




Stanford University
Human-Centered
Artificial Intelligence



Congressional Boot Camp on Artificial Intelligence

August 10-12, 2026


Stanford University



This page is intentionally left blank.

Table of Contents

A Note From the Conveners	4
Agenda	8
Keynote Speakers’ Biographies	12
Session Descriptions	20
Staffer Biographies	53
Thank you to HAI Staff	61



A Note From the Conveners

On behalf of the Stanford Institute for Human-Centered AI (HAI), we would like to welcome you to this year’s Congressional Boot Camp on Artificial Intelligence. We are thrilled to have you here at Stanford University—in the heart of Silicon Valley—to present you with new ideas and perspectives to enhance your work crafting policy around this exciting technology.

This has been another pivotal year for the field of AI. From international summits to national hearings, the call to *action* is clear. Around the world, there is a great sense of urgency to accelerate frontier research and real-world applications. Governments are mobilizing to shape innovation ecosystems, invest in public infrastructure, and define their role in the AI era. While the United States remains the undisputed leader in AI, bipartisan discussions are ongoing in Washington, D.C., about how to sustain and strengthen that leadership.



We started the Congressional Boot Camp on AI in a shared spirit of urgency. In a rapidly advancing field, policymakers need to act with alacrity—yet thoughtfully. A core part of our mission is to supply them with the tools to do so by providing access to timely, evidenced-based insights that help inform forward-thinking, effective policy.

That’s why we deeply believe that lasting leadership in AI requires sustained investment in the academic research ecosystem that underpins it. The breakthroughs that power today’s frontier systems were seeded by decades of publicly funded, curiosity-driven research. Preserving that ecosystem is essential to ensuring AI advances in ways that align with the public interest.

Over the next three days, you will hear from world-class Stanford scholars, as well as industry leaders and civil society pioneers. Their collective expertise will illuminate not just the state of the technology but also the enduring foundations needed to guide it responsibly.

We hope you will also use this opportunity to share your insights and experience in AI policymaking. The exchange of ideas is fundamental to robust scholarship, and we encourage you to ask questions and challenge your perspectives on AI. Through dialogue, we can begin to address the most pressing questions about AI innovation and governance.

Thank you all for traveling here to join us. We would also like to extend a special thank-you to our speakers for taking the time to share their expertise at the boot camp. Finally, we would like to express deep gratitude to our exceptional Stanford HAI staff—Alex Armira, Annie Benisch, Elena Cryst, Chris Ellis, Angie Gieger-Chan, Dana Heinneman, David Levi, Shana Lynch, Jeanina Matias, Caroline Meinhardt, Tracy Navichoque, Vanessa Parli, Harini Sreepathi, Angel Wazin, Madeleine Wright, Caroline Yee, Daniel Zhang, and Maddie Ziff—for all their hard work in organizing this year’s boot camp.

Welcome to Stanford University, aka “the Farm.” We look forward to our time together.

Sincerely,



James Landay
Anand Rajaraman and Venky Harinarayan Professor in the School of Engineering, Professor of Computer Science, Stanford University
Denning Director, Stanford HAI



Fei-Fei Li
Sequoia Professor of Computer Science, Stanford University
Founding Director, Stanford HAI



Russell Wald
Executive Director, Stanford HAI



James Landay, Fei-Fei Li, and John Hennessy at the 2026 HAI Spring Conference
(Photo credit: David Kim - SF Photo Agency)




Daniel Ho presenting at the 2024 Stanford California State Boot Camp on AI, co-hosted by Stanford HAI, the Stanford Institute for Economic Policy Research (SIEPR), and the Stanford Regulation, Evaluation, and Governance Lab (RegLab)
(Photo credit: Ryan Zhang)



James Landay speaking at HAI at Five: Celebrating 5 Years of Impact
(Photo credit: Christine Baker)



Erik Brynjolfsson at the launch of the Digitalist Papers: Artificial Intelligence and Democracy in America
(Photo credit: Patrick Beaudouin)



Boot Camp Agenda

MONDAY, AUGUST 10, 2026

Location: Gates Computer Science Building, 353 Jane Stanford Way, Stanford, CA 94305

11:10 am – 11:15 am	Welcome Speaker: Elena Cryst
11:15 am – 12:15 pm	Session 1: A Call for Human-Centered AI Speaker: James Landay
12:15 pm – 1:15 pm	Session 2: The Fuel of AI: Data (and Its Perils) Speakers: Jennifer King, James Zou, Alex Spangher (moderator)
1:15 pm – 1:45 pm	Break
1:45 pm – 2:45 pm	Session 3: The Role of Compute and Chips Speakers: Azalia Mirhoseini, H.-S. Philip Wong, Drew Spence (moderator)
2:45 pm – 3:45 pm	Session 4: Mitigating Risk: Implementing Safe and Robust AI Speakers: Mykel Kochenderfer, Sanmi Koyejo, Tracy Navichoque (moderator)
3:45 pm – 4:00 pm	Break
4:00 pm – 5:00 pm	Session 5: The Governance of AI Speakers: Rishi Bommasani, Deborah Raji, Rob Reich, Daniel Zhang (moderator)
5:00 pm – 7:30 pm	Dinner Keynote: AI, Automation and the Future of Work Speaker: Erik Brynjolfsson

TUESDAY, AUGUST 11, 2026

Location: Gates Computer Science Building, 353 Jane Stanford Way, Stanford, CA 94305

8:45 am – 9:00 am	Breakfast/Debrief
9:00 am – 10:00 am	Session 1: Transforming Healthcare Through Innovation Speakers: Curtis Langlotz, Michelle Mello, Nigam Shah, Caroline Meinhardt (moderator)
10:00 am – 10:05 am	Break
Location: Stanford Robotics Center, 350 Jane Stanford Way Stanford, CA 94305	
10:05 am – 11:00 am	Stanford Robotics Center Visit Speaker: Allison Okamura
11:00 am – 11:15 am	Break
Location: Gates Computer Science Building, 353 Jane Stanford Way, Stanford, CA 94305	
11:15 am – 12:15 pm	Session 2: AI in Motion: Role of Robotics in the Next Chapter of AI Development Speakers: Monroe Kennedy, Allison Okamura, Drew Spence (moderator)
12:15 pm – 1:30 pm	Lunch
1:30 pm – 2:30 pm	Session 3: Revolutionizing the Classroom: How AI Is Advancing Education Speakers: Emma Brunskill, Victor Lee, Isabelle Hau (moderator)

2:30 pm – 3:30 pm	Session 4: Modernizing a Mammoth: Use-Cases of Public Sector AI Speaker: Daniel Ho
3:30 pm – 4:00 pm	Transition to Li Ka Shing Center
Location: Li Ka Shing Center, 291 Campus Drive, Stanford, CA 94305	
4:00 pm – 5:00 pm	Session 5: The Role of Business: Policy Implications of Industry Leadership in AI Speakers: Sarah Guo, Matt Perault, Andy Song, Vanessa Parli (moderator)
5:00 pm – 6:00 pm	Reception
6:00 pm – 7:30 pm	Keynote Dinner: The New American Industrial Era: Charting a New Age of U.S.-led Innovation Welcome: James Landay Speakers: Jensen Huang, Fei-Fei Li, Condoleezza Rice, Russell Wald (moderator)

WEDNESDAY, AUGUST 12, 2026

Location: SLAC National Accelerator Laboratory, 2575 Sand Hill Rd, Menlo Park, CA 94025

8:00 am – 8:30 am	Shuttle to SLAC
8:30 am – 9:00 am	Breakfast/Debrief
9:00 am – 10:00 am	SLAC National Accelerator Laboratory Visit Speaker: Alberto Salleo
10:00 am – 10:30 am	Shuttle to Gates Building
Location: Gates Computer Science Building, 353 Jane Stanford Way, Stanford, CA 94305	
10:30 am – 11:30 am	Session 1: Unlocking Scientific Discovery: AI’s Potential to Accelerate Breakthroughs Speakers: Russ Altman, Ramesh Johari, Elena Cryst (moderator)
11:30 am – 12:30 pm	Session 2: Powering AI: Meeting the Energy Demands Needed for Frontier AI Speakers: William Chueh, Liang Min
12:30 pm – 1:00 pm	Lunch
1:00 pm – 2:30 pm	Session 3: AI and National Security: Strategic Competition Lessons for the US Speakers: Yejin Choi, Colin Kahl, Amy Zegart, Brad Boyd (moderator)
2:30 pm – 3:30 pm	Session 4: Updating the Law in the Age of AI Speakers: Nathaniel Persily, Riana Pfefferkorn, Irene Liu (moderator)
3:30 pm – 4:00 pm	Break
4:00 pm – 5:00 pm	Session 5: Agents on the Rise: Exploring Agentic AI Speakers: Alex “Sandy” Pentland, Diyi Yang, Andy Zhang (moderator)
5:00 pm – 5:15 pm	Break
5:15 pm – 7:30 pm	Fireside Chat: Cutting Edge: AI Developments Steering the Future Speakers: Mariano-Florentino “Tino” Cuéllar, John Hennessy, Elena Cryst (moderator)



James Zou presenting at the 2025 HAI Hoffman-Yee Symposium
(Photo credit: Christine Baker)



Sanmi Koyejo and Rob Reich speaking at the Stanford 2025 AI + Education Conference co-hosted by Stanford HAI and the Stanford Accelerator for Learning
(Photo credit: Christine Baker)



Russell Wald, Fei-Fei Li, Sarah Friar, and Condoleezza Rice at Stanford HAI's 2025 Congressional Boot Camp on AI
(Photo credit: David Gonzales)



Chris Manning speaking at the 2025 Workshop on Public AI Assistants to World Wide Knowledge (WWK) co-hosted by Stanford HAI and the Stanford Open Virtual Assistant Lab (OVAL)
(Photo credit: Christine Baker)

Plenary Session and Dinner Speakers



Erik Brynjolfsson

Erik Brynjolfsson is the Jerry Yang and Akiko Yamazaki Professor and Senior Fellow at the Stanford Institute for Human-Centered AI and director of the Stanford Digital Economy Lab. He also serves as the Ralph Landau Senior Fellow at the Stanford Institute for Economic Policy Research (SIEPR), professor (by courtesy) at Stanford Graduate School of Business and Stanford Department of Economics, and a research associate at the National Bureau of Economic Research (NBER).

