

HAII

2019

Human-Centered
Artificial Intelligence
Symposium

Welcome to the Launch of HAI.

Today's event is a first glimpse of a long-standing vision shared by many Stanford faculty, students and supporters: the Stanford Institute for Human-Centered Artificial Intelligence. HAI is a historic effort to encourage deep, multidisciplinary engagement in one of the most consequential issues of our time.

The world is undergoing profound change, with a future of both excitement and complexity on the horizon. At the heart of so much of this change is AI, a nascent technology that has rapidly become a fixture in both work and everyday life and a subject of both interest and concern for thinkers in a growing range of fields. Of course, no one can control something so diffuse; but as one of the world's foremost academic institutions, Stanford can play a critical role in guiding it responsibly.

This is the mission of HAI. It will serve as a global hub of dialogue and collaboration and demonstrate AI's future as a shared pursuit bringing to bear expertise from engineering to medicine, ethics, law, economics, neuroscience, the creative arts and more.

HAI is founded on three simple but powerful ideas:

1. **For AI to better serve our needs, it must incorporate more of the versatility, nuance, and depth of the human intellect.**
2. **The development of AI should be paired with an ongoing study of its impact on human society and guided accordingly.**
3. **The ultimate purpose of AI should be to enhance our humanity, not diminish or replace it.**

Realizing these goals will be among the greatest challenges of our time. But it's also an opportunity to share ideas, challenge conventions and explore possibilities on an unprecedented scale. Together, we can ensure the next frontier of AI isn't merely technological, but humanistic as well.

We hope you'll join us.



John Etchemendy
Co-Director
Patrick Suppes Family Professor
in the School of Humanities
and Sciences



Fei-Fei Li
Co-Director
Professor of Computer Science

Agenda

9:15–
9:30

Welcome

John Etchemendy, Co-Director, Stanford Institute for Human-Centered AI, Patrick Suppes Family Professor, School of Humanities and Sciences, Stanford University
Marc Tessier-Lavigne, President, Stanford University

9:30–
10

The Stanford Institute for Human-Centered Artificial Intelligence (HAI)

Fei-Fei Li, Co-Director, Stanford Institute for Human-Centered AI, Professor of Computer Science, Co-Founder and Chairperson, AI4ALL, Stanford University

10–
10:15

Break

INTELLIGENCE INSPIRED BY HUMANS

10:15–
10:30

Lightning Talks

Michael Frank, David and Lucile Packard Foundation Professor in Human Biology, Stanford University
Surya Ganguli, Assistant Professor of Applied Physics and, by courtesy, of Neurobiology, Electrical Engineering and Computer Science, Stanford University
Percy Liang, Assistant Professor of Computer Science and, by courtesy, of Statistics, Stanford University

10:30–
11:30

Panel Discussion

MODERATOR: **Reid Hoffman**, Co-Founder, LinkedIn; Partner, Greylock Partners
Jeff Dean, Senior Google Fellow, Senior Vice President of Google AI, Google; Distinguished Fellow, HAI
Alison Gopnik, Professor of Psychology, Affiliate Professor of Philosophy, University of California, Berkeley
Demis Hassabis, Co-Founder, DeepMind; Distinguished Fellow, HAI
Chris Manning, Thomas M. Siebel Professor in Machine Learning, Professor of Linguistics and of Computer Science, Director, Stanford Artificial Intelligence Laboratory (SAIL), Stanford University

11:30–
12:45

Lunch

12:45–
1:45

Keynote: A Conversation with Bill Gates

Bill Gates, Founder, Microsoft Corporation
Amy Jin, AI4ALL Alum; Student, Harvard University
Stephanie Tena-Meza, AI4ALL Alum; Student, Salinas High School

AI'S HUMAN AND SOCIETAL IMPACT

1:45–
2:05

Lightning Talks

Juliana Bidadanure, Assistant Professor of Philosophy and, by courtesy, of Political Science, Faculty Director, Stanford Basic Income Lab, Stanford University
Mark Duggan, The Trione Director of the Stanford Institute for Economic Policy Research and Wayne and Jodi Cooperman Professor of Economics, Stanford University
Jennifer Pan, Assistant Professor of Communication and, by courtesy, of Political Science and Sociology, Stanford University
David Engstrom, Professor of Law, Associate Dean, and Bernard D. Bergreen Faculty Scholar, Stanford Law School, Stanford University
Daniel Ho, William Benjamin Scott and Luna M. Scott Professor of Law, Stanford Law School, Professor, by courtesy, of Political Science and Senior Fellow, Stanford Institute for Economic Policy Research, Stanford University

