Sayash Kapoor

https://saya.sh

October 2026

Email : sayashk@princeton.edu

Appointments

UC Berkeley School of Information — Berkeley, CA
Incoming Assistant Professor — *Starting July 2027*

Education

Princeton University — Princeton, NJ
Doctor of Philosophy in Computer Science; GPA: 3.96/4.00 — *January 2021 – July 2026*

Indian Institute of Technology Kanpur — Kanpur, India
Bachelor of Technology in Computer Science; GPA: 9.9/10.0 — *July 2015 – June 2019*

École Polytechnique Fédérale de Lausanne — Lausanne, Switzerland
Exchange Student in Computer Science; GPA: 5.7/6.0 — *August 2017 – May 2018*

Awards and Recognition

IP3 Information Policy Award, Public Knowledge
October 2026

MIT Technology Review Innovators Under 35
September 2026

Privacy Papers for Policymakers Award, Future of Privacy Forum
January 2026
For "AI as Normal Technology" (with Arvind Narayanan)

Porter Ogden Jacobus Fellowship
September 2025
The highest honor that the Princeton Graduate School bestows on an enrolled Ph.D. student
One of four awardees across all departments

2026 Siebel Scholar
September 2025
One of 5 Princeton computer science graduate students selected for an award of $35,000

Mozilla Senior Fellowship in Trustworthy AI
January 2025
$130,000 fellowship to advance trustworthy AI through evidence-based tech policy

Princeton School of Engineering and Applied Science Award for Excellence
September 2024

Laurance S. Rockefeller Graduate Prize Fellowship at Princeton University
2024-25
First computer scientist in 20 years to receive the $60,000 graduate prize fellowship from Princeton University

McGraw Center for Teaching and Learning Inaugural Exemplar Mentor Award
May 2024

Featured in the inaugural list: *TIME100 Most Influential People in AI*
September 2023

Impact Recognition Award, ACM CSCW
November 2022

Motorola Gold Medal, IIT Kanpur
June 2019

Best Paper Award, ACM FAccT
January 2019

CMMRS 2018, Pre-Doctoral Research School, Max Planck Institute (Saarbrücken)
August 2018

Bronze Medal, ACM ICPC SWERC, École Normale Supérieure
November 2017

Academic Excellence Award, IIT Kanpur
July 2016, July 2017

Outstanding Freshman Award, IIT Kanpur
March 2016

Books

[1] **AI as Normal Technology**
Arvind Narayanan, **Sayash Kapoor**
Princeton University Press (in preparation)

[2] **AI Snake Oil**
Arvind Narayanan, **Sayash Kapoor**
Princeton University Press (2024)
Peer-reviewed book on what AI can and can't do. Named one of Nature's 10 best books of 2024, Bloomberg's 49 best books of 2024, and Forbes's 10 must-read tech books of 2024

Peer-reviewed publications

An asterisk (*) denotes equal contribution.
Link to Google Scholar · H-Index: 29 · I10-Index: 47 · Citations: 6,040

[1] **Unpacking Open Source Artificial Intelligence: Toward a Framework for Openness in Foundation Models**
Camille François, ..., **Sayash Kapoor** et al.
Communications of the ACM 69(8), 80–87 (2026)

[2] **Holistic Agent Leaderboard: The Missing Infrastructure for AI Agent Evaluation**
Sayash Kapoor* et al.
International Conference on Learning Representations (ICLR 2026)

[3] **The Limits of Inference Scaling Through Resampling**
Benedikt Stroebl, **Sayash Kapoor**, Arvind Narayanan
International Conference on Learning Representations (ICLR 2026)

[4] **Towards a Science of AI Agent Reliability**
Stephan Rabanser, **Sayash Kapoor**, Peter Kirgis, Kangheng Liu, Saiteja Utpala, Arvind Narayanan
International Conference on Machine Learning (ICML 2026)
Also presented at: Workshop on Failure Modes of Agentic AI at ICML 2026

[5] **FLARE-AI: Flaw Reporting for AI · OpenReview · ICML poster**
Shayne Longpre, ..., **Sayash Kapoor** et al.
International Conference on Machine Learning (ICML 2026)

[6] **The 2025 Foundation Model Transparency Index · Project page**
Alexander Wan, ..., **Sayash Kapoor** et al.
Transactions on Machine Learning Research (TMLR 2026)

[7] **AI Agents That Matter**
Sayash Kapoor*, Benedikt Stroebl*, Zachary S. Siegel, Nitya Nadgir, Arvind Narayanan
Transactions on Machine Learning Research (TMLR 2025)

[8] **Establishing Best Practices in Building Rigorous Agentic Benchmarks**
Yuxuan Zhu, ..., **Sayash Kapoor** et al.
Neural Information Processing Systems Datasets and Benchmarks (NeurIPS 2025)

[9] **The Leaderboard Illusion**
Shivalika Singh, ..., **Sayash Kapoor** et al.
Neural Information Processing Systems Datasets and Benchmarks (NeurIPS 2025)

[10] **Build Agent Advocates, Not Platform Agents**
Sayash Kapoor*, Noam Kolt*, Seth Lazar*
International Conference on Machine Learning position paper track (ICML 2025)

[11] **Position: In-House Evaluation Is Not Enough. Towards Robust Third-Party Evaluation and Flaw Disclosure for General-Purpose AI · arXiv**
Shayne Longpre, ..., **Sayash Kapoor** et al.
International Conference on Machine Learning position paper track (ICML 2025 **Spotlight**)

[12] **The Reality of AI and Biorisk**
Aidan Peppin, ..., **Sayash Kapoor** et al.
ACM Conference on Fairness, Accountability, and Transparency (FAccT 2025)