One of the most cited authors on the economics of information, Brynjolfsson was among the first researchers to measure productivity contributions of IT and the complementary role of organizational capital and other intangibles.

He has conducted pioneering research on digital commerce, the Long Tail, bundling and pricing models, intangible assets, and the effects of IT on business strategy, productivity, and performance.

Brynjolfsson speaks globally and is the author of nine books, including (with co-author Andrew McAfee) *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies* and *Machine, Platform, Crowd: Harnessing Our Digital Future*, and more than 100 academic articles. He holds bachelor's and master's degrees in applied mathematics and decision sciences from Harvard University and a PhD from MIT in managerial economics.



Mariano-Florentino “Tino” Cuéllar

Mariano-Florentino (Tino) Cuéllar is the Cameron Schrier Family Professor of Law and (by courtesy) Political Science at Stanford University, Sara Miller McCune Director of the Center for Advanced Study in the Behavioral Sciences, and Senior Fellow at the Stanford Institute for Human-Centered AI (HAI). Starting in September he will also lead Stanford's Knight-Hennessy Scholars — the largest fully endowed graduate fellowship program in the world. A former California Supreme Court justice and president of the Carnegie Endowment for International Peace, he has served in three presidential administrations at the White House and in federal agencies.

He is a member of the American Academy of Arts and Sciences, publishing

widely on problems in American public law and democracy as the U.S. became a global power, artificial intelligence and its impact on institutions and organizations, and managing global challenges such as international security and public health. He chairs the board of the William & Flora Hewlett Foundation.



John Hennessy

John L. Hennessy is co-founder and Shriram Family Director of Knight-Hennessy Scholars, the largest, university-wide, fully-endowed graduate fellowship in the world. He is Chairman of the Board of Alphabet and serves on the Board of Trustees for the Gordon and Betty Moore Foundation. Hennessy has been on the faculty of Stanford University since 1977 and previously served as the President of the university for 16 years after roles including chair of Computer Science, dean of the School of Engineering, and university provost. He co-founded MIPS Computer Systems and Atheros Communications. He and Dave Patterson were awarded the ACM A.M. Turing Prize for 2017 and the National Academy of Engineering Draper Prize in 2022.



Jensen Huang

Jensen Huang founded Nvidia in 1993 and has served as president, chief executive officer, and a member of the board of directors. Since its founding, Nvidia has pioneered accelerated computing. The company's invention of the GPU in 1999 sparked the growth of the PC gaming market, redefined computer graphics, and ignited the era of modern AI. Nvidia is now driving the platform shift of accelerated computing and generative AI, transforming the world's largest industries and profoundly impacting society.

Huang was elected to the National Academy of Engineering and, in 2026, was appointed to the President's Council of Advisors on Science and Technology. He is a recipient of the Semiconductor Industry Association's

highest honor, the Robert N. Noyce Award; the IEEE Founder's Medal; the Dr. Morris Chang Exemplary Leadership Award; and honorary doctorate degrees from Taiwan's National Chiao Tung University, National Taiwan University, Oregon State University, Huazhong University of Science and Technology, and Linköping University. He was named the world's best CEO by Fortune, the Economist, and Brand Finance, as well as one of TIME magazine's 100 most influential people.

Prior to founding Nvidia, Huang worked at LSI Logic and Advanced Micro Devices. He holds a BSEE degree from Oregon State University and an MSEE degree from Stanford University.



James Landay

James Landay is a professor of computer science and the Anand Rajaraman and Venky Harinarayan Professor in the School of Engineering at Stanford University. He specializes in human-computer interaction. Landay is the co-founder and director of the Stanford Institute for Human-Centered Artificial Intelligence. Prior to joining Stanford, he was a professor of information science at Cornell Tech in New York City for one year and a professor of computer science and engineering at the University of Washington for 10 years. From 2003 to 2006, he served as the director of Intel Labs Seattle, a leading research lab that explored various aspects of ubiquitous computing. Landay was also the chief scientist and co-founder of NetRaker, which was acquired by KeyNote Systems in 2004. Before that, he was an

associate professor of computer science at UC Berkeley. Landay received a BS in EECS from UC Berkeley and an MS and PhD in computer science from Carnegie Mellon University. His PhD dissertation was the first to demonstrate the use of sketching in user interface design tools. Landay is a member of the Association for Computing Machinery's (ACM) Special Interest Group on Computer-Human Interaction (SIGCHI) Academy, an ACM Fellow, and an ACM SIGCHI Lifetime Research Award winner. He served for six years on the National Science Foundation's Computer and Information Science and Engineering Advisory Committee.



Fei-Fei Li

Dr. Fei-Fei Li is the inaugural Sequoia Professor in the computer science department at Stanford University and a founding co-director of the Stanford Institute for Human-Centered AI. 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, Dr. Li was vice president at Google and chief scientist of AI/ML at Google Cloud. She has served as a board member or advisor in various public or private companies. She is the co-founder/CEO of World Labs, an AI company focusing on spatial intelligence and generative AI.

Dr. Li obtained a BA in physics with high honors from Princeton and a PhD in electrical engineering from the California Institute of Technology. She received the Distinguished Alumni Award from both Princeton and Caltech. Dr. Li also holds a doctorate (honorary) from Harvey Mudd College.

Dr. Li's current research interests include AI, machine learning, deep learning, computer vision, robotic learning, and ambient intelligence for healthcare delivery. She has also worked on cognitive and computational neuroscience. She has published more than 400 scientific articles in top-tier journals and conferences in science, engineering, and computer science. Dr. Li is the inventor of ImageNet and the ImageNet Challenge, a critical large-scale dataset and benchmarking effort that is widely regarded as one of the three driving forces of the birth of modern AI and deep learning revolution. In 2023, Dr. Li published a memoir, *The Worlds I See: Curiosity, Exploration and Discovery at the Dawn of AI*.

Dr. Li has been working with policymakers locally and nationally to ensure positive and human-centered progress in AI technologies. She has given testimony before the U.S. Senate and House of Representatives and served as a special advisor to the Secretary-General of the United Nations. She was also a member of the California Future of Work Commission for the governor of California in 2019-20 and a member of the National Artificial Intelligence Research Resource Task Force established by the White House Office of Science and Technology Policy and the National Science Foundation (NSF) in 2021-22. She has delivered numerous speeches, including to the president of the United States and the Security Council of the United Nations.

Dr. Li is an elected member of the National Academy of Engineering, the National Academy of Medicine, and American Academy of Arts and Sciences. She is also a Fellow of the Association for Computing Machinery and a member of the Council on Foreign Relations. Among her many recognitions, Dr. Li is a laureate of the VinFuture Prize (2024), a recipient of the Intel Lifetime Achievement Award (2023), the IEEE PAMI Thomas Huang Memorial Prize (2022), the IEEE PAMI Longuet-Higgins Prize (2019), the National Geographic Society Further Award (2019), the IAPR J.K. Aggarwal Prize (2016), the IEEE PAMI Mark Everingham Award (2016), the Alfred Sloan Faculty Award (2011), the NSF Career award (2009), and the Microsoft Research New Faculty Fellowship (2006). Li has been a keynote speaker at many conferences, including the World Economics Forum (Davos), the Grace Hopper Conference 2017, and TED2015 and TED2024. She was named one of the TimeAI100 influencers in 2023, Women in Tech by *Elle* magazine (2017), a Global Thinker of 2015 by Foreign Policy, and one of the "Great Immigrants: The Pride of America" in 2016 by the Carnegie Foundation.



Secretary Condoleezza Rice

Condoleezza Rice is the Tad and Dianne Taube Director of the Hoover Institution and a senior fellow in public policy. She is also the Denning Professor in Global Business at the Stanford Graduate School of Business and a founding partner of the international consulting firm Rice, Hadley, Gates & Manuel.

Rice served as the 66th U.S. secretary of state (2005–2009) and national security advisor (2001–2005), the first Black woman to hold either role. She has been a professor of political science at Stanford University since 1981 and served as provost from 1993 to 1999.

Earlier in her career, she held several roles on the National Security Council under President George H.W. Bush and was a fellow at the Council on Foreign Relations. She has authored numerous books on foreign policy, leadership, and democracy, including multiple bestsellers.

A champion for education, Rice co-founded the Center for a New Generation, which later merged with the Boys & Girls Clubs of America. She currently serves on the boards of C3.ai, Makena Capital, and the Aspen Institute, among others, and is a part owner of the Denver Broncos.

Born in Birmingham, Alabama, Rice holds a BA and PhD in political science from the University of Denver and an MA from Notre Dame. She is a fellow of the American Academy of Arts and Sciences and has received over 15 honorary doctorates.



Russell Wald

Russell Wald is a recognized leader in AI governance and policy. Wald has spent over a decade at the intersection of technology, society, and governance — focused on ensuring that transformative technologies are developed and deployed responsibly to improve the human condition. He works at the forefront of AI policy, balancing innovation, security, and human-centered values in a rapidly evolving landscape.

As executive director of the Stanford Institute for Human-Centered Artificial Intelligence, Wald leads HAI’s efforts to bridge academia, public policy, and civil society. He oversees HAI’s research, education, and outreach missions, guiding initiatives that foster collaboration, innovation, and thoughtful AI governance.

Wald is the co-author of various publications on AI, including *Building a National AI Research Resource* (2021), *Enhancing International Cooperation in AI Research: The Case for a Multilateral AI Research Institute* (2022), and “The Centrality of Data and Compute for AI Innovation: A Blueprint for a National Research Cloud” (2022, *Notre Dame Journal of Emerging Technologies*). Currently, he is part of a HAI seed grant research project titled “Addicted by Design: An Investigation of How AI-Fueled Digital Media Platforms Contribute to Addictive Consumption.” Additionally, he serves as a member of the AI Index Steering Committee hosted by HAI.

During his more than 10 years at Stanford, Wald has strengthened the bridge between academia and policymakers. He managed national security and government affairs at the Hoover Institution, where he worked with leading scholars on tech policy. He also co-founded Stanford’s Congressional Tech Policy Boot Camp, educating lawmakers on emerging technologies, and led key initiatives advancing AI’s societal impact.