2:05–
3:05

Panel Discussion

MODERATOR: **James Manyika**, Chairman and Director, McKinsey Global Institute, Senior Partner, McKinsey & Company; Distinguished Fellow, HAI
Susan Athey, The Economics of Technology Professor, Stanford Graduate School of Business, and Senior Fellow, Stanford Institute for Economic Policy Research, Stanford University
Erik Brynjolfsson, Schussel Family Professor of Management, MIT, and Director, MIT Initiative on the Digital Economy, MIT Sloan School of Management; Visiting Scholar, Stanford University, and Distinguished Fellow, HAI
Kate Crawford, Co-Director and Co-Founder, AI Now Institute, NYU, Distinguished Research Professor, NYU; Principal Researcher, Microsoft Research New York City
Tristan Harris, Co-Founder and Executive Director, The Center for Humane Technology

3:05–
3:35

Break

AUGMENTING HUMAN CAPABILITIES

3:35–
3:55

Lightning Talks

Michael Bernstein, Assistant Professor of Computer Science, Human-Computer Interaction Group, Computer Science Department, Stanford University
Emma Brunskill, Assistant Professor of Computer Science, Stanford University
Serena Yeung, Assistant Professor of Biomedical Data Science and, by courtesy, of Electrical Engineering, Stanford University
Dorsa Sadigh, Assistant Professor of Computer Science and of Electrical Engineering, Stanford University

3:55–
4:55

Panel Discussion

MODERATOR: **Eric Horvitz**, Technical Fellow and Director, Microsoft Research; Distinguished Fellow, HAI
Russ Altman, Kenneth Fong Professor and Professor of Bioengineering, of Genetics, of Medicine, of Biomedical Data Science and, by courtesy, of Computer Science, Stanford University
Justine Cassell, Associate Dean for Technology Strategy and Impact, Director Emerita of the Human-Computer Interaction Institute, and Co-Director of InMind Project on the Future of Personal Assistants, Carnegie Mellon University
Fernanda Viégas, Co-Leader, Google's PAIR (People+AI Research) Initiative
Bob Zhang, Co-Founder and CTO, Didi Chuxing

4:55–
5:10

Keynote: Governor Gavin Newsom

Gavin Newsom, Governor of California

5:10–
6:30

Reception and Student Posters

Acknowledgments & News

HAI Launch Task Force

Steve Denning, Co-Chair
Reid Hoffman, Co-Chair
Jerry Yang, Co-Chair
Jim Breyer
Jeff Dean
John Hennessy
Eric Horvitz
Bob King
James Manyika
Marissa Mayer
Sam Palmisano
Eric Schmidt
Mike Schroepfer
Kevin Scott
Neil Shen
Ram Shriram

Projects and Centers

AI 100
AI Index
Center for the Study of Language and Information
Center for AI Safety

Diversity Partners

AI4ALL
AI4ALL @ Stanford

Fellowship Partners

McCoy Family Center for Ethics in Society
Center for International Security and Cooperation
John S. Knight Journalism Fellowships
Stanford Institute for Economic Policy Research

Upcoming Events

The Coming AI Upheaval
April 22, 2019
Fei-Fei Li
Yuval Harari
MODERATOR:
Nicholas Thompson
Co-sponsored by the McCoy Family Center for Ethics in Society and the Stanford Humanities Center
Conference on AI, People & Society
October 28, 2019

Reid Hoffman
DJ Patil
Eric Schmidt

Executive Education

Congressional AI and Cyber Boot Camp
With Center for International Security and Cooperation and the Hoover Institution
Summer 2019
Harnessing AI for Breakthrough Innovation and Strategic Impact
With Stanford Graduate School of Business
August 2019

Distinguished Fellows

Yoshua Bengio
University of Montreal
Rodney Brooks
MIT
Erik Brynjolfsson
MIT

Jeff Dean
Google

Daniel Dennett
Tufts University

Susan Dumais
Microsoft Research

Edward Feigenbaum
Stanford University

Barbara Grosz
Harvard University

Demis Hassabis
DeepMind

Geoff Hinton
University of Toronto

Eric Horvitz
Microsoft Research

James Manyika
McKinsey & Company

John Markoff
Center for Advanced Study in the Behavioral Sciences

Helen Nissenbaum
Cornell Tech

Judea Pearl
UCLA

Stuart Russell
UC Berkeley

Mustafa Suleyman
DeepMind

Hal Varian
Google

Terry Winograd
Stanford University

Speakers



Alison Gopnik

is a professor of psychology and affiliate professor of philosophy at the University of California at Berkeley. She received her BA from McGill University and her PhD from Oxford University. She is a recognized leader in the study of children's learning and development and was one of the founders of "theory of mind," an originator of the "theory theory" of children's development, and more recently introduced the idea that probabilistic models and Bayesian inference could be applied to children's learning. Gopnik is the author or co-author of over 100 articles and several books including *Words, Thoughts, and Theories*, 1997; *The Scientist in the Crib*, 1999; *The Philosophical Baby*, 2009; and *The Gardener and the Carpenter*, 2016. She has written widely about cognitive science and psychology for *Science*, *The New York Times*, *Scientific American*, and other major publications. Since 2013 she has written the "Mind and Matter" column for *The Wall Street Journal*.



