

Joon Sung Park

Ph.D. Student, Computer Science Department
Stanford University

joonspk@stanford.edu
www.joonsungpark.com
353 Jane Stanford Way, Room #360, Stanford CA 94305

Research Interests

Human-computer interaction; natural language processing; social computing;
generative agents; AI auditing; human-AI interaction

Education

- Stanford University**, Stanford, CA 2020 ~
Ph.D., Computer Science
Advisors: Michael S. Bernstein and Percy Liang
- University of Illinois at Urbana-Champaign (UIUC)**, Urbana, IL 2018 ~ 2020
Master of Science, Computer Science
Advisor: Karrie Karahalios
- Swarthmore College**, Swarthmore, PA 2013 ~ 2017
Bachelor of Arts, Computer Science
Advisor: Tia Newhall

Awards and Honors

- 2022 Microsoft Research Ph.D. Fellowship
Awarded in support of graduate studies at Stanford University. Given to candidates pursuing a Ph.D. in computing-related areas. One of the 13 awardees in North America.
Method Award, ACM CSCW
Best Paper Award, ACM CHI
- 2021 Best Paper Nominated, ACM ASSETS
Hasso Plattner Design Thinking Research Grant
Awarded \$110,000 in support of research activities.
Best Paper Honorable Mention, ACM CHI
- 2020 Terry Winograd Fellowship
Awarded by Stanford School of Engineering in support of graduate studies.
Best Paper Honorable Mention, ACM CHI
- 2019 Siebel Scholar Award
Awarded \$35,000. Given to the top graduate students in the world's leading graduate schools for their academic excellence and demonstrated leadership.
- 2016 Eugene M. Lang Research Fellowship

Academic Publications

Top-tier venues for human-computer interaction research include CHI, UIST, and CSCW. In 2018, CSCW transitioned to a hybrid journal structure. The titles of papers currently under review are edited to maintain authors' anonymity.

Refereed Conference and Journal Papers

- [28] Tony Lee[†], Michihiro Yasunaga[†], Chenlin Meng[†], Yifan Mai, **Joon Sung Park**, Agrim Gupta, Yunzhi Zhang, Deepak Narayanan, Hannah Benita Teufel, Marco Bellagente, Minguk Kang, Taesung Park, Jure Leskovec, Jun-Yan Zhu, Li Fei-Fei, Jiajun Wu, Stefano Ermon, Percy Liang. Holistic Evaluation of Text-to-Image Models (NeurIPS '23 – In Press).
Spotlight
- [27] Mina Lee, Megha Srivastava, Amelia Hardy, Esin Durmus, Ashwin Paranjape, John Thickstun, Ines Gerard-Ursin, Faisal Ladhak, Frieda Rong, Rose E. Wang, Xiang Lisa Li, Minae Kwon, **Joon Sung Park**, Hancheng Cao, Tony Lee, Rishi Bommasani, Michael Bernstein, Percy Liang. Evaluating Human-Language Model Interaction. Transactions on Machine Learning Research (TMLR '23 – In Press).
- [26] **Joon Sung Park**, Joseph C. O'Brien, Carrie J. Cai, Meredith Ringel Morris, Percy Liang, Michael S. Bernstein. Generative Agents: Interactive Simulacra of Human Behavior. The 36th Annual ACM Symposium on User Interface Software and Technology (UIST '23 – In Press).
- [25] Kelly Mack, Maitraye Das, Dhruv Jain, Danielle Bragg, John Tang, Andrew Begel, Erin Beneteau, Josh Urban Davis, Abraham Glasser, **Joon Sung Park**, Venkatesh Potluri. 2023. Mixed Abilities and Varied Experiences: A Group Auto-Ethnography of a Virtual Summer Internship. Communications of the ACM (CACM '23).
- [24] Elise Karinshak, Sunny Liu, **Joon Sung Park**, Jeffrey Hancock. 2023. Working With AI to Persuade: Examining a Large Language Model's Ability to Generate Pro-vaccination Messages. Proceedings of the ACM on Human-Computer Interaction (CSCW '23).
- [23] **Joon Sung Park**, Lindsay Popowski, Carrie J. Cai, Meredith Ringel Morris, Percy Liang, Michael S. Bernstein. 2022. Social Simulacra: Creating Populated Prototypes for Social Computing Systems. Proceedings of the Annual ACM Symposium on User Interface Software and Technology (UIST '22).
- [22] **Joon Sung Park**, Joseph Seering, Michael S. Bernstein. 2022. Measuring the Prevalence of Anti-Social Behavior in Online Communities. Proceedings of the ACM on Human-Computer Interaction (CSCW '22).
Method Award
- [21] **Joon Sung Park**, Karrie Karahalios, Niloufar Salehi, Motahhare Eslami. Power Distribution and Conflicts in Designing, Using, and Maintaining Algorithmic Socio-Technical Systems. 2022. Proceedings of the ACM on Human-Computer Interaction (CSCW '22).
- [20] Mitchell Gordon, Michelle Lam, **Joon Sung Park**, Kayur Patel, Jeffrey Hancock, Tatsunori Hashimoto, Michael S. Bernstein. 2022. Jury Learning: Integrating Dissenting Voices into Machine Learning Models. Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '22).
Best Paper Award