Beyond Stanford, Wald is a Visiting Fellow at the National Security Institute at George Mason University and a former term member of the Council on Foreign Relations. He holds a degree from UCLA in political science and global studies.



Michelle Mello presenting at an HAI seminar on understanding liability risks from healthcare AI tools

(Photo credit: Jeanina Matias)



DJ Patil, Rachel Gillum, Sarah Guo, Michael Krastios, and Navrina Singh in conversation at Stanford HAI's 2024 Congressional Boot Camp on AI

(Photo credit: Daniel Beck)



Sanmi Koyejo speaking at the 2025 Stanford RAISE Health Symposium

(Photo credit: Steve Fisch, Stanford Medicine)



John Etchemendy speaking at the HAI event "Unlocking Public Sector AI Innovation: Next Steps for the National AI Research Resource"

(Photo credit: Jeanina Matias)

Session Descriptions:

MONDAY, AUGUST 11, 2026

11:10 am – 11:15 am

Welcome

This session will welcome staffers to Stanford’s campus and provide an overview of why this boot camp was created and what Stanford HAI hopes participants will gain.

Speaker:

Elena Cryst, Director of Policy and Society, Stanford HAI

Elena Cryst leads HAI’s efforts to bring Stanford’s cutting-edge AI research to policymakers worldwide. She previously served in director roles at both the Stanford Institute for Economic Policy Research and the Freeman Spogli Institute for International Studies, including at that institute’s Cyber Policy Center, Center for International Security and Cooperation, and Center on Democracy, Development and the Rule of Law. She is passionate about bringing informed, empirical research into public discourse.

Elena received her BA with honors in international relations and MA in Latin American studies, both from Stanford University, and her MBA from the UC Berkeley Haas School of Business.



11:15 am – 12:15 pm

Session 1: A Call for Human-Centered AI

AI for Good initiatives recognize the potential impact of AI systems on humans and societies. However, simply recognizing the impact is not enough. To be truly human-centered, AI development must be user-centered, community-centered, and societally centered. Successful human-centered AI requires the early engagement of multidisciplinary teams beyond technologists, including experts in design, the social sciences and humanities, and domains of interest such as medicine or law, as well as community members. This session will explore arguments for an authentic human-centered AI.

Speaker:

James Landay, Anand Rajaraman and Venky Harinarayan Professor in the School of Engineering, Professor of Computer Science, Stanford University; Denning Director, Stanford HAI

See biography on page 14.



12:30 pm – 1:30 pm

Session 2: The Fuel of AI: Data (and Its Perils)

Contemporary AI technologies run on data, but AI developers face significant obstacles to acquiring and cleaning data. In addition, developers must do their best to ensure data’s inherent biases (and their non-obvious proxies) are accounted for in their AI systems. Moreover, different social values around privacy, data ownership, and data creation impact what AI technologies are possible. This session will dive into how the data policies developed today will shape the technologies of tomorrow.

Speakers:

Jennifer King, Privacy and Data Policy Fellow, Stanford HAI

An information scientist by training, Jennifer King is a recognized expert and scholar in information privacy. Sitting at the intersection of human-computer interaction, law, and the social sciences, her research examines the public’s understanding and expectations of online privacy as well as the policy implications of emerging technologies. Most recently, King’s research explored alternatives to notice and consent (with the World Economic Forum), the impact of California’s new privacy laws, and dark patterns. Her past work includes projects focusing on social media, genetic privacy, mobile application platforms, the Internet of Things (IoT), and digital surveillance. Her scholarship has been recognized for its impact on policymaking by the Future of Privacy Forum, and she has been invited to speak before the Federal Trade Commission several times. King has been featured in numerous publications and outlets, including the *New York Times*, the *Washington Post*, the *Los Angeles Times*, *Wired*, *Recode*, National Public Radio, CNBC, Bloomberg, *CNET*, *Vox*, *Consumer Reports*, NBC News, and *MIT Technology Review*.

King completed her doctorate in information management and systems at UC Berkeley. Prior to joining HAI, she was the director of consumer privacy at the Center for Internet and Society at Stanford Law School from 2018 to 2020. Before coming to Stanford, she was a co-director of the Center for Technology, Society, and Policy, a graduate student-led research center at UC Berkeley, and was a privacy researcher at the Samuelson Law, Technology, and Public Policy Clinic at Berkeley Law.

James Zou, Associate Professor of Biomedical Data Science and (by courtesy) of Computer Science and of Electrical Engineering, Stanford University; Faculty Affiliate, Stanford HAI

James Zou’s research focuses on making artificial intelligence more reliable, human-compatible, and statistically rigorous, with a particular emphasis on applications in human disease and health. He received his PhD from Harvard University and has served as a member of Microsoft Research, a Gates Scholar at Cambridge, and a Simons Fellow at UC Berkeley. Zou joined the Stanford faculty in 2016 and is also a Chan-Zuckerberg Investigator. His lab is part of the Stanford AI Lab, and his work has been recognized with two Chan-Zuckerberg Biohub Investigator Awards, a Sloan Research Fellowship, an National Science Foundation Career Award, a Top Ten Clinical Achievement Award, and faculty awards from Google, Adobe, and Amazon.



12:30 pm – 1:30 pm Session 2, continued



Alexander Spangher (moderator), Postdoctoral Researcher, Stanford University
Alex Spangher is a postdoctoral scholar at Stanford University, working with Diyi Yang, Sanmi Koyejo, and Daniel Ho. His research lies at the intersection of natural language processing, planning, and creativity. He models how humans execute creative and open-ended tasks to build AI systems that align better with societal goals. He received his PhD from the University of Southern California, co-advised by Jonathan May, Nanyun Peng, and Emilio Ferrara. Before that, he was a data scientist and journalist at the *New York Times*, where he worked with Dr. Chris Wiggins. Spangher studied double bass at the Juilliard School of Music and worked as a professional musician in off-Broadway productions, including *The Dybbuk* and *Spears: The Gospel According to Brittnay*. While working on his PhD, he organized charity concerts and raised money for humanitarian aid for Ukraine.

1:30 pm – 1:45 pm Break

1:45 pm – 2:45pm Session 3: The Role of Compute and Chips

This session highlights how computational power directly influences the capabilities and efficiency of AI systems, impacting everything from machine learning model training times to the sophistication of AI applications. Policymakers are introduced to key concepts such as high-performance computing chips, trade-offs between computational demands and energy consumption, and the strategic importance of compute in national competitiveness in AI.



Speakers:
Azalia Mirhoseini, Assistant Professor of Computer Science, Stanford University
Azalia Mirhoseini’s research interest is in developing capable, reliable, and efficient AI systems for solving high-impact, real-world problems. Her work includes generalized learning-based methods for decision-making problems in systems and chip design, self-improving AI models through interactions with the world, and scalable deep learning optimization. Prior to Stanford, she spent several years in industry AI labs, including Anthropic and Google Brain. At Anthropic, Mirhoseini worked on advancing the capabilities and reliability of large language models. At Google Brain, she co-founded the ML for Systems team, with a focus on automating and optimizing computer systems and chip design. She received her BSc in electrical engineering from Sharif University of Technology and her PhD in electrical and computer engineering from Rice University. Her work has been recognized with an *MIT Technology Review*’s 35 Under 35 Award, the Best ECE Thesis Award at Rice University, publications in flagship venues such as *Nature*, and coverage by various media outlets, including *MIT Technology Review*, *IEEE Spectrum*, *The Verge*, *The Times*, *ZDNet*, *VentureBeat*, and *Wired*.

1:45 pm - 2:45 pm Session 3, continued



H.-S. Philip Wong, Willard R. and Inez Kerr Bell Professor in the School of Engineering, Stanford University
H.-S. Philip Wong joined Stanford as professor of electrical engineering in 2004. From 1988 to 2004, he was with the IBM T.J. Watson Research Center. From 2018 to 2020, he took a leave from Stanford to serve as vice president of corporate research at TSMC, the largest semiconductor foundry in the world, and, since 2020, he remains the chief scientist of TSMC in a consulting, advisory role.

Wong is a Fellow of the IEEE and received the IEEE Andrew S. Grove Award, the IEEE Technical Field Award to honor individuals for outstanding contributions to solid-state devices and technology, as well as the IEEE Electron Devices Society J.J. Ebers Award, the society’s highest honor for outstanding technical contributions to the field of electron devices that have made a lasting impact.

He is the founding faculty co-director of the Stanford SystemX Alliance — an industrial affiliate program focused on building systems — and the faculty director of the Stanford Nanofabrication Facility, which is a shared facility for device fabrication on the Stanford campus that serves academic, industrial, and governmental researchers across the United States and around the globe, sponsored in part by the National Science Foundation. He is the principal investigator of the Microelectronics Commons California-Pacific-Northwest AI Hardware Hub, a consortium of over 40 companies and academic institutions funded by the CHIPS Act. He is a member of the U.S. Department of Commerce Industrial Advisory Committee on microelectronics.



Drew Spence (moderator), Policy Program Manager, Stanford HAI
Drew Spence leads HAI’s policy education programs, student fellowships, and outreach to policymakers in Washington and beyond. Prior to joining the team at HAI, Spence held multiple roles in government, most recently on Vice President Kamala Harris’ legislative affairs team. He also worked for the U.S. Department of the Treasury, in the U.S. Senate, and in the California Attorney General’s office. Spence is a graduate of George Washington University.

2:45 pm – 3:45 pm

Session 4: Mitigating Risk: Implementing Safe & Robust AI

The consequences of deploying robust AI and decision-making technologies in safety-critical systems such as driverless vehicles and autonomous aircraft are enormous. Challenges for AI developers range from biased inputs to constantly evolving conditions, explainability issues and more. This session will discuss the obstacles developers face, as well as the difficult — and often politically fraught — decisions they make around operational efficiency and how they define acceptable risk parameters.

Speakers:

Mykel Kochenderfer, Associate Professor of Aeronautics and Astronautics and (by courtesy) of Computer Science, Stanford University; Senior Fellow and Associate Director, Stanford HAI
Prior to joining the Stanford faculty, Mykel Kochenderfer was at MIT Lincoln Laboratory, where he worked on airspace modeling and aircraft collision avoidance, with his early work leading to the establishment of the ACAS X program. He received a PhD from the University of Edinburgh and BS and MS degrees in computer science from Stanford University. Kochenderfer is the director of the Stanford Intelligent Systems Laboratory (SISL), where he conducts research on advanced algorithms and analytical methods for the design of robust decision-making systems. Research at SISL focuses on efficient computational methods for deriving optimal decision strategies from high-dimensional, probabilistic problem representations. He is an author of *Decision Making Under Uncertainty: Theory and Application* (2015), *Algorithms for Optimization* (2019), and *Algorithms for Decision Making* (2022), all from MIT Press. He is a third generation pilot.