Amy Jin

is a 19-year-old student from San Jose, CA. Ever since she was introduced to computer science and artificial intelligence, she has been fascinated by their ability to serve as interfaces and tools to explore different disciplines. In 2015, she attended the inaugural AI4ALL summer program, where she was exposed to various branches of AI. She later began working on her research project, "Tool Detection and Operative Skill Assessment in Surgical Videos Using Region-Based Convolutional Neural Networks," under the mentorship of Serena Yeung and Jeffrey Jopling. Now a first-year student studying computer science at Harvard University, Jin hopes to combine her interests in CS, AI, and ethics for social good and to continue tackling complex problems as a socially conscious innovator. She would like to thank her mentors and professors, the Stanford Artificial Intelligence Lab, and the Clinical Excellence Research Center for the incredible research opportunity.



Bill Gates

is co-chair of the Bill & Melinda Gates Foundation. In 1975, Bill Gates founded Microsoft with Paul Allen and led the company to become the worldwide leader in business and personal software and services. In 2008, Bill transitioned to focus full-time on his foundation's work to expand opportunity to the world's most disadvantaged people. Along with co-chair Melinda Gates, he leads the foundation's development of strategies and sets the overall direction of the organization. In 2010, Bill, Melinda, and Warren Buffett founded the Giving Pledge, an effort to encourage the wealthiest families and individuals to publicly commit more than half of their wealth to philanthropic causes and charitable organizations during their lifetime or in their will. In 2015, Bill created the Breakthrough Energy Coalition, a group of individuals and entities committed to clean energy innovation, followed by Breakthrough Energy Ventures in 2016, an investor-led fund focused on providing patient capital to support cutting-edge clean energy companies.



Bob Zhang

is the co-founder and chief technology officer of Didi Chuxing, the world's leading mobile transportation platform. Zhang leads the creation and continuous development of DiDi's overall product, technology, and big-data analytical framework and oversees the company's evolution from a taxi-hailing app into an app-based transportation platform providing Taxi, Express, Premier, Luxe, Bus, Designated Driving, Enterprise Solutions, Bike-Sharing, E-bike Sharing, Car Sharing, and food delivery services, serving more than 550 million riders and 31 million drivers. Previously, he was a senior technology leader at Baidu, responsible for the creation of over 10 software applications with more than 100 million users each. Zhang obtained his BS in software engineering from Wuhan University in 2005 and his MS from the Chinese National Academy of Sciences in 2008, with research focusing on human-machine interaction and artificial intelligence (AI).

Speakers



Christopher Manning

is the inaugural Thomas M. Siebel Professor in Machine Learning in the departments of computer science and linguistics at Stanford University. His research goal is computers that can intelligently process, understand, and generate human language material. Manning is a leader in applying deep learning to natural language processing. He also focuses on computational linguistic approaches to parsing, robust textual inference, and multilingual language processing, including being a principal developer of Stanford Dependencies and Universal Dependencies. Manning has co-authored leading textbooks on statistical approaches to Natural Language Processing (NLP) (Manning and Schütze 1999) and information retrieval (Manning, Raghavan, and Schütze, 2008). His research has won ACL, Coling, EMNLP, and CHI best paper awards. He has a BA from The Australian National University and a PhD from Stanford and held faculty positions at Carnegie Mellon University and the University of Sydney before returning to Stanford. He is the founder of the Stanford NLP group.



Demis Hassabis

is the co-founder and CEO of DeepMind, a neuroscience-inspired AI company that develops general-purpose learning algorithms and uses them to help tackle some of the world's most pressing challenges. Since its founding in London in 2010, DeepMind has published over 200 peer-reviewed papers, five of them in the scientific journal *Nature*, an unprecedented achievement for a computer science lab. DeepMind's groundbreaking work includes the development of deep reinforcement learning, combining the domains of deep learning and reinforcement learning. In 2014, DeepMind was acquired by Google and is now part of Alphabet. After leading successful technology startups for a decade and, prior to founding DeepMind, Hassabis completed a PhD in cognitive neuroscience at University College London, followed by postdocs at MIT and Harvard. He is a five-time World Games Champion, recipient of the Royal Society's Mullard Award, and a fellow of the Royal Society of Arts and the Royal Academy of Engineering.