[13] **The 2024 Foundation Model Transparency Index**
Rishi Bommasani, ..., **Sayash Kapoor** et al.
Transactions on Machine Learning Research (TMLR 2025)

[14] **The 2023 Foundation Model Transparency Index**
Rishi Bommasani, ..., **Sayash Kapoor** et al.
Transactions on Machine Learning Research (TMLR 2025 **Featured certification**)

[15] **CORE-Bench: Fostering the Credibility of Published Research Through a Computational Reproducibility Agent Benchmark**
Zachary S. Siegel, **Sayash Kapoor**, Nitya Nadgir, Benedikt Stroebl, Arvind Narayanan
Transactions on Machine Learning Research (TMLR 2024)

[16] **REFORMS: Consensus-based Recommendations for Machine-learning-based Science · Blog post**
Sayash Kapoor et al.
Science Advances (2024)

[17] **Considerations for governing open foundation models**
Rishi Bommasani, **Sayash Kapoor** et al.
Science (2024)

[18] **On the Societal Impact of Open Foundation Models · Blog post**
Sayash Kapoor* et al.
International Conference on Machine Learning position paper track (ICML 2024 **Oral**)

[19] **A Safe Harbor for AI Evaluation and Red Teaming · Blog post**
Shayne Longpre, **Sayash Kapoor** et al.
International Conference on Machine Learning position paper track (ICML 2024 **Oral**)
Our open letter to AI companies calling for a safe harbor was signed by over 350 academics, researchers, and civil society members.

[20] **Against Predictive Optimization: On the Legitimacy of Decision-Making Algorithms that Optimize Predictive Accuracy · Blog post**
Angelina Wang*, **Sayash Kapoor***, Solon Barocas, Arvind Narayanan
ACM Journal on Responsible Computing (JCR 2024)
Also presented at: Philosophy, AI, and Society (2023); Data (Re)Makes the World (2023); ACM FAccT (2023)

[21] **How large language models can reshape collective intelligence**
Jason W. Burton, ..., **Sayash Kapoor** et al.
Nature Human Behaviour (2024)

[22] **The Responsible Foundation Model Development Cheatsheet: A Review of Tools & Resources**
Shayne Longpre, ..., **Sayash Kapoor** et al.
Transactions on Machine Learning Research (TMLR 2024 **Survey certification**)

[23] **Foundation Model Transparency Reports · Blog post**
Rishi Bommasani, ..., **Sayash Kapoor** et al.
AAAI/ACM Conference on Artificial Intelligence, Ethics, and Society (AIES 2024)

[24] **Leakage and the reproducibility crisis in ML-based science**
Sayash Kapoor, Arvind Narayanan
Patterns (2023)

[25] **Weaving Privacy and Power: On the Privacy Practices of Labor Organizers in the U.S. Technology Industry**
Sayash Kapoor*, Matthew Sun*, Mona Wang*, Klaudia Jaźwińska*, Elizabeth Anne Watkins*
ACM Conference on Computer-Supported Cooperative Work & Social Computing (CSCW 2022 **Impact Recognition Award**)

[26] **The worst of both worlds: A comparative analysis of errors in learning from data in psychology and machine learning**
Jessica Hullman, **Sayash Kapoor**, Priyanka Nanayakkara, Andrew Gelman, Arvind Narayanan
AAAI/ACM Conference on Artificial Intelligence, Ethics, and Society (AIES 2022)

[27] **Controlling polarization in personalization: an algorithmic framework**
L. Elisa Celis, **Sayash Kapoor**, Farnood Salehi, and Nisheeth K. Vishnoi
ACM Conference on Fairness, Accountability, and Transparency (FAccT 2019 **Best Paper Award**)

[28] **Corruption-tolerant bandit learning**
Sayash Kapoor, Kumar Kshitij Patel, and Purushottam Kar
Machine Learning (2019)

[29] **A dashboard for controlling polarization in personalization**
L. Elisa Celis, **Sayash Kapoor**, Farnood Salehi, Vijay Keswani, and Nisheeth K. Vishnoi
AI Communications (2019)
Also presented at: International Joint Conference on Artificial Intelligence Demos Track (IJCAI 2018)

[30] **Balanced news using constrained bandit-based personalization**
Sayash Kapoor, Vijay Keswani, Nisheeth K. Vishnoi, and L. Elisa Celis
International Joint Conference on Artificial Intelligence Demos Track (IJCAI 2018)

Preprints, manuscripts, and comments

[1] **Births are difficult to predict even with rich survey and full-population register data**
Elizaveta Sivak, ..., **Sayash Kapoor** et al.
Preprint (2026)

[2] **Can AI agents conduct open-ended AI research? Early evidence from two case studies · Project page**
Peter Kirgis, **Sayash Kapoor** et al.
Preprint (2026)

[3] **Bridging Predictions and Interventions: An Integrated Framework for Automated Decision-Systems**
Inioluwa Deborah Raji*, ..., **Sayash Kapoor** et al.
Preprint (2026)

[4] **Life After Benchmark Saturation: A Case Study of CORE-Bench · OpenReview**
Nitya Nadgir, **Sayash Kapoor** et al.
Second Workshop on Agents in the Wild: Safety, Security, and Beyond at ICML 2026

[5] **Log analysis is necessary for credible evaluation of AI agents · OpenReview**
Peter Kirgis, **Sayash Kapoor** et al.
Workshop on Failure Modes of Agentic AI at ICML 2026

[6] **Seven simple steps for log analysis in AI systems**
Magda Dubois, ..., **Sayash Kapoor** et al.
Preprint (2026)

[7] **Open-World Evaluations for Measuring Frontier AI Capabilities · CRUX · Blog post · OpenReview**
Sayash Kapoor et al.
Second Workshop on Agents in the Wild: Safety, Security, and Beyond at ICML 2026