- [19] Danaë Metaxa, **Joon Sung Park**, Ronald E. Robertson, Karrie Karahalios, Christo Wilson, Jeffrey Hancock, Christian Sandvig. 2021. Auditing Algorithms: Understanding Algorithmic Systems from the Outside In. Foundations and Trends in Human-Computer Interaction (FnT' 21).
- [18] Kelly Mack, Maitraye Das, Dhruv Jain, Danielle Bragg, John Tang, Andrew Begel, Erin Beneteau, Josh Urban Davis, Abraham Glasser, **Joon Sung Park**, Venkatesh Potluri. 2021. Mixed Abilities and Varied Experiences: A Group Auto-Ethnography of a Virtual Summer Internship. Proceedings of the ACM Conference on Computers and Accessibility (ASSETS '21).
Best Paper Nominated
- [17] **Joon Sung Park**, Michael S. Bernstein, Robin N. Brewer, Ece Kamar, Meredith Ringel Morris. 2021. Understanding the Representation and Representativeness of Age in AI Data Sets. Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society (AIES '21).
- [16] Vinay Koshy, **Joon Sung Park**, Ti-Chung Cheng, Karrie Karahalios. 2021. "We Just Use What They Give Us": Understanding Passenger User Perspectives in Smart Homes. Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '21).
Best Paper Honorable Mention
- [15] **Joon Sung Park**, Danielle Bragg, Ece Kamar, Meredith Ringel Morris. 2021. Designing an Online Infrastructure for Collecting AI Data From People With Disabilities. Proceedings of the Conference on Fairness, Accountability, and Transparency (FAccT '21).
- [14] Mufan Luo, Tiffany Hsu, **Joon Sung Park**, Jeffrey Hancock. 2020. Emotional Amplification During Live-Streaming: Evidence from Comments During and After News Events. Proc. ACM Hum.-Comput. Interact. (CSCW '20).
- [13] **Joon Sung Park**. 2020. Re-Examining the Role of Time in Human-Algorithm Interaction. University of Illinois at Urbana-Champaign (Thesis '20).
- [12] Sijia Xiao, Danaë Metaxa, **Joon Sung Park**, Karrie Karahalios, Niloufar Salehi. 2020. Random, Messy, Funny, Raw: Finstas as Intimate Reconfigurations of Social Media. Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '20).
Best Paper Honorable Mention
- [11] Danaë Metaxa, **Joon Sung Park**, James Landay, Jeffrey Hancock. 2019. Search Media and Elections: A Longitudinal Investigation of Political Search Results in the 2018 U.S. Elections. Proc. ACM Hum.-Comput. Interact. 3, Article 129, 17 pages (CSCW '19).
- [10] **Joon Sung Park**, Rick Barber, Alex Kirlik, Karrie Karahalios. 2019. A Slow Algorithm Improves Users' Assessments of the Algorithm's Accuracy. Proc. ACM Hum.-Comput. Interact. 3, Article 102, 15 pages (CSCW '19).