Daniel E. Ho

is the William Benjamin Scott and Luna M. Scott Professor of Law, a professor (by courtesy) of political science, and a senior fellow at the Stanford Institute for Economic Policy Research (SIEPR). His scholarship centers on quantitative empirical legal studies, with a substantive focus on administrative law, regulatory policy, anti-discrimination law, and courts. He received his JD from Yale Law School and a PhD in political science from Harvard University. He clerked for Judge Stephen F. Williams on the U.S. Court of Appeals, District of Columbia Circuit. Ho has served as president of the Society of Empirical Legal Studies (2011-12), co-editor of *The Journal of Law, Economics, & Organization* (2013-16), and is a faculty fellow at the Center for Advanced Study in the Behavioral Sciences. He directs the Regulation, Evaluation, and Governance Lab (RegLab) at Stanford.



Dorsa Sadigh

is an assistant professor in computer science and electrical engineering at Stanford University. Her research interests lie in the intersection of robotics, learning and control theory, and algorithmic human-robot interaction. Specifically, she works on developing efficient algorithms for autonomous systems that safely and reliably interact with people. She received her bachelor's degree in electrical engineering and computer sciences (EECS) at the University of California, Berkeley, in 2012 and her doctoral degree in EECS at UC Berkeley, in 2017. She is a recipient of the Amazon Faculty Research Award and National Science Foundation (NSF) and National Defense Science and Engineering Graduate (NDSEG) research fellowships as well as the Leon O. Chua Award at UC Berkeley.



David Freeman Engstrom

is a nationally recognized expert in civil procedure, administrative law, and constitutional law. Current work includes a project for the Administrative Conference of the United States on AI use by federal agencies and a project on the effect of advances in "legal tech" on the civil justice system. He is also serving as an associate dean at Stanford Law School and is leading an initiative charting the school's future work around digital technology. Beyond teaching and research, Engstrom has served as counsel or consultant to a wide range of public and private entities and is a frequent amicus before the U.S. Supreme Court. He is an elected member of the American Law Institute and a faculty affiliate at CodeX: The Stanford Center for Legal Informatics. He holds a JD from Stanford Law School, an MSc from Oxford University, and a PhD from Yale University.



Emma Brunskill

is an assistant professor in the computer science department at Stanford University where she leads the AI for Human Impact (@ai4hi) group. Her work focuses on reinforcement learning in high-stakes scenarios—how can an agent learn from experience to make good decisions when experience is costly or risky, such as in educational software, health care decision making, robotics or people-facing applications. She was previously on the faculty at Carnegie Mellon University. She is the recipient of multiple early faculty career awards, including from the National Science Foundation, the Office of Naval Research, and Microsoft Research. Her group has received several best research paper nominations and awards in top machine learning for education conferences and AI conferences.

Speakers



Eric Horvitz

is technical fellow and director of Microsoft Research Labs. He pursues research on principles and applications of AI, including efforts in learning, inference, planning, and human-AI collaboration. His efforts and collaborations have led to fielded AI technologies in health care, transportation, e-commerce, productivity applications, and computing systems. He received the AAAI-ACM Allen Newell Award and the AAAI Feigenbaum Prize for contributions in AI. He has been elected a fellow of the National Academy of Engineering (NAE), the Association for the Advancement of AI (AAAI), the Association of Computing Machinery (ACM), and the American Academy of Arts and Sciences. He chairs Microsoft's Aether board, Microsoft's committee focused on the responsible fielding of AI. He earned his PhD and MD degrees at Stanford University.



Erik Brynjolfsson

is director of the MIT Initiative on the Digital Economy, Schussel Family Professor at the MIT Sloan School, and a research associate at NBER. His research examines the effects of information technologies on business strategy, productivity and performance, digital commerce, and intangible assets. Brynjolfsson was among the first researchers to measure the productivity contributions of IT and the complementary role of organizational capital and other intangibles. His research also provided the first quantification of the value of online product variety, often known as the "Long Tail." Author of several books including, with co-author Andrew McAfee, *NYTimes* best-seller *The Second Machine Age: Work, Progress and Prosperity in a Time of Brilliant Technologies* (2014) and *Machine, Platform, Crowd: Harnessing Our Digital Future* (2017). He holds bachelor's and master's degrees from Harvard University in applied mathematics and decision sciences and a PhD from MIT in managerial economics and has taught at Harvard and Stanford.