[8] **International AI Safety Report 2026**
Yoshua Bengio, ..., **Sayash Kapoor** et al. (2026)
A report on the state of advanced AI capabilities and risks written by 100 AI experts

[9] **AI Won't Automatically Make Legal Services Cheaper**
Justin Curl, **Sayash Kapoor**, Arvind Narayanan
Lawfare (2026)

[10] **Bridging Prediction and Intervention Problems in Social Systems**
Lydia T. Liu, ..., **Sayash Kapoor** et al.
Preprint (2025)

[11] **A Different Approach to AI Safety: Proceedings from the Columbia Convening on Openness in Artificial Intelligence and AI Safety**
Camille François, ..., **Sayash Kapoor** et al.
Proceedings from the Columbia Convening on Openness in Artificial Intelligence and AI Safety (2025)

[12] **AI as Normal Technology · PDF**
Arvind Narayanan, **Sayash Kapoor**
Knight First Amendment Institute (2025). Received press coverage in New York Times, Economist, New Yorker etc.

[13] **Why an overreliance on AI-driven modelling is bad for science**
Arvind Narayanan, **Sayash Kapoor**
Nature (2025)

[14] **International AI Safety Report**
Yoshua Bengio, ..., **Sayash Kapoor** et al. (2025)
A report on the state of advanced AI capabilities and risks written by 100 AI experts

[15] **Towards a Framework for Openness in Foundation Models: Proceedings from the Columbia Convening on Openness in Artificial Intelligence**
Adrien Basdevant, ..., **Sayash Kapoor** et al.
Preprint (2024)

[16] **Promises and pitfalls of artificial intelligence for legal applications · Blog post**
Sayash Kapoor, Peter Henderson, Arvind Narayanan
Journal of Cross-disciplinary Research in Computational Law (CRCL 2024)

[17] **The limitations of machine learning models for predicting scientific replicability**
M. J. Crockett, Xuechunzi Bai, **Sayash Kapoor**, Lisa Messeri, and Arvind Narayanan
Proceedings of the National Academy of Sciences Comment (PNAS 2023)

[18] **How to Prepare for the Deluge of Generative AI on Social Media**
Sayash Kapoor, Arvind Narayanan
Knight First Amendment Institute (2023)

Public Writing

In addition to the texts below, I write extensively on the AI as Normal Technology newsletter, which has over 85,000 readers.

[1] **A big-tent or small-tent AI safety movement?**
Arvind Narayanan, **Sayash Kapoor**
AI as Normal Technology (October 2026)

[2] **The AI-as-Normal-Technology view of loss-of-control incidents**
Sayash Kapoor, Arvind Narayanan
AI as Normal Technology (2026)

[3] **Why AI hasn't replaced software engineers, and won't**
Arvind Narayanan, **Sayash Kapoor**
AI as Normal Technology (2026)

[4] **Could AI slow science?**
Sayash Kapoor, Arvind Narayanan
AI as Normal Technology (2025)

[5] **AGI is not a milestone**
Sayash Kapoor, Arvind Narayanan
AI as Normal Technology (2025)

[6] **AI as Normal Technology**
Arvind Narayanan, **Sayash Kapoor**
AI as Normal Technology (2025)

[7] **Worry About Misuse of AI, Not Superintelligence**
Arvind Narayanan, **Sayash Kapoor**
WIRED (2024)

[8] **We Looked at 78 Election Deepfakes. Political Misinformation Is Not an AI Problem.**
Sayash Kapoor, Arvind Narayanan
Knight First Amendment Institute (2024)

[9] **Is AI too dangerous to release openly?**
Sayash Kapoor, Arvind Narayanan
Princeton Engineering Magazine (2024)

[10] **A Safe Harbor for AI Evaluation and Red Teaming**
Shayne Longpre, **Sayash Kapoor** et al.
Knight First Amendment Institute (2024)

[11] **Does AI Pose an Existential Risk to Humanity? Two Sides Square Off**
Arvind Narayanan, **Sayash Kapoor**
The Wall Street Journal, November 2023

[12] **How to report better on artificial intelligence**
Sayash Kapoor, Hilke Schellmann, Ari Sen
Columbia Journalism Review (2023)

[13] **Generative AI companies must publish transparency reports**
Arvind Narayanan, **Sayash Kapoor**
Knight First Amendment Institute (2023)

[14] **A Checklist of Eighteen Pitfalls in AI Journalism**
Sayash Kapoor, Arvind Narayanan
Reporting on artificial intelligence: a handbook for journalism educators, UNESCO (2023)

[15] **The LLaMA is out of the bag. Should we expect a tidal wave of disinformation?**
Arvind Narayanan, **Sayash Kapoor**
Knight First Amendment Institute (2023)

[16] **Through the Wire**
Klaudia Jaźwińska, Matthew Sun, **Sayash Kapoor**, Mona Wang
Logic Mag (2022)

[17] **The platform as the city**
Mac Arboleda, Palak Dudani, **Sayash Kapoor**, Lorna Xu
ACM Interactions Mag (2021)

Policy Input

[1] **In Support of Mandatory Nucleic Acid Synthesis Screening and Recordkeeping**
Signed open letter to policymakers (2026)

[2] **AI Agents and Reliability: Princeton AI Policy Precepts · Program page**
Sayash Kapoor, Arvind Narayanan, Mihir Kshirsagar, Jaime Fernández Fisac
Speaker, Princeton SPIA DC Center (2026)

[3] **Response to Request for Information Regarding Security Considerations for Artificial Intelligence Agents**
Stephan Rabanser, **Sayash Kapoor**, Peter Kirgis, Mihir Kshirsagar, Arvind Narayanan
Submitted to the Center for AI Standards and Innovation, NIST, U.S. Department of Commerce (2026)