Oluwasanmi “Sanmi” Koyejo, Associate Professor of Computer Science, Stanford University; Faculty Affiliate, Stanford HAI

Sanmi Koyejo is an associate professor in the computer science department at Stanford University. He was previously an associate professor in the computer science department at the University of Illinois at Urbana-Champaign. Koyejo’s research interests are in developing the principles and practice of trustworthy machine learning, focusing on applications to neuroscience and healthcare. He completed a PhD at the University of Texas at Austin and postdoctoral research at Stanford University. Koyejo has been the recipient of several honors, including a best paper award from the conference on uncertainty in artificial intelligence, a Skip Ellis Early Career Award, a Sloan Fellowship, a Terman faculty fellowship, a National Science Foundation Career award, a Kavli Fellowship, an IJCAI early career spotlight, and a trainee award from the Organization for Human Brain Mapping. Koyejo spends time at Google as a part of the Brain team and serves on the Neural Information Processing Systems Foundation Board, the Association for Health Learning and Inference Board, and as president of Black in AI.

Tracy Navichoque (moderator), Program Manager for AI and Society, Stanford HAI
Before she joined HAI, Tracy Navichoque worked at the Stanford Global Digital Policy Incubator, providing human rights defenders with analysis around the governance of emerging technologies. She holds an MA in public diplomacy from USC and a BA in history and international studies from Northwestern University. She was a Fulbright Scholar to Uruguay.

3:45 pm – 4:00 pm

Break

4:00 pm – 5:00 pm

Session 5: The Governance of AI

Governments around the world are exploring governance initiatives regarding the ethical and responsible deployment of AI. From institutes focused on building government capacity to evaluate AI to laws that seek to place guardrails around model development, policymakers are experimenting with different approaches to regulating AI. This session will analyze the evolving landscape of regulations, guidelines, and best practices and provide insights into key considerations for effective governance.

Speakers:

Rishi Bommasani, Senior Research Scholar, Stanford HAI

Rishi Bommasani broadly focuses on the societal impact of AI, often by leading large-scale collaborations and building interdisciplinary teams, including an introduction to the foundation model landscape and the holistic evaluation of foundation models. More specifically, he works on questions of transparency, harm, governance, policy, and power in relation to foundation models.

Bommasani earned a BA in mathematics and computer science and an MS in computer science from Cornell University and a PhD in computer science from Stanford University. His research focuses on understanding AI systems and their societal impact, as well as using natural language processing to further scientific inquiry.

4:00 pm – 5:00 pm

Session 5, continued



Deborah Raji, Incoming Center Fellow at Stanford HAI

Deborah Raji is an Incoming Center Fellow at Stanford HAI. She is currently an Academic Fellow at the Leadership Conference on Civil and Human Rights, and was formerly a Senior Trustworthy AI Fellow at the Mozilla Foundation. She has worked closely with industry, civil society and within academia to push forward various projects to operationalize ethical considerations in machine learning practice, and push forward benchmarking and model evaluation norms in the field. She works on various topics related to the legal and institutional accountability required for machine learning systems to be deployed safely.

She is on the advisory boards for the Center for Democracy and Technology AI Governance Lab, the Health AI Partnership, TeachAI, REALML and the Center for Civil Rights and Technology, and others. For her efforts, she has been named to Forbes 30 Under 30, MIT Tech Review 35 Under 35 Innovators and the TIME 100 Most Influential in AI, is the recipient of the 2024 Tech For Humanity Prize, and the 2024 Mozilla Rise 25 award, as well as the co-recipient of the EFF Pioneer Barlow Award with Joy Buolamwini and Timnit Gebru. She received her Bachelors of Applied Science in Engineering Science from the University of Toronto and her PhD in computer science from the University of California, Berkeley.

Her work applies to a broad range of machine learning deployments — including automated decision systems, recommendation systems, and large pre-trained models. Her interests operate at the intersection of law and policy, applied economics and computer science.



Rob Reich, McGregor-Girand Professor of Social Ethics of Science and Technology, Stanford University; Senior Fellow and Associate Director, Stanford HAI

Rob Reich’s scholarship in political theory engages with the work of social scientists and engineers. His newest work is on governance of frontier science and technology. His most recent books are *System Error: Where Big Tech Went Wrong and How We Can Reboot* (with Mehran Sahami and Jeremy M. Weinstein) and *Digital Technology and Democratic Theory*. He has also written widely about philanthropy, including *Just Giving: Why Philanthropy is Failing Democracy and How It Can Do Better* and *Philanthropy in Democratic Societies: History, Institutions, Values*. His early work focused on democracy and education, including *Bridging Liberalism and Multiculturalism in American Education* and *Education, Justice, and Democracy*. He has testified before Congress and written for the *New York Times*, *Washington Post*, *Wired*, *Time*, *The Atlantic*, *The Guardian*, and the *Stanford Social Innovation Review*, among others outlets.

Reich is the recipient of the Walter J. Gores Award, Stanford’s highest honor for teaching. He was a sixth-grade teacher at Rusk Elementary School in Houston before attending graduate school. He is a board member of the Boston Review and the Spencer Foundation. He helped to create the global movement #GivingTuesday and served as its inaugural board chair. Reich was on public service leave in 2024-25 as senior adviser to the United States Artificial Intelligence Safety Institute.

4:00 pm – 5:00 pm

Session 5, continued



Daniel Zhang (moderator), Chief of Staff, Stanford HAI

Daniel Zhang is the chief of staff at the Stanford Institute for Human-Centered Artificial Intelligence (HAI). Previously, he led the institute’s policy research, outreach, and education initiatives. Zhang’s goal is to develop evidence-based AI policy recommendations, and his research interests lie at the intersection of technology policy, governance, and societal impact, including translational and original research on AI regulation and standards, the geopolitical implication of emerging technology, and the governance of large-scale ML models.

Zhang is also a member of the High-Level Expert Group on AI Ethics at UNESCO, where he advises the agency on the implementation of its Recommendation on the Ethics of AI. As the manager of the AI Index, he lead-authored the 2021 and 2022 annual reports that measure and evaluate the rapid rate of AI advancement.

Before Stanford, Zhang worked on global AI talent flows and security risks at the Center for Security and Emerging Technology and public education policy at the Riley Institute Center for Education and Leadership. He holds a master’s degree in security studies from Georgetown University’s Walsh School of Foreign Service, where he concentrated on technology policy, and a bachelor’s degree from Furman University.

5:00 pm – 7:30 pm

Dinner Keynote: Cutting Edge: AI, Automation, and the Future of Work

AI and automation will have a ripple effect on today’s workforce and the future of work. Mainstream narratives forecast that AI will displace workers and funnel profits up to a select few. Viewed differently, AI has the potential to augment and supercharge labor, ensuring the benefits of AI are spread and enjoyed widely. This session dives into deeper detail regarding what exactly we should expect as AI and automation integrate into the economy and the consequences for the workforce. The session will also explore how policies can shape and guide what the future holds.

Speaker:

Erik Brynjolfsson, Director, Stanford Digital Economy Lab; Jerry Yang and Akiko Yamazaki Professor and Senior Fellow and Associate Director, Stanford HAI; Ralph Landau Senior Fellow, Stanford Institute for Economic Policy Research

See biography on page 12.



TUESDAY, AUGUST 11, 2026

8:45 am – 9:00 am **Breakfast/Debrief**
Stanford HAI staff will lead a discussion debriefing the key concepts that staffers learned in earlier sessions. They will also offer a preview of Day Two and leave ample time for questions.

9:00 am – 10:00 am **Session 1: Transforming Healthcare Through Innovation**
Some of the most exciting advances produced by this technological wave are focused on healthcare: faster and better diagnoses, enhanced therapies, increased hospital standards that reduce patient harms, and protein folding and its potential to cure debilitating diseases. Healthcare is on the cusp of a revolution that will improve human well-being. At the same time, the United States faces a notable shortage of qualified healthcare workers, lacks proper evaluation of medical devices, and struggles to define liability risk and clinician responsibility. These mounting challenges raise the question: Can AI help “save” the U.S. healthcare system? This session will highlight the coming changes in healthcare, the opportunities and risks AI presents, and how policies can ensure safe and robust health systems.



Speakers:
Curtis Langlotz, Senior Associate Vice Provost for Research, Professor of Radiology, Medicine, and Biomedical Data Science, Stanford University; Director, Center for Artificial Intelligence in Medicine and Imaging; Senior Fellow, Stanford HAI
Curtis Langlotz’s laboratory investigates the use of deep neural networks and other machine learning technologies to detect disease and eliminate diagnostic errors through analysis of medical images and clinical notes. As associate chair for information systems and medical informatics director for Stanford Health Care, he is responsible for the computer technology that supports the Stanford radiology practice, including 20 million imaging studies that occupy two petabytes of storage.

Raised in St. Paul, Minnesota, Langlotz received his undergraduate degree in human biology, a master’s degree in computer science, an MD in medicine, and a PhD in medical information science, all from Stanford University. He is a founder and past president of the Radiology Alliance for Health Services Research and has served as president of the Society for Imaging Informatics in Medicine (SIIM) and the College of SIIM Fellows. He is a former board member of the Association of University Radiologists, the American Medical Informatics Association, and the Society for Medical Decision Making. He currently serves as president of the Radiological Society of North America.

He is a recipient of the Lee B. Lusted Research Prize from the Society of Medical Decision Making and the Career Achievement Award from the Radiology Alliance for Health Services Research. Langlotz and his trainees have received numerous scientific awards, including seven best paper awards and five research career development grants. He has founded several healthcare information technology companies, including Montage Healthcare Solutions, which was acquired by Nuance Communications in 2016.