Fei-Fei Li

is a professor in Stanford's computer science department and co-director of the Stanford Institute for Human-Centered Artificial Intelligence. She served as the director of Stanford's AI Lab from 2013 to 2018. During her sabbatical from Stanford from January 2017 to September 2018, she was a VP at Google and served as chief scientist of AI/ML at Google Cloud. Li obtained her BA in physics from Princeton in 1999 and her PhD in electrical engineering from Caltech in 2005. She joined Stanford in 2009 as an assistant professor. Her main research areas are in machine learning, deep learning, computer vision, and cognitive and computational neuroscience. Li is the inventor of ImageNet and the ImageNet Challenge, a critical large-scale dataset and benchmarking effort that has contributed to the latest developments in deep learning and AI. She is a fellow of the Association for Computing Machinery (ACM), recipient of the 2017 Athena Award for Academic Leadership and the 2016 IEEE PAMI Mark Everingham Award, among many others.



Fernanda Viégas

is a senior researcher at Google, where she co-leads the PAIR (People+AI Research) initiative, part of Google Brain. Her work in machine learning, with longtime collaborator Martin Wattenberg, focuses on improving human/AI interaction with a broader agenda of democratizing AI technology. Their team has open-sourced tools for ML interpretability and fairness, created educational materials for non-experts, and helped pioneer the use of ML on the web with the launch of TensorFlow.js. She is well known for her contributions to social and collaborative visualization, and the systems she and her team have created are used daily by millions of people. Her visualization-based artwork has been exhibited worldwide and is part of the permanent collection of the Museum of Modern Art in New York. She holds a PhD from the MIT Media Lab.



Gavin Newsom

is the governor of California, formerly lieutenant governor of California and mayor of San Francisco. Newsom is recognized for his bold willingness to lead — repeatedly developing, advocating, and implementing innovative and groundbreaking solutions to some of our most challenging issues. His leadership on a wide range of issues, including same-sex marriage, gun safety, marijuana, the death penalty, universal health care, access to preschool, technology, criminal justice reform, and the minimum wage, has often resulted in sweeping changes when his policies were ultimately accepted and replicated. Governor Newsom's top priorities for his administration are tackling the state's affordability crisis, creating inclusive economic growth and opportunity for every child, and standing up for California values that are under attack from Washington — from civil rights to immigration to environmental protection, education and increasing affordable access to quality schools at all levels, and justice.



James Manyika

is a senior partner at McKinsey & Company and chair and director of the McKinsey Global Institute (MGI), the firm's business and economics research arm. Based in Silicon Valley for over 20 years, he has worked with the chief executives and founders of many of the world's leading technology companies. At MGI, he has led research on technology and its impact. He has published a book on AI and robotics as well as numerous articles and reports and recently co-authored *No Ordinary Disruption* (2016). He also serves on various academic advisory boards, including the Oxford Internet Institute, MIT's Initiative on the Digital Economy, UC Berkeley's School of Information, Harvard's Hutchins Center, and the W.E.B. Du Bois Institute for African and African-American Research. He is a member of the standing committee for the Stanford-based 100 Year Study on Artificial Intelligence, a fellow at DeepMind, a non-resident senior fellow of the Brookings Institution in Economic Studies, and a fellow of the Royal Society of Arts.

Speakers



Jeff Dean

joined Google in 1999 and is currently a senior fellow and SVP for Google AI. He oversees Google's research division and health division. His teams are working on systems for speech recognition, computer vision, language understanding, and other machine-learning tasks. He has co-designed/implemented many generations of Google's crawling, indexing, and query-serving systems and co-designed/implemented major pieces of Google's early advertising systems. He is a co-designer and co-implementor of Google's distributed computing infrastructure, including the MapReduce, BigTable, and Spanner systems, protocol buffers, the open-source TensorFlow system for machine learning, and a variety of internal and external libraries and developer tools. He received a PhD in computer science from the University of Washington in 1996. He is a member of the National Academy of Engineering, a fellow of the Association for Computing Machinery (ACM), the American Association for the Advancement of Sciences (AAAS), and a winner of the ACM Prize in Computing.