White Papers and Reports

- [9] **Joon Sung Park**, Chris Donahue, Mina Lee, Siddharth Karamcheti, Dorsa Sadigh, Michael S. Bernstein. 2021. On the Opportunities and Risks of Foundation Models: §2.5 Interaction (Report '21).
- [8] Pingjing Yang, Sajjadur Rahman, Mangesh Bendre, **Joon Sung Park**, Kevin Chang, Karrie Karahalios, Aditya Parameswaran. 2018. "It's a grid, it's a form, it's a spreadsheet!": Revisiting

Sensemaking with Spreadsheets for Presentation and Data Analysis. Technical Report. September 2018.

Organized Workshops and SIGS

- [7] Michael S. Bernstein, **Joon Sung Park**, Meredith Ringel Morris, Saleeman Amershi, Lydia Chilton, Mitchell L. Gordon. 2023. Architecting Novel Interactions With Generative AI Models. The Annual ACM Symposium on User Interface Software and Technology – Workshop (UIST Workshop '23).
- [6] Hong Shen, Tianshi Li, Toby Jia-Jun Li, **Joon Sung Park**, Diyi Yang. 2023. Shaping the Emerging Norms of Using Large Language Models in Social Computing Research. Computer-Supported Cooperative Work & Social Computing – Special Interest Groups (CSCW SIG '23).

Other Juried Papers (Nonarchival Papers and Posters)

- [5] Elise Karinshak, Sunny Liu, **Joon Sung Park**, Jeffrey Hancock. 2022. Can AI Persuade? Examining a large language model's ability to generate pro-vaccination messages. In the Annual International Communication Association Conference (ICA '22).
- [4] Mufan Luo, Tiffany Hsu, **Joon Sung Park**, Jeffrey Hancock. 2020. Live Streaming Intensifies Emotions: Evidence from Live and Retrospective Comments on YouTube. In the Annual International Communication Association Conference (ICA '20).
- [3] Danaë Metaxa, **Joon Sung Park**, James Landay, Jeffrey Hancock. 2019. Search Media and Elections: Investigating Partisanship in Political Search Results. In the Annual International Communication Association Conference (ICA '19).

Media Publications

- [2] **Joon Sung Park**, Rick Barber, Alex Kirlik, Karrie Karahalios. (2019, December) Faster is not always better: benefits of slower interaction with our algorithmic systems. ACM CSCW Medium.
- [1] Jeffrey Hancock, Danaë Metaxa, **Joon Sung Park**. (2018, September) Are Google search results politically biased? The Guardian.

Art Exhibitions

- 2017 **Evanescant Spaces**. Solo Exhibition at the List Gallery
- 2016 **Refinery and Sunflowers**. Group Exhibition at the Kitao Gallery

Research Experience

- 2020/9 ~ **Stanford University**, Stanford, CA
Research Assistant
Mentors: Michael S. Bernstein and Percy Liang
- 2020/5 ~ 8 **Microsoft Research**, Redmond, WA
Research Intern
Mentors: Meredith Ringel Morris, Ece Kamar, and Danielle Bragg
- 2018 ~ 2020 **University of Illinois at Urbana-Champaign**, Urbana, IL

- Research Assistant
Mentors: Karrie Karahalios
- 2018/2 ~ 8 **Stanford University**, Stanford, CA
Research Assistant
Mentors: James Landay and Jeff Hancock
- 2018/6 ~ 8 **Swarthmore College**, Swarthmore, PA
Research Assistant
Mentors: Bryce Wiedenbeck

Selected Software

- 2018/9 ~ 10 **commPAIR.io**, Eugene, OR
Full-stack Developer
- 2015 ~ 2016 **Willow-Cutting Distributed Hash Table (WCDHT)**, Swarthmore, PA
Main Developer