9:00 am – 10:00 am **Session 1, continued**



Michelle Mello, Professor of Law and Health Policy, Stanford University; Faculty Affiliate, Stanford HAI
Michelle Mello is a leading empirical health law scholar whose research focuses on understanding the effects of law and regulation on healthcare delivery and population health outcomes. She is the author of more than 250 articles on medical liability, public health law, the public health response to COVID-19, pharmaceuticals and vaccines, artificial intelligence, data ethics and privacy, biomedical research ethics and governance, and other topics. Her publications appear in medical, health policy, and law journals, and she serves on the editorial boards of *JAMA Health Forum* and the *Journal of Health Politics, Policy and Law*.

In 2013, Mello was elected to the National Academy of Medicine (then called the Institute of Medicine), among the highest honors in the fields of health and medicine, in recognition of outstanding professional achievement and commitment to service. Her work garnered the Alice S. Hersh New Investigator Award from AcademyHealth, the leading professional organization for health services and health policy research in the United States; a Greenwall Faculty Scholars Award in Bioethics; a Robert Wood Johnson Foundation Investigator Award in Health Policy Research; an Open Science Champion Prize; and the John A. Benson Jr., MD Professionalism Article Prize.

Mello teaches Health Law: Improving Public Health and Torts. Prior to joining Stanford in 2014, she was a professor at the Harvard School of Public Health and director of its Program in Law and Public Health.



Nigam Shah, Professor of Medicine (Biomedical Informatics) and of Biomedical Data Science, Stanford University; Chief Data Scientist, Stanford Health Care; Faculty Affiliate, Stanford HAI
Nigam Shah’s research group analyzes multiple types of health data (EHR, claims, wearables, weblogs, and patient blogs) to answer clinical questions, generate insights, and build predictive models for the learning health system. At Stanford Health Care, he leads artificial intelligence and data science efforts for advancing the scientific understanding of disease, improving the practice of clinical medicine, and orchestrating the delivery of healthcare.

Shah is an inventor on eight patents and patent applications, has authored over 200 scientific publications, and has co-founded three companies. He was elected to the American College of Medical Informatics in 2015 and inducted into the American Society for Clinical Investigation in 2016. He holds an MBBS from Baroda Medical College, India, and a PhD from Penn State University, and he completed postdoctoral training at Stanford University.

9:00 am – 10:00 am **Session 1, continued**



Caroline Meinhardt (moderator), Policy Research Manager, Stanford HAI
Caroline Meinhardt develops and oversees policy research initiatives. She is passionate about harnessing AI governance research to inform policies that ensure the safe and responsible development of AI around the world — with a focus on research on the privacy implications of AI development, the implementation challenges of AI regulation, and the governance of large-scale AI models. Prior to joining HAI, Meinhardt worked as a China-focused consultant and analyst, managing and delivering in-depth research and strategic advice regarding China’s development and regulation of emerging technologies, including AI. She holds a master’s degree in international policy from Stanford University and a bachelor’s degree in Chinese studies from the University of Cambridge.

10:00 am – 10:05 am **Break**

10:05 am – 11:00 am **Stanford Robotics Center Visit**

This session offers a unique opportunity to witness firsthand the latest advancements in robotics research and development. Participants will learn the history of automation; explore innovative applications of robotics in various fields, from healthcare to autonomous systems; engage in demonstrations; and understand robotics’ potential impact on society.



Speaker:
Allison Okamura, Richard W. Weiland Professor in the School of Engineering, Professor of Mechanical Engineering, and (by courtesy) of Computer Science, Stanford University; Faculty Affiliate, Stanford HAI
Allison Okamura is director of graduate studies for Stanford Mechanical Engineering, deputy director of the Wu Tsai Neurosciences Institute, science fellow of the Hoover Institution, and executive committee member of the Stanford Robotics Center. She is an IEEE Fellow and was previously editor-in-chief of the journal *IEEE Robotics and Automation Letters*. Her honors include the IEEE Engineering in Medicine and Biology Society Technical Achievement Award, IEEE Robotics and Automation Society Distinguished Service Award, and Duca Family University Fellow in Undergraduate Education. She received her BS from UC Berkeley and PhD from Stanford. Her interests include haptics, teleoperation, mixed reality, medical robotics, soft robotics, and rehabilitation. Outside academia, she enjoys spending time with her family, running, and playing ice hockey.

11:00 am – 11:15 am **Walk back to Gates Building**

11:15 am – 12:15 pm

Session 2: AI in Motion: Role of Robotics in the Next Chapter of AI Development

As AI systems grow more advanced, their integration with robotics opens up new possibilities for how machines perceive, navigate, and act within the physical world. While large language models have accelerated progress in AI, embodied AI offers the potential to embed this intelligence into our daily lives in transformative ways. Could AI-powered robotics expand access to healthcare? Will autonomous vehicles fundamentally reshape our transportation systems? This session will explore the pivotal intersection of AI and robotics and examine the far-reaching implications across a range of policy domains.

Speakers:
Monroe Kennedy, Assistant Professor of Mechanical Engineering, Stanford University; Faculty Affiliate, Stanford HAI



Monroe Kennedy is the recipient of the NSF Career Award. He received his PhD in mechanical engineering and applied mechanics, and a master’s degree in robotics from the University of Pennsylvania, where he was a recipient of both the NSF and GEM graduate research fellowships. His area of expertise is in collaborative robotics, specifically the development of theoretical and experimental approaches to enhance robotic autonomy and robotic effectiveness in decentralized tasks toward human-robot collaboration. Kennedy applies expertise in machine learning, computer vision, collaborative robot teammate intent estimation, dynamical systems analysis, control theory (classical, non-linear, and robust control), state estimation and prediction, and motion planning.



Allison Okamura, Richard W. Weiland Professor in the School of Engineering, Professor of Mechanical Engineering, and (by courtesy) of Computer Science, Stanford University; Faculty Affiliate, Stanford HAI
See *biography on page 30*.



Drew Spence (moderator), Policy Program Manager, Stanford HAI
See *biography on page 23*.

12:15 pm – 1:30 pm Lunch

1:30 pm – 2:30 pm

Session 3: Revolutionizing the Classroom: How AI Is Advancing Education

AI has the potential to dramatically improve education. From teacher support to personalized student engagement, AI could democratize extraordinary teaching and learning. But dangers and concerns loom. Collecting data from children raises privacy concerns, and current inequities in the education system might be exacerbated by the introduction of AI. This session will look into how AI can be leveraged to improve the education system without causing harm to teachers or students.

Speakers:

Emma Brunskill, Associate Professor of Computer Science, Stanford University; Associate Director, Stanford HAI

Emma Brunskill’s goal is to create AI systems that learn from few samples to robustly make good decisions, motivated by applications to healthcare and education. Brunskill’s lab is part of the Stanford AI Lab, the Stanford Statistical ML group, and AI Safety @Stanford.

Brunskill was previously an assistant professor at Carnegie Mellon University. Her work has been honored by early faculty career awards from the National Science Foundation, Office of Naval Research, and Microsoft Research and, together with her lab members, she has received several best research paper nominations (CHI, EDMx3) and awards (UAI, RLDM, ITS). Brunskill is privileged to serve on the International Machine Learning Society Board, the Khan Academy Research Advisory Board, and the Stanford Faculty Women’s Forum Steering Committee. She previously served on the Women in Machine Learning board.

1:30 pm – 2:30 pm

Session 3, continued



Victor R. Lee, Associate Professor, Graduate School of Education, Stanford University; Faculty Lead, Stanford Accelerator for Learning Initiative on AI and Education; Faculty Affiliate, Stanford HAI

Victor R. Lee is the Khosla Family Professor in the Graduate School of Education at Stanford University and faculty lead for the Stanford Accelerator for Learning’s initiative on AI and Education. Through his research, he asks what future-facing STEM knowledge, tools, and practices are important to know to enable active participation and critical engagement with our increasingly digitally infused lives. He then uses the tools of educational design research to create examples for how we could get there. Currently, this involves researching and designing learning experiences and resources for data literacy, K-12 data science education, and artificial intelligence literacy for both students and teachers. He also maintains a portfolio of research related to elementary computer science education, maker education, and science cognition. His research is most often done through research-practice partnerships and involves design, implementation, analysis, and continual revision of new learning experiences in actual learning settings, such as schools, districts, or libraries. Lee has co-authored multiple national reports for the National Research Council related to computing and data in education. His work has been featured in the *New York Times*, CNN, *Forbes*, *Politico*, and other national media outlets.

Lee completed his undergraduate studies at UC San Diego with an emphasis in cognitive science, human computer interaction, and mathematics. He earned his doctorate in learning sciences at Northwestern University, where he was supported for several years through a fellowship with the NSF-funded Center for Curriculum Materials in Science. Since leaving the Midwest and beginning his academic career, Lee has received the National Science Foundation Career award, the Jan Hawkins Award, a postdoctoral fellowship from the National Academy of Education and the Spencer Foundation, and various best paper awards. His book, *Learning Technologies and the Body*, is the first compendium of current research of embodied technologies for learning. With Abigail Phillips, he published *Reconceptualizing Libraries: Perspectives from the Information and Learning Sciences* (2018). His most recent book, released in 2025, is *Advancing Data Science Education in K-12*. He is a past president and elected fellow of the International Society of the Learning Sciences.

1:30 pm – 2:30 pm

Session 3, continued



Isabelle Hau (moderator), Executive Director, Stanford Accelerator for Learning
Isabelle Hau partners with scholars, students, and external collaborators to leverage the latest advances in brain and learning sciences, data, and technology to create more effective and equitable solutions in learning for all.

Prior to Stanford, Hau spent a decade with the Omidyar Group, the philanthropic investment group of Pierre Omidyar, the founder of eBay, and his wife, Pam. Hau founded and led the U.S. Education initiative at Omidyar Network and became a founding partner of its spinoff, Imaginable Futures, the education venture of the Omidyar Group. There, she invested in organizations and projects that have reached millions of learners with impactful solutions. Behind these numbers are countless human stories of positive change, upward mobility, empowerment, and hope.

Hau joined the impact investing field in the early 2010s, when the concept was still being developed. Before joining Omidyar Network, she worked at ImpactAssets, one of the first nonprofit organizations dedicated to aligning capital with meaning and democratizing access to impact investing. She began her career as a global financier at Morgan Stanley, supporting the growth of organizations via financial capital tools.