[4] **Princeton AI Policy Precepts: Fall 2025, Session 1 · Program page**
Sayash Kapoor, Mihir Kshirsagar, Hilke Schellmann
Speaker and facilitator, Princeton SPIA DC Center (2025)

[5] **Response to ASAP Request for Information**
Sayash Kapoor, Arvind Narayanan
Submitted to the offices of Senator Heinrich and Senator Rounds (2025)

[6] **Safeguarding Third-Party AI Research**
Kevin Klyman, Shayne Longpre, **Sayash Kapoor**, Rishi Bommasani, Percy Liang, Peter Henderson
Stanford HAI Policy Brief (2025)

[7] **Panel on recent AI advances · Program page**
Panelist, Brookings AI Policy Idea Incubator, Center on Regulation and Markets (2025)

[8] **Response to NIST's Request for Comments on Managing Misuse Risk for Dual-Use Foundation Models**
Sayash Kapoor, Mihir Kshirsagar, Arvind Narayanan, Benedikt Stroebl
Submitted to NIST / U.S. AI Safety Institute (2024)

[9] **Generative AI Companies: Safe Harbor and Whistleblower Protections**
Sayash Kapoor, Arvind Narayanan
Testimony to the New Jersey Assembly Science, Innovation and Technology Committee (2024)

[10] **Response to the EU AI Office's Consultation on the AI Act**
Varun Nagaraj Rao, Kyler Zhou, **Sayash Kapoor**, Arvind Narayanan
Submitted to the EU AI Office (2024)

[11] **Princeton Dialogues in AI: Predictive AI**
Arvind Narayanan, **Sayash Kapoor**, Peter Henderson
Senate AI Caucus (2024)

[12] **Princeton Dialogues in AI: AI Safety**
Sayash Kapoor, Mihir Kshirsagar
Senate AI Caucus and House AI Caucus (2024)

[13] **A Safe Harbor For AI Researchers: Promoting Safety And Trustworthiness Through Good-Faith Research**
Kevin Klyman, **Sayash Kapoor**, Shayne Longpre
Federation of American Scientists: Policy memo (2024)

[14] **Reducing harm from deepfakes**
Sayash Kapoor, Arvind Narayanan
Testimony to the New Jersey Assembly Science, Innovation and Technology Committee (2024)

[15] **Response to Request for Comment on Dual Use Foundation Artificial Intelligence Models With Widely Available Model Weights**
Alondra Nelson, ..., **Sayash Kapoor** et al.
Submitted to the National Telecommunications and Information Administration (2024)

[16] **Comment to the Copyright Office in Support of a Safe Harbor Exemption for Generative AI Research**
Kevin Klyman, Shayne Longpre, **Sayash Kapoor**, Arvind Narayanan, Aleksandra Korolova, Peter Henderson
Submitted to the U.S. Copyright Office (2024)

[17] **Beyond the AI hype**
Sayash Kapoor, Arvind Narayanan
Government of Canada's Federal Foresight Network (2024)

[18] **Intro to AI/ML for Regulators**
Sayash Kapoor, Mihir Kshirsagar
Consumer Financial Protection Bureau (2024)

[19] **How to Prepare for the Deluge of Generative AI on Social Media**
Sayash Kapoor, Arvind Narayanan
Federal Trade Commission Division of Advertising Practices Tech Speaker Series (2023)

[20] **Considerations for governing open foundation models · Blog post**
Rishi Bommasani, **Sayash Kapoor** et al.
Stanford HAI Issue Brief (2023)

[21] **The urgent need for accountability in predictive AI**
Arvind Narayanan, **Sayash Kapoor**
Congressional Forum (2023)

[22] **Three Ideas for Regulating Generative AI · Blog post**
Rishi Bommasani, **Sayash Kapoor**, Daniel Zhang, Arvind Narayanan, Percy Liang
Submitted to the National Telecommunications and Information Administration (2023)

[23] **CITP Comments on AI Accountability · Blog post**
Archana Ahlawat, ..., **Sayash Kapoor** et al.
Submitted to the National Telecommunications and Information Administration (2023)

[24] **Calling for Investing in Equitable AI Research in Nation's Strategic Plan · Blog post**
Solon Barocas, **Sayash Kapoor**, Mihir Kshirsagar, Arvind Narayanan
Response to RFI, White House Office of Science and Technology Policy (2022)

[25] **National AI Research Infrastructure Needs to Support Independent Evaluation of Performance Claims · Blog post**
Sayash Kapoor, Mihir Kshirsagar, Arvind Narayanan
Response to RFI, White House Office of Science and Technology Policy and National Science Foundation (2021)

Work Experience

Facebook London, UK
Software Engineer *July 2019 - December 2020*

Developed machine learning models to detect non-consensual intimate imagery across Facebook and Instagram. Interned from May – August 2018; developed machine learning models to detect child sexual abuse material across Facebook and Instagram.

Service and Workshops

Workshop organizer
The Future of Third-Party AI Evaluation (Stanford & Princeton)
Over 400 registrations. Video recordings seen over 2,000 times.

Workshop on Useful and Reliable AI Agents (Princeton)
Over 500 registrations. Video recordings seen over 3,300 times.

Responsible and open foundation models (Princeton & Stanford)
Over 900 registrations. Video recordings seen over 3,200 times.

The Reproducibility Crisis in ML-based Science (Princeton)
Over 1,700 registrations. Video recordings seen over 6,500 times.