Other Experience

- 2016 ~ 2017 **Galamingo.com**, Palo Alto, CA
Co-founder
- 2014/6 ~ 8 **National Assembly of the Republic of Korea**, Seoul, South Korea
Secretary to a Member of National Assembly
- 2014/1 ~ 5 **Philadelphia Museum of Art**, Philadelphia, PA
Student Ambassador, Marketing and Communications Division

Invited Presentations and Panels

Research talks

- [r.31] *Simulating Human Behavior With Generative Agents*
TED Talk (TED AI), San Francisco, CA. October 17, 2023
- [r.30] *Generative Agents: Interactive Simulacra of Human Behavior*
University of Sydney. Internal Seminar. Virtual, CA. September 14, 2023
- [r.29] Apple. Internal Seminar. Cupertino, CA. August 24, 2023
- [r.28] Riot Games. Internal Seminar. San Francisco, CA. August 11, 2023
- [r.27] London Machine Learning Meetup, Virtual. August 2, 2023
- [r.26] Microsoft Research, Montreal Seminar. August 1, 2023
- [r.25] Tencent. Internal Seminar. Virtual. July 11, 2023
- [r.24] SelectStar. Public Seminar. Virtual. July 6, 2023
- [r.23] The Bank of England. New Methods, Perspectives, and Approaches. June 20, 2023
- [r.22] Simulation Summit, San Francisco, CA. June 8, 2023
- [r.21] InfoSec Seminar, University College, London. Virtual. May 25, 2023
- [r.20] University of California, Berkeley. MARL Seminar. Berkeley, CA, May 19, 2023
- [r.19] Reddit, LLM-Powered NPCs: Chat with Stanford's Joon Sung Park. Virtual. May 17, 2023
- [r.18] IBM, Accelerated Discovery Seminar. Virtual, May 3, 2023

- [r.17] Google, Internal Seminar. Mountain View, CA, May 3, 2023
- [r.16] Chat LLM. Sydney, Australia, April 27, 2023
- [r.15] Deep Learning Seminar. Krafton, Seoul, South Korea, April 27, 2023
- [r.14] University of California, Berkeley. BiD Seminar. Berkeley, CA, April 11, 2023
- [r.13] HAI-Google Lightning Talk. Stanford, CA, March 23, 2023
- [r.12] Interactive Comp. Semi., University of California, San Diego. La Jolla, CA, February 15, 2023
- [r.11] Generated Worlds & Fictional Embodiments. San Francisco, CA, December 15, 2022
- Social Simulacra: Creating Populated Prototypes for Social Computing Systems***
- [r.10] Google Brain, People + AI Research (PAIR). Virtual, October 19, 2022
- [r.9] A New Era of HAI-Google Collaborations. Virtual, October 18, 2022
- [r.8] HPI-HPDTRP Community Workshop. Potsdam, Germany, September 15, 2022
- [r.7] Simons Institute for the Theory of Computing. Berkeley, CA, July 27, 2022
- [r.6] HPDTRP Community Workshop. Stanford, CA, May 24, 2022
- [r.5] CRFM Research Update. Virtual, May 3, 2022
- [r.4] HAI-Google Research Token Update. Virtual, April 21, 2022
- [r.3] CRFM Spotlight Talk. Virtual, February 8, 2022
- Creating Inclusive AI Datasets for People With Disabilities and the Oldest Old***
- [r.2] Microsoft Research HCI Seminar. Redmond, WA, August 14, 2020
- Modeling a Multi-Agent Environment with Varying Degrees of Competitiveness***
- [r.1] Sigma Xi Research Presentation. Swarthmore, PA, August 12, 2016

