendor offering citations for purchase is latest bad actor in scholarly publishing [3]
2023	Scientific American: ChatGPT can get good grades. What should educators do about it? [8]

- 2023 **The Times:** Students “double cheat” to hide AI in their answers [8]
- 2023 **New Scientist:** ChatGPT gets better marks than students in some university courses [8]
- 2023 **The Independent:** AI can write better university assignments than students [8]
- 2023 **The Daily Beast:** OpenAI’s ChatGPT surpasses college students in writing homework [8]
- 2023 **Daily Mail:** AI can achieve better grades than university students [8]
- 2023 **The National:** AI models ‘achieve higher university grades than students’ [8]
- 2023 **Nature Middle East:** ChatGPT sparks debate among students and educators [8]
- 2023 **Nature Asia:** ChatGPT’s answers to essay questions can surpass those of university students [8]
- 2023 **Neuroscience News:** Can ChatGPT rival university students in academic performance? [8]
- 2023 **PsyPost:** YouTube’s recommendation system exhibits left-leaning bias, new study suggests [6]

Additional coverage in *The Telegraph*, *The Herald Sun*, *Breitbart News*, *Truthout*, *DerStandard*, *Government Technology*, *Phys.org*, *American Council on Science and Health*, and *Daily Caller*.

REFERENCES

Available upon request.