Jennifer Pan

is an assistant professor of communication and an assistant professor, by courtesy, of political science and sociology at Stanford University. Her research focuses on the intersection of information, technology, and politics: how autocrats enforce social control through online censorship, propaganda, and responsiveness; how new media technologies allow regime insiders to manipulate information; how public preferences are arranged and formed as a result. Pan combines experimental and computational methods with large-scale datasets on political activity in China and other non-democratic regimes to examine these questions. Her work has appeared in peer-reviewed publications including *American Political Science Review*, *American Journal of Political Science*, *Comparative Political Studies*, *Journal of Politics*, and *Science* and has been covered in numerous media outlets, including *The New York Times*, *The Wall Street Journal*, *the Washington Post*, and *The Economist*. Pan received her PhD from Harvard in 2015. Until 2009, she was a consultant at McKinsey & Company based in New York and Beijing.



John Etchemendy

is the Patrick Suppes Family Professor in the Stanford School of Humanities and Sciences and is now the faculty co-director of the Stanford Institute for Human-Centered Artificial Intelligence (HAI). He was Stanford's twelfth—and longest serving – provost. He is currently a professor of philosophy and symbolic systems. In his 35 years as a Stanford faculty member, Etchemendy has also served as director of the Center for the Study of Language and Information, chair of the Philosophy Department, and associate dean for the School of Humanities and Sciences. He earned his undergraduate and master's degrees in philosophy from University of Nevada, Reno and his PhD from Stanford.



Juliana Bidadanure

is an assistant professor in the Department of Philosophy at Stanford University. Her research is located at the intersection of political philosophy and public policy. She is currently writing a book on issues of justice between generations—asking how we should distribute resources between age groups and what it means for young and old to stand as equals. Bidadanure has written on a number of public policies, including youth quotas in parliaments, basic income and basic capital. She is the founder and faculty director of the Stanford Basic Income Lab—a center on campus that promotes research on the philosophy, politics, and economics of universal basic income and informs policymakers and practitioners about latest best practices.



Justine Cassell

is associate dean of technology strategy and impact in the School of Computer Science at Carnegie Mellon University and director emerita of the Human Computer Interaction Institute. She co-directs the Yahoo-CMU InMind partnership on the future of personal assistants and was a founding co-director of the Simon Initiative on Technology-Enhanced Learning. Previously, Cassell was at Northwestern University where she founded the Technology and Social Behavior Doctoral Program and Research Center. Before that she was a tenured professor at the MIT Media Lab. In 2011 she was named to the World Economic Forum Global Agenda Council on AI and Robotics, which she chaired for two years before moving to chair the World Economic Forum's Council on the Future of Computing. Cassell has spoken at the World Economic Forum in Davos for the past eight years on topics concerning the impact of AI and robotics on society.

Speakers



Kate Crawford

is a researcher, academic, and author who has spent more than a decade studying large-scale data systems, machine learning, and artificial intelligence. She is the co-founder and co-director of the AI Now Institute at NYU, which conducts research across computer science, social science, and law to better understand and address the social implications of artificial intelligence. In 2016, she co-chaired the Obama White House symposium on the social and economic implications of AI. Crawford has published in academic journals, including *Nature*, *New Media & Society*, and *Information, Communication & Society*, and she has written for *The New York Times*, *Harper's Magazine*, and *the Washington Post*. She has advised policymakers at the European Commission, the United Nations, the Federal Trade Commission, and the City of New York. In 2018, she was selected for a Richard von Weizsäcker Fellowship at the Robert Bosch Academy in Berlin.



Marc Tessier-Lavigne

President of Stanford University since 2016, was born in Canada. He received degrees in physics from McGill University and in philosophy and physiology from Oxford University. He earned a PhD in physiology from University College London (UCL). Tessier-Lavigne's research has focused on the cause and treatment of degenerative brain diseases. He and his colleagues revealed how neural circuits in the brain form during embryonic development by identifying molecules that direct the formation of connections among nerve cells. His contributions have been recognized by numerous prizes and honors, including his election as a member of the National Academy of Sciences, the National Academy of Medicine, and the American Philosophical Society and as a fellow of the Royal Society (UK), the Royal Society of Canada, the Academy of Medical Sciences (UK), the American Association for the Advancement of Science, and the American Academy of Arts and Sciences.

He has held faculty positions at the University of California, San Francisco (UCSF) and at Stanford where from 2001 to 2005 he was a professor of biological sciences and held the Susan B. Ford Professorship. In 2003 Tessier-Lavigne was recruited to biotechnology company Genentech, where he served as executive vice president for research and chief scientific officer, directing 1,400 scientists in disease research and drug discovery for cancer, immune disorders, infectious diseases, and neurodegenerative diseases, while maintaining an active research laboratory. In 2011, he became president of The Rockefeller University, a leading biomedical research university in New York City, before returning to the West Coast and rejoining Stanford.