Panels

- AI Town***
- [p.9] Andreessen Horowitz (a16z). San Francisco, CA, September 21, 2023
- New Horizons: Science, Creativity, and Society in the Age of GenAI***
- [p.8] 2023 HAI Fall Conference, Stanford, CA. October 24, 2023
- Will the next generation LLM borrow insights from game research?***
- [p.7] TED AI – After Party, San Francisco, CA. October 18, 2023
- Path to Generative Agents (That Work)***
- [p.6] Zhen x Google. San Francisco, CA, August 10, 2023
- Opening Panel for Founder Summit***
- [p.5] Gradient Ventures. Redwood City, CA, June 1, 2023
- Games x AI Roundtable***
- [p.4] Andreessen Horowitz (a16z). San Francisco, CA, May 23, 2023
- Simulations, AI Memory and AI Agents***
- [p.3] AISF – an AI Week for Builders. San Francisco, CA, May 19, 2023
- The Future of Social Chatbots***
- [p.2] Foundation Capital, Generative AI Unconference. San Francisco, CA, May 10, 2023
- Ph.D. Research and Career***
- [p.1] Stanford University CURIS Program. Stanford, CA, June 30, 2021

Guest Lectures

Generative Agents

- [gl.5] Stanford Computational Sociology Workshop. October 19, 2023
- [gl.4] Stanford Artificial Intelligence and Deliberative Democracy. September 14, 2023
- [gl.3] Stanford HAI x Commonwealth Bank of AUS. September 13, 2023
- [gl.2] Stanford HAI: Accenture Foundation Model Scholar Program. July 10, 2023

Architecting Proxies of Our Lives in Simulations

- [gl.1] CS329X Human-Centered NLP, Stanford University. Stanford, CA, May 17, 2023

Other Appearances

- [oa.10] ***Fireside Chat: Generative Agents and Open Source***
Virtual, October 20, 2023
- [oa.9] ***TED AI Hackathon: Judge***
TED AI. San Francisco, CA, October 15, 2023
- [oa.8] ***Speaker Series: Generative Agents***
Sutter Hill Ventures. Palo Alto, CA, September 28, 2023
- [oa.7] ***Fireside Chat: AI Studio***
Unusual Ventures. Menlo Park, CA, July 20, 2023
- [oa.6] ***Fireside Chat: The Agents Are Coming***
Scale Ventures. San Francisco, CA, June 27, 2023
- [oa.5] ***Guest Appearance: Generative Agents***
The AutoML Podcast with Adam Boaz Becker. June 15, 2023
- [oa.4] ***Guest Appearance: Generative Agents and Human-Computer Interaction***
The Gradient with Daniel Bashir. June 15, 2023
- [oa.3] ***Guest Appearance: : Modeling Human Behavior with Generative Agents***
The TWIML AI Podcast with Sam Charrington. June 5, 2023
- [oa.2] ***Guest Appearance: SBS News***
SBS News. May 19, 2023
- [oa.1] ***Guest Appearance: NBC News***
NBC News with Jessica Aguirre. April 12, 2023

Professional Activities and Service**Program Committee**

- 2022 FAccT
- ACL, Workshop on Intelligent and Interactive Writing Assistants 2022

Reviewer

- CHI 2020, 2021, 2022, 2023
- CSCW 2020, 2021, 2022, 2023
- UIST 2022, 2023
- NeurIPS 2021, 2022, 2023
- ICML, AI & HCI Workshop 2023
- ICWSM 2021, 2022, 2023

CHI Case Studies	2021
The Web Conference	2021
IMWUT	2022
New Media & Society	2021

Student Volunteer

CSCW 2019 (also won the SV t-shirt design competition)

Departmental Service at Stanford University

CS Ph.D. Admissions Committee

2021-2022, 2022-2023

Center for Research on Foundation Models (CRFM) Designer

2021

CS Ph.D. Mentor for Underrepresented Students

2020-2021, 2021-2022

CS Ph.D. Student-Applicant Support Program (SASP) Reviewer

2020, 2021

CSLI Application Reviewer

2021, 2022

Students Mentored

Ph.D. Students

- 2022 **Bobby Yan.** Stanford University.
Worked on the interactive simulacra project.
- 2021 **Lindsay Popowski.** Stanford University.
Co-authored [23]. Worked on the social simulacra project.