Program committee member
AIES 2022, FAccT 2022, FAccT 2023, AIES 2024, FAccT 2024

Reviewer
Nature, Science Advances, PLoS ONE, JMLR, Patterns, ICML 2022, ICML 2025, ICLR 2026

Mentor
Princeton Computer Science Pre-Application Graduate Mentoring Program

Teaching

Open Source AI: Benefits and Risks
Wilson Center Advanced AI Labs. Four courses taught in Fall 2024 and Spring 2025

Princeton Policy Precepts: AI's impact on labor
Princeton University School of Public and International Affairs, Fall 2025

Princeton Policy Precepts: Predictive AI
Princeton University School of Public and International Affairs, Fall 2024

CS 5382: Practical Principles for Designing Fair Algorithms
Cornell University. Guest Lecturer. Spring 2024.

COS 350: Ethics of Computing
Princeton University. Preceptor and teaching assistant. Fall 2023.

COS 324: Introduction to Machine Learning
Princeton University. Preceptor and teaching assistant. Spring 2023.

PHI 543: Machine Learning: A Practical Introduction for Humanists and Social Scientists
Princeton University. Guest Lecturer. Fall 2023.

SOC 306: Machine Learning with Social Data: Opportunities and Challenges
Princeton University. Guest Lecturer. Spring 2022, Spring 2023.

Selected Press

State of AI Report 2026
State of AI Report, October 2026

He tests whether AI is living up to its promise
MIT Technology Review, Innovators Under 35, September 2026

How the Clash Between Money and Safety Created a Monumental Crisis for AI
The Wall Street Journal, September 2026

Why Andon Labs Puts AI Agents in Charge of Real Businesses
IEEE Spectrum, September 2026

AI insiders fear extinction. Security experts see a familiar fight
Scientific American, September 2026

AI's recursive self-improvement might not come so quickly after all
MIT Technology Review, August 2026

AI isn't ready to research itself
Nature, August 2026

Shaping AI policy as an academic
Horizon Institute, Launchpad, July 2026

The AI Shift: How capable is AI? It depends what you're worried about
Financial Times, July 2026

You Can Now Sound the Alarm on AI Behaving Badly
WIRED, July 2026

Researchers caught in the crossfire as companies and government grapple over AI safety
Science, June 2026

Anthropic's New Fable AI Model Is Met With User Backlash Over Restrictions
The Wall Street Journal, June 2026

How (un)reliable are AI agents?
Financial Times, April 2026

AI agents are getting more capable, but reliability is lagging—and that's a problem
Fortune, March 2026

The State of AI: life in 2030
Financial Times, December 2025

Ben's Book of the Month: AI Snake Oil
RSA Conference, December 2025

How investors should place bets on AI
Financial Times, September 2025

AI as Normal Technology, revisited
Platformer, September 2025

"AI will kill everyone" is not an argument. It's a worldview.
Vox, September 2025

What if artificial intelligence is just a "normal" technology?
The Economist, September 2025

How ChatGPT surprised me
The New York Times, August 2025

A.I. May Be Just Kind of Ordinary
The New York Times, August 2025

Silicon Valley Is Drifting Out of Touch With the Rest of America
The New York Times, August 2025

Developers Say GPT-5 Is a Mixed Bag
WIRED, August 2025

Is AI hitting a wall?
Financial Times, August 2025

AI isn't magic. It's just weird
Prospect Magazine, June 2025

Depending on AI alone to make decisions affecting people's lives is worrisome: Sayash Kapoor, co-author, AI Snake Oil
The Indian Express, June 2025

Toward a Realpolitik for AI
Public Books, June 2025

China is gaining ground in the global race to develop AI agents
Rest of World, June 2025

How to build a better AI benchmark
MIT Technology Review, May 2025

Experts Explain: AI as 'normal' technology
The Indian Express, May 2025

Two Paths for A.I.
The New Yorker, May 2025

We need to start thinking of AI as "normal"
MIT Technology Review, April 2025

Researchers Propose a Better Way to Report Dangerous AI Flaws
WIRED, March 2025

What works in AI, and what's just hype
The Hindu, February 2025

Artificial intelligence dominates annual convention
Al Jazeera, January 2025

'AI Snake Oil': A conversation with Princeton AI experts
Princeton Engineering, January 2025

AI Snake Oil with Sayash Kapoor
BCG Henderson Institute, December 2024

Thoughtless obedience and the healing power of trees: 2024's best Books in brief
Nature, December 2024

The Year of the AI Election Wasn't Quite What Everyone Expected
WIRED, December 2024

AI Snake Oil: What Artificial Intelligence Can Do, What It Can't, and How to Tell the Difference
Harvard Gazette, October 2024 (original page unavailable)

Seeing the Forest Through the A.I. Trees
Air Mail, October 2024

Popping the AI Hyperbole Bubble
The Deal, October 2024

Why AI isn't as clever – or as dangerous – as we think
The Telegraph, October 2024

AI Snake Oil: Exposing The Truth Behind Overhyped Claims
NDTV, October 2024

AI Snake Oil (excerpt)
Stanford Social Innovation Review, October 2024

Is the future of AI in health care just a lot of hype?
STAT, October 2024

Generative AI Hype Feels Inescapable. Tackle It Head On With Education
WIRED, September 2024

Professor Arvind Narayanan and Sayash Kapoor Explain AI
Princeton Alumni Weekly, September 2024

Snake Oil: Don't believe the artificial intelligence hype
Financial Review, September 2024

A new book tackles AI hype – and how to spot it
Science News, September 2024

Arvind Narayanan and Sayash Kapoor on AI Snake Oil
Princeton University Press, September 2024

Princeton SPIA AI Experts Separate Hype from Substance in New Book
Princeton SPIA, September 2024