2:30 pm – 3:30 pm

Session 4: Modernizing a Mammoth: Use-Cases of Public Sector AI

The U.S. government is in great need of a technological upgrade. From streamlining administrative processes to providing personalized services to constituents, there is ample opportunity for AI to help government agencies achieve their missions. However, integrating AI into the government is not as easy as obtaining and deploying the technology. Talent, infrastructure, public trust, and morale play equally important roles in ensuring the successful modernization of government. This session will dive into current use-cases of AI in government, the challenges and successes of these cases, and how to improve the integration of new technologies that will help the government serve its citizens.

Speakers:

Daniel Ho, William Benjamin Scott and Luna M. Scott Professor of Law, Stanford Law School; Director, Stanford Regulation, Evaluation, and Governance Lab (RegLab); Senior Fellow, Stanford Institute for Economic Policy Research; Senior Fellow and Associate Director, Stanford HAI



Daniel Ho served on the National Artificial Intelligence Advisory Committee, advising the White House on AI policy; as senior adviser on responsible AI at the U.S. Department of Labor; on the Committee on National Statistics of the National Academies of Science, Engineering, and Medicine; as a public member of the Administrative Conference of the United States; and as special adviser to the American Bar Association Task Force on Law and Artificial Intelligence. He is an elected member of the American Academy of Arts and Sciences.

His scholarship focuses on administrative law, regulatory policy, and antidiscrimination law. With RegLab, his work has developed high-impact demonstration projects of data science and machine learning in public policy through partnerships with a range of government agencies, including the Internal Revenue Service, the Treasury Department, the Environmental Protection Agency, the Department of Labor, the Santa Clara County Public Health Department, and Seattle and King County Public Health. The collaboration with Santa Clara County was awarded the Innovative Practice Gold Award for “the highest level of program innovation” to serve “community during the COVID-19 pandemic” by the National Association of County and City Health Officials.

Ho received his JD from Yale Law School and PhD from Harvard University, and he clerked for Judge Stephen F. Williams on the U.S. Court of Appeals, District of Columbia Circuit. He is the recipient of numerous awards, including the John Bingham Hurlbut Award for Excellence in Teaching at Stanford Law School; the Carole Hafner Award for the best paper at the International Conference on Artificial Intelligence and Law; the Best Empirical Paper Prize from the American Law and Economics Review; a Best Paper Award at the ACM Conference on Fairness, Accountability, and Transparency; the Best Paper Award at the AAAI/ACM Conference on AI, Ethics, and Society; and the Warren Miller prize for the best paper published in Political Analysis.

3:30 pm - 4:00 pm Walk to Li Ka Shing

4:00 pm – 5:00 pm **Session 5: The Role of Business: Policy Implications of Industry Leadership in AI**
Silicon Valley, home to venture capital, startups, and leading tech firms, is a global center of tech innovation. From the startup lab to the boardrooms of major companies, this session will bring together startup founders and tech executives to map out Silicon Valley’s innovation ecosystem, discuss its vibrancy, and think critically about the consequences of AI developments on society. Panelists will offer their perspectives on starting, funding, and running successful companies as well as providing counsel to companies on ensuring ethical business practices.

Speakers:

Sarah Guo, Founder and Managing Partner, Conviction

Sarah Guo was a general partner at Greylock before founding Conviction, an AI-native venture capital firm, in 2022. She has been an early investor and partner to more than 50 companies, including Figma, Harvey, Sierra, HeyGen, Mistral, Cognition, and Stackblitz. Guo is from Wisconsin, has four degrees from the University of Pennsylvania, and lives in the Bay Area with her husband and three kids. She co-hosts the AI podcast “No Priors” with Elad Gil.

Matt Perault, Head of Artificial Intelligence Policy, Andreessen Horowitz

Matt Perault oversees Andreessen Horowitz’s policy strategy on AI and helps portfolio companies navigate the AI policy landscape. He previously served as the director of the Center on Technology Policy at the University of North Carolina-Chapel Hill, as a consultant on technology policy at Open Water Strategies, and as the head of global policy development at Facebook. Perault holds a degree from Harvard Law School, a master’s degree from Duke University’s Sanford School of Public Policy and a bachelor’s degree in political science from Brown University.

Andy Song, Senior Director & Special Advisor, Google DeepMind

Andy Song is Senior Director & Special Advisor at Google DeepMind, where he leads initiatives across frontier AI strategy, governance, law, and policy. He previously served as Chief Operating Officer at Manifold, deploying AI to accelerate drug design and clinical trials for research organizations like the Broad Institute and American Cancer Society.

Before Manifold, Andy managed licensing, mergers & acquisitions, intellectual property, and litigation at Apple. He began his career at Munger, Tolles & Olson, the firm co-founded by Charlie Munger, after clerking on the U.S. Court of Appeals for the Ninth Circuit. He earned his undergraduate degree from Yale and J.D. from Harvard Law School.

4:00 pm – 5:00 pm Session 5, continued



Vanessa Parli (moderator), Managing Director of Programs and External Engagement, Stanford HAI

As managing director of programs and external engagement at the Stanford Institute for Human-Centered Artificial Intelligence (HAI), Vanessa Parli leads initiatives that foster interdisciplinary collaboration and connect academic research with real-world impact. Her work includes managing HAI’s industry research partnerships, executive education, policy engagement and the AI Index, creating opportunities for leaders from diverse backgrounds to engage with cutting-edge AI research, its applications, and its potential for positive social impact.

Prior to joining Stanford, Parli worked in management consulting, using statistics, machine learning, and data science to advise government agencies, biotech firms, and nonprofit organizations. She holds an MS in engineering management and computational mathematics from Johns Hopkins University and a BA in industrial engineering from Arizona State University.

5:00 pm – 6:00 pm Reception

6:00 pm – 7:30 pm

**Keynote Dinner: The New American Industrial Era:
Charting a New Age of U.S.-led Innovation**

As we look to a new era of AI development key conversations are happening around how we can build an AI stack that is designed and assembled in the United States. The discussion will examine how government, academia, the private sector, and civil society can collectively shape the future of AI with US leadership at the forefront. It will examine how today’s policy decisions may chart the course for innovation in the years ahead.

Speakers:



James Landay, Anand Rajaraman and Venky Harinarayan Professor in the School of Engineering, Professor of Computer Science, Stanford University; Denning Director, Stanford HAI
See biography on page 14.



Jensen Huang, Founder and Chief Executive Officer, Nvidia
See biography on page 13.



Fei-Fei Li, Sequoia Professor of Computer Science, Stanford University; Founding Co-Director, Stanford HAI
See biography on page 15.



Secretary Condoleezza Rice, Tad and Dianne Taube Director, Hoover Institution; Denning Professor in Global Business and the Economy, Stanford Graduate School of Business; Advisory Council Member, Stanford HAI
See biography on page 16.



Russell Wald (moderator), Executive Director, Stanford HAI
See biography on page 17.

WEDNESDAY, AUGUST 12, 2026

8:00 am – 8:30 am

Shuttle to SLAC National Accelerator Laboratory

8:30 am – 9:00 am

Breakfast/Debrief

Stanford HAI staff will lead a discussion reviewing key concepts from the first two days of the boot camp. They will also offer a preview of Day Three and leave ample time for questions.

9:00 am – 10:00 am

SLAC National Accelerator Laboratory Visit

This session will provide participants with the opportunity to see how researchers at a premier national laboratory on the Stanford campus are using technology to further scientific discovery. The visit will include stops at key SLAC facilities where materials and methods are developed for advances across physics, chemistry, biology, and medicine.



Speaker:

Alberto Salleo, Deputy Director for Science & Technology and Chief Research Officer, SLAC National Accelerator Laboratory
Alberto is responsible for advancing the lab’s Science & Technology (S&T) Strategy, identifying and realizing opportunities and strengthening the coordination and execution of S&T initiatives. He also serves as the lab’s chief research officer and oversees the Office of Strategic Planning and the Office of Technology Transfer and Strategic Partnerships. Prior to joining SLAC, he was an assistant professor at Stanford University, and was the Hong Seh and Vivian W. M. Lim Professor in the School of Engineering. He also served as chair of its Department of Materials Science and Engineering.

10:00 am - 10:30 am

Shuttle to Gates Building

10:30 am - 11:30am

Session 1: Unlocking Scientific Discovery: AI’s Potential to Accelerate Breakthroughs

AI is poised to fundamentally reshape the way scientific research is conducted — from predicting protein structures and modeling climate systems to uncovering new materials and advancing quantum research. Building on the SLAC facility tour, this session will highlight how AI is being applied across disciplines such as physics, chemistry, biology, and cosmology. Speakers will explore how these tools are accelerating the pace of discovery, enabling more precise experimentation, and opening new frontiers in our understanding of the natural world. Participants will gain insight into the transformative potential of AI to fuel the next era of scientific breakthroughs.

Speakers:

Russ B Altman, Kenneth Fong Professor and Professor of Bioengineering, of Genetics, of Medicine, of Biomedical Data Science, Stanford University; Senior Fellow and Associate Director, Stanford HAI

Russ B. Altman is the Kenneth Fong Professor of Bioengineering, Genetics, Medicine, Biomedical Data Science and (by courtesy) Computer Science and past chairman of the bioengineering department at Stanford University. His primary research interests are in the application of computing (AI, data science, and informatics) 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 (e.g., <http://www.pharmgkb.org/>). Other work focuses on the analysis of biological molecules to understand the actions, interactions, and adverse events of drugs (e.g., <http://helix.stanford.edu/>). He helps lead an FDA-supported Center of Excellence in Regulatory Science & Innovation.

Dr. Altman holds a BA from Harvard College, an MD from Stanford Medical School, and a PhD in medical information sciences from Stanford. He received the U.S. Presidential Early Career Award for Scientists and Engineers and a National Science Foundation Career Award. He is a fellow of the American College of Physicians, the American College of Medical Informatics, the American Institute of Medical and Biological Engineering, and the American Association for the Advancement of Science. He is a member of the National Academy of Medicine. He is a past president, founding board member, and fellow of the International Society for Computational Biology and a past president of the American Society for Clinical Pharmacology & Therapeutics. He chaired the Science Board advising the FDA commissioner and served on the NIH Director’s Advisory Committee and as co-chair of the IOM Drug Forum. He is an organizer of the annual Pacific Symposium on Biocomputing and a founder of Personalis. Dr. Altman is board certified in internal medicine and clinical informatics. He received the Stanford Medical School graduate teaching award in 2000 and 2020 and the mentorship award in 2014. He is the founding editor of the *Annual Review of Biomedical Data Science* and hosts the “Future of Everything” podcast.