Mark Duggan

is the Trione Director of the Stanford Institute for Economic Policy Research and Wayne and Jodi Cooperman Professor of Economics at Stanford. His research focuses on health care and the effects of government programs including Medicare, Medicaid, and Social Security. His research has been published in leading academic outlets including the *American Economic Review*, *the Journal of Political Economy*, and *the Quarterly Journal of Economics* and has been featured in *The Economist*, *The New York Times*, and *the Wall Street Journal*. Duggan was the 2010 recipient of the ASHEcon Medal (awarded once every two years to the leading health economist in the U.S. under age 40). He has testified about his research to committees in the U.S. Senate and U.S. House of Representatives and he served from 2009 to 2010 as the senior economist for health care policy at the White House Council of Economic Advisers. He teaches "Econ 1" at Stanford and advises dozens of undergraduate and graduate students.



Michael Bernstein

is an assistant professor of computer science at Stanford University, where he is a member of the Human-Computer Interaction group. His research focuses on the design of crowdsourcing and social computing systems. His work has received numerous best paper awards at premier computing venues, and his PhD students have gone on both to industry (e.g., Adobe Research, Facebook Data Science) and faculty positions (e.g., Carnegie Mellon, UC Berkeley). He has been recognized as a Robert N. Noyce Family Faculty Scholar and has received an NSF CAREER award, an Outstanding Academic Title citation from the American Library Association, and an Alfred P. Sloan Fellowship. He holds a bachelor's degree in symbolic systems from Stanford University, as well as a master's degree and a PhD in computer science from MIT.



Michael C. Frank

is the David and Lucile Packard Professor of Human Biology at Stanford University. He received his PhD from MIT in brain and cognitive sciences in 2010. He studies language use and language learning and how these interact with social cognition, focusing especially on early childhood. He is the organizer of the ManyBabies Consortium, a collaborative replication network for infancy research, and has led open-data projects including Wordbank and MetaLab. He has been recognized as a "rising star" by the Association for Psychological Science. His dissertation received the Glushko Prize from the Cognitive Science Society, and he is recipient of the FABBS Early Career Impact award and a Jacobs Advanced Research Fellowship. He has served as associate editor for the journal *Cognition*, member and chair of the governing board of the Cognitive Science Society, and was a founding executive committee member of the Society for the Improvement of Psychological Science.

Speakers



Percy Liang

is an assistant professor of computer science at Stanford University (BS from MIT, 2004; PhD from UC Berkeley, 2011). His research spans machine learning and natural language processing, with the goal of developing trustworthy agents that can communicate effectively with people and improve over time through interaction. Specific topics include question answering, dialogue, program induction, interactive learning, and reliable machine learning. His awards include the IJCAI Computers and Thought Award (2016), an NSF CAREER Award (2016), a Sloan Research Fellowship (2015), and a Microsoft Research Faculty Fellowship (2014).



Reid Hoffman

an accomplished entrepreneur, executive, and investor, has played an integral role in building many of today's leading consumer technology businesses. In 2003 he co-founded LinkedIn, the world's largest professional networking service. In 2009 he joined Greylock Partners. He currently serves on the boards of Airbnb, Apollo Fusion, Aurora, Coda, Convoy, Entrepreneur First, Gixo, Microsoft, Nauto, Xapo, and a few early-stage companies still in stealth. In addition, he serves on a number of not-for-profit boards, including Biohub, CZI, Do Something, Endeavor, Kiva, Stanford's Human-Centered AI Initiative, and the MacArthur Foundation's 100&Change. He is the co-author of two *New York Times* best-selling books: *The Start-Up of You* and *The Alliance*. His new book, *Blitzscaling*, is based on his Stanford course of the same name. He is an Aspen Institute Crown Fellow, a Marshall Scholar at Oxford, and a graduate of Stanford University.



Russ Altman

is a professor of bioengineering, genetics, medicine, and biomedical data science (and of computer science, by courtesy) and past chairman of the bioengineering department at Stanford University. His primary research interests are in the application of computing and informatics technologies to problems relevant to medicine. He is particularly interested in methods for understanding drug action at molecular, cellular, organism, and population levels. His lab studies how human genetic variation impacts drug response. Other work focuses on the analysis of biological molecules to understand the actions, interactions, and adverse events of drugs. He helps lead an FDA-supported Center of Excellence in Regulatory Science & Innovation. Altman holds an AB from Harvard College, an MD from Stanford School of Medicine, and a PhD in medical information sciences from Stanford. He is a past-president, founding board member, and a fellow of the International Society for Computational Biology (ISCB). He is an organizer of the annual Pacific Symposium on Biocomputing.