Undergraduate Students

- 2023 **Carolyn Zou.** Northwestern University.
Worked on the GABM project.
- Garbo Chung.** Stanford University.
Worked on the GABM project.
- 2022 **Joseph Coleman O'Brien.** Stanford University.
Co-authored [26]. Worked on the generative agents project.
- Phuc Tran.** Stanford University.
Worked on the self-fulfilling prophecy project.
- 2021 **Fern Limprayoon** (co-mentored w/ Ranjay Krishna and Mitchell Gordon), Carnegie Mellon University.
Afterwards: Ph.D at Yale University advised by Brian Scassellati
Worked on the self-fulfilling prophecy project.
- Nicole Tong.** Stanford University.

- 2019 Worked on the concept decomposition project.
Elizabeth (Liza) Gak. Washington University in St. Louis.
Afterwards: Ph.D. at UC Berkeley advised by Niloufar Salehi
 Worked on Finstagram grant and research protocol design.

Selected Media Coverage

- 2023 [Computational Agents Exhibit Believable Humanlike Behavior](#)
Stanford HAI. Katharine Miller. September 21, 2023
- [‘They went to the bar at noon’: what this virtual AI village is teaching researchers](#)
Nature. Matthew Hutson. NEWS Q&A. September 20, 2023
- [How AI agents are already simulating human civilization](#)
Venture Beat. Ben Dickson. September 11, 2023
- [What can a virtual village made up of AI chatbots tell us about human interaction?](#)
Financial Times. Oliver Roeder. August 24, 2023
- [ChatGPT Invented This Fake Epidemic. Can It Help Prevent the Next One?](#)
Daily Beast. Tony Ho Tran. August 18, 2023
- [Researchers from Stanford and Google have made an entire AI village. The 25 bots that live there gossip, work, and plan Valentine’s Day parties.](#)
Business Insider. Kai Xiang Teo. August 10, 2023
- [AI Can Give You an NPC That Remembers. It Could Also Get Your Favorite Artist Fired](#)
Wired. Fernanda Seavon. August 7, 2023
- [The NeverEnding Game: How AI Will Create a New Category of Games?](#)
Andreessen Horowitz. Jonathan Lai. July 18, 2023
- [Guinea Pigbots: Doing research with human subjects is costly and cumbersome. Can AI chatbots replace them?](#)
Science. Matthew Hutson. July 13, 2023
- [Cognitive twins will put generative AI to use for businesses. Here’s how it could work](#)
Wired. Susheel Sethumadhavan and Akshay Goel. July 3, 2023
- [ChatGPT Inside a Game](#)
The Kyungbyang Shinmun. Jaeduk Lee. June 5, 2023
- [No more ‘I took an arrow to the knee’: could AI write super-intelligent video game characters?](#)
The Guardian. Will McCurdy. May 28, 2023
- [AI Canon](#)
Andreessen Horowitz. Derrick Harris, Matt Bornstein, and Guido Appenzeller. May 25, 2023
- [Would You Come to the Valentine’s Day Party? An Experiment With ChatGPT Agents](#)
Seoul Broadcasting System (SBS) 8 PM News. May 19, 2023
- [An Experiment With 25 AI Agents](#)
The Nihon Keizai Shimbun (Nikkei). May 16, 2023
- [Exploring Autonomous Agents: A Semi-Technical Dive](#)
Sequoia. Dan Chen. May 11, 2023

[25 ChatGPT AIs Play A Game - So What Happened?](#)

Two Minute Papers. April 24, 2023

[ChatGPT agents are better at simulated role-play than humans](#)

New Scientist. Jeremy Hsu. April 21, 2023

[The lead researcher behind those Sims-like 'generative agents' on the future of AI NPCs](#)

PC Gamer. Tyler Wilde. April 19, 2023

[Hard Fork: How A.I. is bringing us closer to "Westworld."](#)

The New York Times. Kevin Roose and Casey Newton. April 14, 2023

[AI 'humans' in virtual world plot and gossip like us](#)