AI Snake Oil: Separating Hype from Reality
Tech Policy Press, September 2024

In the Age of A.I., What Makes People Unique?
The New Yorker, August 2024

'AI Snake Oil' Sorts Promise from Hype
Practical Ecommerce, August 2024

Chatbots Are Primed to Warp Reality
The Atlantic, August 2024

AI agent benchmarks are misleading, study warns
VentureBeat, July 2024

Ray Kurzweil Still Says He Will Merge With A.I.
The New York Times, July 2024

Is AI too dangerous to release openly?
Princeton Engineering, May 2024

Science has an AI problem: This group says they can fix it
ScienceDaily, May 2024

Science has an AI problem. This group says they can fix it.
UC San Diego Today, May 2024

Experts call for legal 'safe harbor' so researchers, journalists and artists can evaluate AI tools
VentureBeat, March 2024

Top AI researchers say OpenAI, Meta and more hinder independent evaluations
Washington Post, March 2024

Researchers, legal experts want AI firms to open up for safety checks
Computerworld, March 2024

Stanford study outlines risks and benefits of open AI models
Axios, March 2024

A Mistral chills European regulators
Politico, March 2024

What are LLMs, and how are they used in generative AI?
Computerworld, February 2024

Princeton University's 'AI Snake Oil' authors say generative AI hype has 'spiraled out of control'
VentureBeat, August 2023

Computer Science Researchers Call Out AI Hype as 'Snake Oil'
Princeton Alumni Weekly, December 2023

OpenAI's ChatGPT turns one year old; what it did (and didn't do)
Computerworld, November 2023

Artificial intelligence is not a silver bullet
NPR, December 2023

AI's Spicy-Mayo Problem
The Atlantic, November 2023

AI Is Becoming More Powerful—but Also More Secretive
WIRED, October 2023

How Does AI 'Think'? We Are Only Starting to Understand That
The Wall Street Journal, October 2023

The world's biggest AI models aren't very transparent
The Verge, October 2023

Maybe We Will Finally Learn More About How A.I. Works
The New York Times, October 2023

We Don't Actually Know If AI Is Taking Over Everything
The Atlantic, October 2023

Klobuchar Says AI Regulation Still Possible Before End of Year
Bloomberg, October 2023

Why everyone seems to disagree on how to define artificial general intelligence
Fast Company, October 2023

OpenAI Is Human After All: Sharing Is Caring, Researchers Tell Model Developers
The Information, October 2023

How transparent are AI models? Stanford researchers found out
VentureBeat, October 2023

Arvind Narayanan & Sayash Kapoor
TIME, September 2023

Newsletter helped us dissect fake claims about AI in real-time
The Indian Express, September 2023

Prominent AI fairness advocates among Princeton AI luminaries
The Daily Princetonian, September 2023

OpenAI Worries About What Its Chatbot Will Say About People's Faces
The New York Times, July 2023

GPT-4: Is the AI behind ChatGPT getting worse?
New Scientist, July 2023

Tips for Investigating Algorithm Harm and Avoiding AI Hype
Global Investigative Journalism Network, July 2023

Six tips for better coding with ChatGPT
Nature News, June 2023

The White House AI R&D Strategy Offers a Good Start. Here's How to Make It Better
Tech Policy Press, May 2023

The AI backlash is here. It's focused on the wrong things
Washington Post, April 2023

What is needed instead of an AI moratorium
Tagesspiegel Background, March 2023

Here are 5 reasons people are dunking on that call for a 6-month A.I. development pause
Fortune, March 2023

The AI factions of Silicon Valley
Semafor, March 2023

Why exams intended for humans might not be good benchmarks for LLMs like GPT-4
VentureBeat, March 2023

ChatGPT: AI bot 'saw test papers before acing exams'
The Times (UK), March 2023

ChatGPT is used to automate bullshit
Süddeutsche Zeitung, March 2023

What if ChatGPT could make us less gullible?
World Economic Forum, February 2023

7 problems facing Bing, Bard, and the future of AI search
The Verge, February 2023

The reproducibility issues that haunt health-care AI
Nature, January 2023

Hello, ChatGPT—Please Explain Yourself!
IEEE Spectrum, December 2022

OpenAI's ChatGPT Is Seen as a Path-breaking Chatbot for AI–But Experts Are Not Impressed
Indian Express, December 2022

The Artificial Intelligence Field Is Infected With Hype
Los Angeles Times, October 2022

Will AI Make Artists Obsolete?
Prospect Magazine, October 2022

Scientists are sloppy with machine learning
NRC, August 2022

Dangerous overoptimism
Frankfurter Allgemeine Zeitung, August 2022

Sloppy Use of Machine Learning Is Causing a 'Reproducibility Crisis' in Science
WIRED, August 2022

Could Machine Learning Fuel a Reproducibility Crisis in Science?
Nature, July 2022

Selected Talks

The Single Smartest Case Against AI Doom
Plain English with Derek Thompson. Podcast. September 2026.

More AI researchers warn of AI's threat to humanity
NBC News. Interview. September 2026.

AI doomerism: will artificial intelligence spiral out of control?
Studio 2 / WHYY. Interview. September 2026.

Hugging Face attack: Can 'rogue agents' bring about AI Armageddon? · Video
In Focus / The Hindu. Podcast. September 2026.

Discussion of autonomous agents and open-ended AI research · Episode recap
MTS Live. Interview. August 2026.

The Missing Science of AI Evaluation · Recording
Princeton University / CITP. Doctoral defense. July 2026.

Computer scientist on why he believes mass layoffs due to AI is a 'convenient excuse'
Morning Edition / NPR. Interview. June 2026.

The peril and opportunity of artificial superintelligence
1A / WAMU / NPR. Podcast. June 2026.

'Hard Fork' Live, Part 3: Differing Visions of an A.I. Future
Hard Fork / The New York Times. Podcast. June 2026.

AI as Normal Technology
SEC AI CoP Seminar. Invited talk. May 2026.