10:30 am - 11:30am

Session 1, continued




Ramesh Johari, Professor of Management Science and Engineering and (by courtesy) of Electrical Engineering, Stanford University; Associate Director, Stanford HAI

Ramesh Johari is a professor at Stanford University, with a full-time appointment in the department of management science and engineering (MS&E) and a courtesy appointment in the department of electrical engineering (EE). He is an associate director of Stanford Data Science. He is a member of the Operations Research group and the Social Algorithms Lab (SOAL) in MS&E, the Information Systems Laboratory in EE, and the Institute for Computational and Mathematical Engineering. He received a BA in mathematics from Harvard, a certificate of advanced study in mathematics from Cambridge, and a PhD in electrical engineering and computer science from MIT.

He is the recipient of a British Marshall Scholarship, First Place in the INFORMS George E. Nicholson Student Paper Competition, the George M. Sprowls Award for the best doctoral thesis in computer science at MIT, Honorable Mention for the ACM Doctoral Dissertation Award, the Okawa Foundation Research Grant, the MS&E Graduate Teaching Award, the INFORMS Telecommunications Section Doctoral Dissertation Award, the NSF Career Award, and the Cisco Faculty Scholarship. He has served on the program committees of ACM Economics and Computation (including co-chairing the conference in 2019), ACM SIGCOMM, IEEE Infocom, and ACM SIGMETRICS, as the track chair for the Internet Economics and Monetization Track at WWW, and as a founding co-organizer of the Marketplace Innovation Workshop. He served as founding Area Co-Editor of the Revenue Management and Market Analytics Area for Operations Research and as associate editor for Management Science (in the Stochastic Models and Simulation area) and Stochastic Systems.

Elena Cryst (moderator), Director of Policy and Society, Stanford HAI
See biography on page 20.





11:30 am – 12:30 pm

Session 2: Powering AI: Meeting the Energy Demand Needed for Frontier AI
AI continues to increasingly rely on compute power, which is an incredibly energy-intensive resource. There is mounting concern around AI’s environmental impact, drawing attention to the cost-benefit analysis of AI advancements. Might it be possible to find ways to power AI research more efficiently? Could technological advances in energy storage and transmission help reduce the intense energy demands that frontier research currently requires? This session will unpack the realities of AI’s energy needs and how we can develop technology more responsibly in the future.



Speakers:
William Chueh, Director, Precourt Institute for Energy; Kimmelman Professor, Professor of Materials Science and Engineering, of Energy Science and Engineering, of Photon Science, Stanford University
Will Chueh is a Kimmelman Professor in the department of materials science & engineering and energy science & engineering at Stanford University, department of photon science at SLAC National Accelerator Laboratory, and the director and senior fellow of the Precourt Institute for Energy. He leads a group of more than 30 researchers who aim to understand reactions and transport involving ions and electrons and to decarbonize various energy transformation pathways. He received his BS in applied physics and his MS and PhD in materials science from Caltech. Prior to joining Stanford in 2012, he was a Distinguished Truman Fellow at Sandia National Laboratories. Chueh has received numerous honors, including the David A. Shirley Award (2023), Friedrich Wilhelm Bessel Research Award (2022), MRS Outstanding Young Investigator Award (2018), Volkswagen/BASF Science Award Electrochemistry (2016), Camille Dreyfus Teacher-Scholar Award (2016), Sloan Research Fellowship (2016), NSF Career Award (2015), Solid State Ionics Young Scientist Award (2013), and Caltech Demetriades-Tsafka-Kokkalis Prize in Energy (2012). In 2012, he was named one of the “Top 35 Innovators Under the Age of 35” by *MIT Technology Review*. He serves on the editorial boards of numerous journals, including *Energy & Environmental Science*, and is the editor-in-chief of *Solid State Ionics*.

11:30 am – 12:30 pm

Session 2, continued



Liang Min, Managing Director, Bits & Watts Initiative, Precourt Institute for Energy
Liang Min leads a multidisciplinary program advancing the digital transformation of the electric grid. Under his direction, Bits & Watts has launched several pioneering research efforts, including 100% Clean Electric Grid, EV50, AI for Climate and Energy, and the Digital Grid — an open platform designed to integrate distributed energy resources at scale. Most recently, he launched the Powering AI Sustainably program, a cross-sector program addressing the surging energy demands of artificial intelligence while accelerating the transition to a clean and reliable power system.

Min also serves as the founding managing director of the Stanford Net-Zero Alliance, a global platform that brings together industry leaders, faculty, and students to collaborate on research and education aligned with a net-zero emissions future. He also established the Stanford Energy Executive Education Program, which equips senior energy leaders with the strategic insight and tools needed to navigate the rapidly evolving energy landscape.

His career spans leadership roles across academia, national laboratories, and industry. At the Electric Power Research Institute, Min led the grid operations and planning team and directed phasor measurement unit research, resulting in multiple U.S. patents deployed at utilities through the support of the American Recovery and Reinvestment Act. At Lawrence Livermore National Laboratory, he founded the Energy Delivery Group and served as associate program leader for the lab’s national cyber and infrastructure resilience program. He also directed the Electric Operations Program for California’s Energy Systems for the 21st Century — a \$150 million public-private initiative leveraging high-performance computing to enhance grid reliability, security, and value to ratepayers.

Outside of work, Min is an avid runner with a personal best marathon time of 2 hours and 50 minutes.

12:30 pm – 1:00 pm

Lunch Break

1:00 pm – 2:30 pm

Session 3: AI and National Security: Strategic Competition Lessons for the United States

AI is introducing new opportunities to strengthen U.S. intelligence and national security capabilities. However, it is crucial for policymakers to grasp the inner workings of both the U.S. national security framework and AI technologies to leverage these advancements effectively. Modern espionage, for example, is more pervasive yet less understood than ever, leading to misinformation and policy missteps. This session will separate fact from fiction as panelists discuss the past, present, and future of American espionage and how AI is creating an adapt-or-fail moment for international security and U.S. intelligence.

Speakers:

Yejin Choi, Dieter Schwarz Foundation Professor of Computer Science, Stanford University; Senior Fellow, Stanford HAI

Yejin Choi is the Dieter Schwarz Foundation HAI Professor, a senior fellow and associate director at HAI, and a professor of computer science at Stanford University. She is a MacArthur Fellow (class of 2022) and an AI2050 Senior Fellow (class of 2024), and she was named among the Time100 Most Influential People in AI in 2023. In addition, Choi is the co-recipient of two Test-of-Time awards and eight Best and Outstanding Paper Awards at top AI conferences, including ACL, ICML, NeurIPS, ICCV, CVPR, and AAAI. She also won the Borg Early Career Award in 2018, the inaugural Alexa Prize Challenge in 2017, and IEEE AI's 10 to Watch in 2016. Choi was a main stage speaker at TED 2023 and a keynote speaker for conferences across several AI disciplines, including ACL, CVPR, ICLR, MLSys, VLDB, WebConf, and AAAI. Her current research interests include fundamental limits and capabilities of large language models, alternative training recipes for language models, symbolic methods for neural networks, reasoning and knowledge discovery, moral norms and values, pluralistic alignment, and AI safety.

1:00 pm – 2:30 pm

Session 3, continued



Colin Kahl, Director, Freeman Spogli Institute for International Studies; Steven C. Házzy Senior Fellow, Center for International Security and Cooperation (CISAC); Professor, by courtesy, of Political Science, Stanford University

From April 2021 to July 2023, Colin Kahl served as the under secretary of defense for policy at the U.S. Department of Defense (DOD). In that role, he was the principal adviser to the secretary of defense for all matters related to national security and defense policy and represented the DOD as a standing member of the National Security Council Deputies' Committee. He oversaw the writing of the 2022 National Defense Strategy, which focused the Pentagon's efforts on the "pacing challenge" posed by China, and he led the DOD's response to Russia's invasion of Ukraine and other international crises. In June 2023, Secretary of Defense Lloyd Austin III presented Kahl with the Department of Defense Distinguished Public Service Medal, the highest civilian award presented by the secretary of defense.

Kahl served as deputy assistant to President Barack Obama and national security advisor to Vice President Joe Biden from October 2014 to January 2017. He also served as deputy assistant secretary of defense for the Middle East from February 2009 to December 2011, a role for which he received the Outstanding Public Service Medal in July 2011.

Kahl is the co-author of *Aftershocks: Pandemic Politics and the End of the Old International Order* and the author of *States, Scarcity, and Civil Strife in the Developing World*. He previously taught at Georgetown University and the University of Minnesota, and he held fellowship positions at Harvard University, the Council on Foreign Relations, CNAS, and the Penn Biden Center for Diplomacy and International Engagement.

Kahl received his BA in political science from the University of Michigan and his PhD in political science from Columbia University.

1:00 pm – 2:30 pm

Session 3, continued



Amy Zegart, Morris Arnold and Nona Jean Cox Senior Fellow, Hoover Institution; Senior Fellow, Freeman Spogli Institute of International Studies; Senior Fellow and Associate Director, Stanford HAI

The author of multiple books, Amy Zegart specializes in U.S. intelligence, emerging technologies and national security, grand strategy, and global political risk management.

Zegart’s award-winning research includes the leading academic study of intelligence failures before 9/11: *Spying Blind: The CIA, the FBI, and the Origins of 9/11*. Her most recent book is the bestseller *Spies, Lies, and Algorithms: The History and Future of American Intelligence*, which was nominated for the Pulitzer Prize. She co-authored *Political Risk: How Businesses and Organizations Can Anticipate Global Insecurity* with Condoleezza Rice and co-edited *Bytes, Bombs, and Spies: The Strategic Dimensions of Offensive Cyber Operations* with Herbert Lin. Her op-eds and essays have appeared in *Foreign Affairs*, *Politico*, the *New York Times*, the *Washington Post*, and the *Wall Street Journal*.

Zegart has advised senior officials about intelligence and foreign policy for more than two decades. She served on the National Security Council staff and as a presidential campaign foreign policy adviser and testified before the House and Senate Intelligence Committees.