The Times. Rhys Blakely. April 14, 2023

[These Are Your Sims. And These Are Your Sims on ChatGPT](#)

Yahoo Finance. Jose Antonio Lanz. April 14, 2023

[Stanford Scientists Put ChatGPT Into Video Game Characters And It's Incredible](#)

IFLScience. Jack Dunhill. April 13, 2023

[A Closer Look: A.I. Simulation Takeaways](#)

NBC. Jessica Aguirre. April 12, 2023

[Stanford U & Google's Generative Agents Produce Believable Proxies of Human Behaviours](#)

Synched. Hecate He. April 12, 2023

[This is what happened when 25 AI avatars were let loose in a virtual town](#)

Business Insider. Aaron Mok. April 12, 2023

[AI-driven NPCs are living lives and planning parties all on their own](#)

PC Gamer. Katie Wickens. April 12, 2023

[This Tiny Town Created by ChatGPT Is Better Than Reality TV](#)

Daily Beast. Tony Ho Tran. April 12, 2023

[Exploring the Potential of Generative Agents: Simulating Human Behavior with AI](#)

Hackernoon. stephanbv. April 12, 2023

[Inside Smallville, the wholesome village populated solely by AIs](#)

Dazed. Thom Waite. April 12, 2023

[What if someone mixed The Sims with ChatGPT bots? It would look like this](#)

The Register. Thomas Claburn. April 11, 2023

[Surprising things happen when you put 25 AI agents together in an RPG town](#)

Ars Technica. Benj Edwards. April 11, 2023

[Google Tells AI Agents to Behave Like 'Believable Humans' to Create 'Artificial Society'](#)

Vice. Chloe Xiang. April 11, 2023

[Sul Sul: Study Mixes ChatGPT With 'The Sims' to Delightful Effect](#)

Gizmodo. Thomas Germain. April 11, 2023

[Researchers populated a tiny virtual town with AI \(and it was very wholesome\)](#)

Tech Crunch. Devin Coldewey. April 10, 2023

[Stanford AI Researchers Develop New Algorithm for More Inclusive Content Moderation](#)

The Stanford Daily. Richard Xue. October 4, 2022

- Using Large Language Models to Generate Trolling at Scale
Statistical Modeling, Causal Inference, & Social Science. Jessica Hullman. August 12, 2022
- Who Decides? Dealing with Online Toxic Speech by Selecting the Decision Makers
Stanford HAI. Katharine Miller. June 1, 2022
- Stanford Researchers Propose ‘Jury Learning’ as a Way to Mitigate Bias in AI
VentureBeat. Kyle Wiggers. February 10, 2022
- 2021 The Big Question
Nature Machine Intelligence. Editorials. September 17, 2021
- The Tech and Social Impact of AI’s Emerging Foundation Models
Communications of ACM. The Register. August 26, 2021
- Inside the Mind of Big AI
Financial Times. Richard Waters. August 19, 2021
- 100+ Stanford Researchers Publish 200+ Page Paper on the AI Paradigm Shift Introduced by Large-Scale Models Synced Review.
Synced. August 19, 2021
- AI’s Shaky Foundations
Axios. Bryan Walsh. August 18, 2021
- Foundation Models Risk Exacerbating AI’s Ethical Challenges
VentureBeat. Kyle Wiggers. August 18, 2021
- Tech-Industry AI Is Getting Dangerously Homogenized, Say Stanford Experts
FastCompany. Mark Sullivan. August 18, 2021
- 2019 Search Results Not Biased Along Party Lines, Stanford Scholars Find
Communications of ACM. November 27, 2019
- Search Results Not Biased Along Party Lines, Stanford Scholars Find
Stanford News. Melissa De Witte. November 26, 2019
- Siebel Scholars Foundation Announces Class of 2020
Bloomberg. Jennifer Stern. September 25, 2019
- World’s Brightest Business, Computer Science, Bioengineering and Energy Science Students Join Distinguished Leaders
Business Wire. Jennifer Stern. September 25, 2019