AI as Normal Technology
Massey Cancer Center. Invited talk. May 2026.

The Missing Science of AI Evaluation
Stanford Computer Science. Job talk. April 2026.

Is AI Good for Science?
DC Industrial Policy Salon. Invited talk. April 2026.

Sayash Kapoor on Claude Mythos as normal technology
AI Summer. Podcast. April 2026.

The Missing Science of AI Evaluation
University of Maryland Computer Science. Invited talk. March 2026.

Debunking AI's "Existential Risk"
Factually! with Adam Conover. Podcast. March 2026.

The Missing Science of AI Evaluation
NYU Computer Science Colloquium. Invited talk. February 2026.

The Missing Science of AI Evaluation
Chicago Booth Center for Applied AI. Job talk. February 2026.

The Missing Science of AI Evaluation
UC Berkeley School of Information. Job talk. February 2026.

AI as Normal Technology
Future Knowledge / Internet Archive. Podcast. February 2026.

The Missing Science of AI Evaluation
Johns Hopkins Computer Science. Invited talk. January 2026.

The Missing Science of AI Evaluation
Johns Hopkins School of Government and Policy. Job talk. January 2026.

Holistic Agent Leaderboard
Capital One. Invited talk. November 2025.

AI as Normal Technology
George Washington University. Guest lecture. November 2025.

Holistic Agent Leaderboard
UK AI Security Institute / DSIT. Invited talk. November 2025.

Legal AI Snake Oil
Yale Law School. Invited talk. October 2025.

How To AI: A Practical Business Q&A With Three Experts
On with Kara Swisher. Podcast. September 2025.

AI as Normal Technology
Mastercard. Invited talk. September 2025.

AI Snake Oil
AmCham Lab Global. Invited talk. August 2025.

Moving the AGI Goal Posts: AI Skepticism with Sayash Kapoor
Scaling Laws / Lawfare. Podcast. July 2025.

Talk and discussion on AI as Normal Technology, HAL, and AI research automation
Google DeepMind Foundational Research Unit. Invited talk. July 2025.

Beyond Superintelligence: A Realist's Guide to AI
Interpreting India / Carnegie India. Podcast. July 2025.

Artificial Intelligence as a Normal Technology · Forum website
Global Forum at Casa de Campo / FUNGLODE. Invited talk. July 2025.

AI as a Normal Technology
Google DeepMind AGI & Society speaker series. Invited talk. June 2025.

AI Snake Oil
Princeton University Press Author Talk. Book Talk. June 2025.

AI agents and the law
Vista Institute for AI Policy. Invited talk. June 2025.

AI as Normal Technology
Montclair. Invited talk. June 2025.

AI as Normal Technology
World Innovation, Technology and Services Alliance (WITSA). Invited talk. June 2025.

AI as Normal Technology
American Enterprise Institute. Invited talk. May 2025.

AI Snake Oil
Science, Technology, and Policy fellows, Washington DC. Invited talk. May 2025.

AI and National Security
Department of Defense. Panel. May 2025.

AI as Normal Technology
Two Sigma. Invited talk. May 2025.

AI Snake Oil
Washington D.C. Book Talk. May 2025.

AI as Transformative or Normal Technology
Americans for Responsible Innovation. Debate with Eli Lifland. May 2025.

AI as Normal Technology
RAND. Panel. May 2025.

Building and evaluating AI Agents That Matter
AI Safety Institute (AISI). Invited talk. April 2025.

AI Snake Oil
NASA Goddard. Invited talk. April 2025.

AI Snake Oil
NPR. Interview. April 2025.

AI Snake Oil
Zelen Symposium. Invited talk. April 2025.

AI Snake Oil
Baltimore Library. Book Talk. April 2025.

AI's Impact on Science, Law, and Society
Berkman Klein Center (BKC), Harvard. Invited talk. April 2025.

Open source AI: Risks, Benefits, Interventions
Wilson Center. Policy course. March 2025.

AI Snake Oil
Section School. Book Talk. March 2025.

AI Snake Oil
Loyola University. Invited talk. March 2025.

AI Snake Oil
Schmidt Sciences. Invited talk. March 2025.

Princeton GradFutures science and policy careers
Princeton University. Panel. March 2025.

AI Snake Oil
Oxford Internet Institute (OII). Invited talk. March 2025.

Building and evaluating AI Agents That Matter
AWS. Invited talk. February 2025.

AI Snake Oil
RightsCon. Panel. February 2025.

Building and evaluating AI Agents That Matter
Snowflake. February 2025.

Building and evaluating AI Agents That Matter
Keynote, AI Engineer Conference. February 2025.

Open source AI: Risks, Benefits, Interventions
Wilson Center course for federal staffers in the U.S. government, February 2025.

Can AI automate science?
University of Illinois Generative AI and Future of Research Speaker Series. Invited talk. February 2025.

AI Snake Oil
Brussels Privacy Hub Winter Academy: AI & Law Crossroads. Keynote. February 2025. · Program

Is Generative AI a Threat to Democracy?
GETTING-Plurality, January 2025.

AI Snake Oil
Sony AI Ethics Talk, January 2025.

AI Snake Oil
9^{th} Harris Miller Book Talk, Washington D.C. January 2025.

Health Wanted: AI in Health Care
Health Wanted / WABE / Emory University. Podcast. November 2024.

AI Snake Oil
The Prompt Podcast, Denmark. Podcast. November 2024.

On the Societal Impact of Open Foundation Models
University of Rochester AI Policy and Regulation Workshop. Invited talk. November 2024.

AI Snake Oil
Fidelity. Book Talk. November 2024.

AI Snake Oil
Princeton University GradFutures Responsible AI course. Book Talk. November 2024.