Serena Yeung

is an assistant professor of biomedical data science and, by courtesy, of electrical engineering at Stanford University starting in the fall of 2019. She is currently a Technology for Equitable and Accessible Medicine (TEAM) Research Fellow at Harvard University. She received her PhD from Stanford in 2018, where she was a member of the Artificial Intelligence Lab and the lead graduate student in the Partnership in AI-Assisted Care (PAC), a collaboration with Stanford's Clinical Excellence Research Center. Her research has been broadly in the areas of computer vision, machine learning, and deep learning, with application to AI-assistance in health care. She was also co-instructor of Stanford's CS231N "Convolutional Neural Networks for Visual Recognition" course in 2017 and 2018. Yeung spent time at Facebook AI Research in 2016 and Google Cloud AI in 2017.



Stephanie Tena-Meza

was born and raised in Salinas, Calif., and is currently a junior at Salinas High School. Her passion is using her knowledge about STEM-related fields including artificial intelligence (AI) and computer science (CS) to create technology opportunities that can uplift her community. One of her proudest accomplishments is starting a CS and AI after-school club for seventh- and eighth-grade students at her former middle school. She serves as the instructor and teaches students about both fields. She has received support from AI4ALL, a nonprofit organization, through impact grants to fund her club. She also enjoys participating in her underwater robotics team and advocating for the youth of Salinas. In the future, she plans to attend college in California and pursue computer science as a major. As a future computer scientist, she hopes to use her CS skills to combat social issues. Lastly, she plans on returning to her hometown to give back to her community.



Surya Ganguli

triple-majored in physics, mathematics, and electrical engineering and computer science at MIT, completed a PhD in string theory at Berkeley and a postdoc in theoretical neuroscience at UCSF. He is now an assistant professor of applied physics at Stanford where he leads the Neural Dynamics and Computation Lab and is also a consulting professor on the Google Brain Research Team. His research spans the fields of neuroscience, machine learning, and physics, focusing on understanding and improving how both biological and artificial neural networks learn striking emergent computations. He has been awarded a Swartz-Fellowship in computational neuroscience, a Burroughs-Wellcome Career Award at the scientific interface, a Terman Award, a NIPS Outstanding Paper Award, an Alfred P. Sloan foundation fellowship, a James S. McDonnell Foundation scholar award in human cognition, a McKnight Scholar award in neuroscience, and a Simons Investigator Award in the mathematical modeling of living systems.

Speakers



Susan Athey

is The Economics of Technology Professor at Stanford Graduate School of Business and a SIEPR senior fellow. She received her bachelor's degree from Duke University (1991) and her PhD from Stanford (1995). She previously taught in the economics departments at MIT, Stanford, and Harvard. Her current research focuses on the economics of the internet, marketplace design, auction theory, the statistical analysis of auction data, and the intersection of econometrics and machine learning. She has focused on several applications, including timber auctions, internet search, online advertising, the news media, and virtual currency. She advises governments and businesses on the design of auction-based marketplaces. She has served as a long-term consultant for Microsoft Corporation since 2007, including a period as chief economist. At the age of 36, Athey received the John Bates Clark Medal, which honors an "American economist under the age of forty who is adjudged to have made the most significant contribution to economic thought and knowledge."



Tristan Harris

has been called "the closest thing Silicon Valley has to a conscience" by *The Atlantic* magazine. Prior to founding the new Center for Humane Technology, he was Google's design ethicist, developing a framework for how technology should "ethically" steer the thoughts and actions of billions of people from screens. He has spent a decade understanding the invisible influences that hijack human thinking and action. Drawing on literature from addiction, performative magic, social engineering, persuasive design, and behavioral economics, he is currently developing a framework for ethical persuasion, especially as it relates to the moral responsibility of technology companies. His work has been featured widely, including on TV and in *The Atlantic*, *WIRED*, *The New York Times*, *Der Spiegel*, *The Economist*, and elsewhere. He has briefed heads of state, technology company CEOs, and members of Congress about the attention economy. He is a visiting fellow at the Harvard Kennedy School of Government and an advisor to Open Markets Institute.



The colored threads that make up the distinctive HAI typeface symbolize human diversity in all its forms, and their weaving together reminds us of our common humanity.

Stanford Visitor Wireless Access

How to get online:

1. Select the SSID Stanford Visitor from your device's list of available wireless networks.
2. Open a browser and load any URL; you will be redirected to an access page (on a mobile phone, the access page should open automatically).
3. Click Accept to acknowledge the terms of use. Your browser will be sent to a confirmation page, and from there you can get online.

Stanford
Human-Centered
Artificial Intelligence