Promises and Pitfalls of AI in law
Law and Technology Centre, HKU. Invited talk. October 2024.

AI Snake Oil
Airbnb. Book Talk. October 2024.

Is AI-generated disinformation a threat to democracy?
Global Summit on the Future of Free Speech. Invited talk. October 2024.

AI Snake Oil
Princeton Public Library. Book Talk. October 2024.

Types of AI and AI Snake Oil
AAAS Center for Scientific Evidence in Public Issues. Policy seminar. October 2024.

Open source AI and its policy implications
Wilson Center (Executive staffers). Policy course. October 2024.

AI Snake Oil
Adam Conover's Factually! Podcast. Book Podcast. October 2024.

AI Agents That Matter
Weaviate Podcast. Podcast. September 2024.

AI Snake Oil
AI & Social Sciences Seminar, Paris. Book Talk. September 2024.

AI Snake Oil
Eric Topol's Ground Truths Podcast. Book Podcast. September 2024.

AI Snake Oil
City Lights. Book Talk. September 2024.

The threat of existential risk from AI
Machine Learning Street Talk Podcast. Podcast. August 2024.

Types of AI and AI Snake Oil
AAAS Center for Scientific Evidence in Public Issues. Policy seminar. August 2024.

Open source AI and its policy implications
Wilson Center (Legislative staffers). Policy course. August 2024.

A Safe Harbor for AI Evaluation and Red Teaming
Federation of American Scientists. Congressional briefing. July 2024.

On the Societal Impact of Open Foundation Models
ICML Oral. Conference oral presentation. July 2024.

AI agents that matter
Meta (Core Applied Sciences). Invited talk. May 2024.

AI and disinformation
Dutch Ministry of Interior and Kingdom Relations workshop. Invited talk. May 2024.

On the Societal Impact of Open Foundation Models
Toronto AI Safety group. Invited talk. May 2024.

Understanding and Unlocking AI's Economic Potential
World Bank Measuring Development 2024. Panel. May 2024.

Princeton Dialogues in AI
Senate AI Caucus. April 2024.

Princeton Dialogues in AI
House AI Caucus. April 2024.

On the Societal Impact of Open Foundation Models
Stanford RegLab. Invited talk. April 2024.

On the Societal Impact of Open Foundation Models
Mechanism Design For Social Good Speaker Series. Invited talk. April 2024.

On the Societal Impact of Open Foundation Models
World Innovation, Technology and Services Alliance. Invited talk. April 2024.

On the Societal Impact of Open Foundation Models
Tech Policy Press. Podcast. April 2024.

Assessing the risks of open models
This Week in Machine Learning. Podcast. March 2024.

On the Societal Impact of Open Foundation Models
Safe Mode. Podcast. March 2024.

Intro to AI/ML for Regulators
Consumer Financial Protection Bureau. Invited talk. March 2024.

On the Societal Impact of Open Foundation Models
Princeton Alignment Reading Group. Invited talk. February 2024.

Against Predictive Optimization
Cornell University. Guest lecture. February 2024.

Understanding AI Hype
Symphony AI. Invited talk. February 2024.

Against Predictive Optimization
Stanford University Fairness Lunch Speaker Series. Invited talk. February 2024.

On the Societal Impact of Open Foundation Models
Stanford Workshop on Governance of Open Foundation Models. Panel. February 2024.

Beyond the AI hype
Government of Canada's Federal Foresight Network. Panel. March 2024.

How to Prepare for the Deluge of Generative AI on Social Media
Federal Trade Commission. Invited talk. December 2023.

Launch of NTIA's Public Consultation Process on Widely Available AI Foundation Model Weights
Center for Democracy and Technology. Panel. December 2023.

Data Governance in the Age of AI
Washington D.C. Panel. December 2023.

National Association of Attorneys General
Washington D.C. Panel. November 2023.

AI and its hazards for science
ScienceWriters Conference, University of Colorado, Boulder. Invited talk. October 2023.

How to detect AI hype
Princeton University Press. Invited talk. October 2023.

Tigers on Strike
Princeton University. Panel. September 2023.

Responsible and Open Foundation Models
Princeton-Stanford. Workshop organizer and panel moderator. September 2023.

Improving Reproducibility, Trustworthiness and Fairness in Machine Learning
ICIAM Minisymposium, Tokyo. Invited talk. August 2023.

Investigating algorithmic harm: Best practices and hard-learned lessons
Investigative Reporters and Editors, Orlando. Panel. June 2023.

Against Predictive Optimization
ACM FAccT, Chicago. Paper talk. June 2023.

CITP Digital Investigators Conference
Princeton University. Invited talk. May 2023.

Critical voices on AI
Birkbeck Institute of Data Analytics. Invited talk. May 2023.

Co-opting AI: Language
New York University. Invited talk. April 2023.

Reproducibility in AI-based science
Royal Society / UK Reproducibility Network (UKRN). Panel. April 2023.

Data (Re)Makes the World
Yale Law School. Panel. April 2023.

How to spot AI hype
Yale Quantum Institute, Physics Open Mic. Invited talk. March 2023.

Reproducibility Crisis in Machine Learning Based Science
Keynote. November 2022.

Leakage and the Reproducibility Crisis in ML-based Science
University of Exeter, Institute for Data Science and Artificial Intelligence. Invited talk. November 2022.

Data Science Institute seminar
Lawrence Livermore National Lab. Invited talk. October 2022.

5th Annual conference of the Massive Analysis and Quality Control Society
FDA headquarters. Invited talk. September 2022.

The reproducibility crisis in ML-based science
ML reproducibility workshop. Workshop. August 2022.

Workshop on The Reproducibility Crisis in ML-based Science
Princeton University. Opening talk. July 2022.