Zegart led Stanford’s Center for International Security and Cooperation, founded the Stanford Cyber Policy Program, and served as chief academic officer of the Hoover Institution. Before coming to Stanford, she was professor of public policy at UCLA and a McKinsey & Company consultant.

She is the recipient of a Fulbright Fellowship, the American Political Science Association’s Leonard D. White Dissertation Prize, and research grants from the Carnegie Corporation of New York, the Hewlett Foundation, the Smith Richardson Foundation, and the National Science Foundation.

A native of Louisville, Kentucky, Zegart received a BA in East Asian studies, magna cum laude, from Harvard and an MA and a PhD in political science from Stanford. She serves on the boards of the Council on Foreign Relations, Kratos Defense & Security Solutions, and the American Funds/Capital Group.

1:00 pm – 2:30 pm

Session 3, continued



Brad Boyd (moderator), Senior Advisor and Founder, The Future of Decision-making Project; Senior Military Fellow, The Gordian Knot Center (GKC), Stanford University

A former visiting fellow at the Hoover Institution, Brad Boyd focuses his research on the integration of emerging technology into warfare and national security. He is particularly interested in the rise of automation and autonomy in military systems, including in decision making, planning, information operations, enterprise operations, and weaponry. Boyd also looks at the way U.S., Chinese, and Russian military integration of emerging technology affects military, economic, and social stability. Prior to joining the Hoover Institution, Boyd served as a defense and foreign policy advisor to Senator Angus King and as Senator King’s representative to the Cyberspace Solarium Commission. Boyd also served as the director of AI-enabled warfighting capability development at the U.S. Department of Defense’s Joint Artificial Intelligence Center and as the director of General Mark A. Milley’s coordination group. Boyd was a senior military fellow at Stanford’s Center for International Security and Cooperation and a Seminar XXI Fellow at the Massachusetts Institute of Technology.

2:45 pm – 3:45 pm

Session 4: Updating the Law in the Age of AI

Tech policymaking has long faced the challenge of using old legal frameworks to regulate new technologies. Now, AI has the potential to supercharge that problem. How do we take legal frameworks written before the internet was invented and make them relevant in the AI era. This panel will explore how we should think about evolving legal structures to respond to new and emerging technologies.

Speakers:

Nathaniel Persily, James B. McClatchy Professor of Law, Stanford Law School; Co-Director of the Stanford Law AI Initiative; Faculty Affiliate, Stanford HAI

Nate Persily is the James B. McClatchy Professor of Law at Stanford Law School and the Co-Director of the Stanford Law AI Initiative. He also holds appointments in the departments of political science, communication, and at the Freeman Spogli Institute of International Studies. Persily’s work focuses on two distinct areas: the law of democracy and governance of technology. He is co-author, with Samuel Issacharoff, Pamela Karlan, Richard Pildes, and Franita Tolson, of *The Law of Democracy* (2020), which addresses issues such as voting rights, political party regulation, campaign finance, redistricting, and election administration. He has served as a special master or court-appointed expert to craft congressional or legislative districting plans for Georgia, Maryland, Connecticut, New Hampshire, New York, North Carolina, and Pennsylvania. He also served as the nonpartisan senior research director for the Presidential Commission on Election Administration. His current work, for which he has been honored as a Guggenheim Fellow, Andrew Carnegie Fellow, and a Fellow at the Center for Advanced Study in the Behavioral Sciences, examines the governance of technology and technology’s impact on democracy. He is co-editor, with Joshua Tucker, of *Social Media and Democracy: The State of the Field and Prospects for Reform* (2020). He is also co-editor of the *Digitalist Papers: Artificial Intelligence and Democracy in America* (2024) with Erik Brynjolfsson, Alex Pentland and Condoleezza Rice. Persily co-chairs the Presidential Task Force on Artificial Intelligence of the American Political Science Association. He is a member of the American Academy of Arts and Sciences. He received a BA and an MA in political science from Yale; a JD from Stanford, where he was president of the *Stanford Law Review*; and a PhD in political science from UC Berkeley.



2:45 pm – 3:45 pm

Session 4, continued

Riana Pfefferkorn, Policy Fellow, Stanford HAI

Riana Pfefferkorn focuses on investigating and analyzing U.S. and other governments’ policies and practices for forcing decryption and/or influencing crypto-related design of online platforms and services, devices, and products, both via technical means and through the courts and legislatures. She also researches the benefits and detriments of strong encryption on free expression, political engagement, economic development, and other public interests.

Pfefferkorn was a research scholar at the Stanford Internet Observatory. She was also the associate director of surveillance and cybersecurity at the Stanford Center for Internet and Society. Prior to joining Stanford, Pfefferkorn was an associate in the internet strategy and litigation group at Wilson Sonsini Goodrich & Rosati, where she worked on litigation and counseling matters involving online privacy, internet intermediary liability, consumer protection, copyright, trademark, and trade secrets and was actively involved in the firm’s pro bono program. Before that, she clerked for the Hon. Bruce J. McGiverin of the U.S. District Court for the District of Puerto Rico. Pfefferkorn interned during law school for the Hon. Stephen Reinhardt of the U.S. Court of Appeals for the Ninth Circuit. She earned her law degree from the University of Washington School of Law and her undergraduate degree from Whitman College.



2:45 pm – 3:45 pm

Session 4, continued



Irene Liu, Executive Director, Stanford Law School AI Initiative

Irene Liu is the executive director of the Stanford Law School AI Initiative, which serves as a hub for the law school’s interdisciplinary work in artificial intelligence. She brings experience across academia, government, and technology to guide the initiative’s work in teaching, training, research, partnerships, and public engagement in artificial intelligence.

Prior to joining Stanford Law, she served as an AI adviser to the California Senate and as an adviser to UC Berkeley School of Law’s AI Institute, where she also taught on emerging topics in generative AI. She is the founder and CEO of Hypergrowth GC, where she advises technology CEOs and executive teams on strategic and legal decisions to support company growth. Through these and other roles, Liu has long focused on the governance, business, and policy implications of technology, with an emphasis on responsible AI development and deployment.

She has held a number of senior leadership roles at technology companies. Liu was the chief financial and legal officer of Hopin, where she oversaw legal, finance, policy, and trust and safety functions during a period of rapid global expansion. She also served as general counsel and senior vice president at Checkr, where she led legal, compliance, policy, and customer education teams and advised on corporate governance, regulatory compliance, and company growth. She led global law enforcement compliance at BlackBerry and held senior legal roles at Lookout, an enterprise mobile security company, where she worked on international privacy, security, and regulatory compliance issues.

Before entering the private sector, Liu served as a trial attorney in the U.S. Department of Justice Antitrust Division and the Federal Trade Commission’s Bureau of Consumer Protection.

Her work has been recognized with honors including General Counsel of the Year by Corporate Counsel, Top 50 Women Lawyer by the National Diversity and Leadership Conference, and In-House Leader by the *Los Angeles Times*. She has spoken at academic institutions, government hearings, and industry conferences, including Fortune’s Most Powerful Women Next Gen Conference, the Federal Reserve, and a Federal Trade Commission public hearing on artificial intelligence.

She has hosted the Hypergrowth GC video series on Luminate+ and “Coffee Break with Irene” on Berkeley Boost. She also served as a foundational adviser to Women Defining AI.

Liu received her JD from UC Berkeley School of Law and her BA from Mount Holyoke College.

3:45 pm – 4:00 pm

Break

4:00 pm – 5:00 pm

Session 5: Agents on the Rise: Exploring Agentic AI

What if AI could act on our behalf to make our travel arrangements, order our prescription refills, and take care of several other tedious daily tasks? AI agents represent the promise of AI to enhance productivity and reduce friction in our daily routines. Yet, this emerging capability also raises concerns around labor disruption, security vulnerabilities, and delegation of control. This session will examine how agentic AI works, how rapidly the technology is evolving, and what it might mean for the economy, the workforce, and society at large.

Speakers:

Alex “Sandy” Pentland, HAI Fellow and Faculty Lead for Digital Platforms and Society, Stanford Digital Economy Lab



Alex Pentland is an HAI Fellow at Stanford and the Toshiba Professor at MIT. He is a member of the U.S. National Academy of Engineering and has won numerous awards and prizes, including the 40th Anniversary of the Internet from DARPA, the Brandeis Privacy Award, and AI Influencer Lifetime Achievement Award, as well as many scientific awards. He is faculty lead for Digital Societies in the Stanford Digital Economy Lab.

Diyi Yang, Assistant Professor of Computer Science, Stanford University



Diyi Yang’s research interests are computational social science and natural language processing. Her research goal is to understand the social aspects of language and then build socially aware NLP systems to better support human-human and human-computer interaction. Yang received her PhD from the School of Computer Science at Carnegie Mellon University and her bachelor’s degree from Shanghai Jiao Tong University in China. Her work has received multiple best paper nominations or awards at ICWSM, EMNLP, SIGCHI, ACL, and CSCW. She is a recipient of Forbes 30 under 30 in Science, IEEE “AI 10 to Watch,” the Intel Rising Star Faculty Award, Microsoft Research Faculty Fellowship, and NSF Career Award.

Andy Zhang (moderator), Graduate Student Fellow, Stanford Regulation, Evaluation, and Governance Lab



Andy Zhang is a researcher at Stanford University focused on developing and benchmarking cybersecurity agents. He previously led Cybench, a framework to evaluate cybersecurity capabilities and risk of language model agents, and he investigated the train-test overlap transparency of language models.

5:00 pm – 5:15 pm

Break

5:15 pm – 7:30 pm

Fireside Chat

Cutting Edge: AI Developments That Are Steering the Future

The recent, rapid advancements in AI have shocked the world. From models generating photorealistic images to ambient technologies that enhance the human condition, the possibilities of what AI can do for humanity are endless. Understanding today’s cutting-edge AI will help steer tomorrow’s innovation. This session will dive into what is on the horizon of AI advancements and how these technologies can be leveraged to benefit society.

Speakers:

Mariano-Florentino “Tino” Cuéllar, Sara Miller McCune Director, Center for Advanced Study in the Behavioral Sciences (CASBS); Shriram Family Director, Knight-Hennessy Scholars (KHS); Senior Fellow, Stanford HAI
See biography on page 12.

John Hennessy, President Emeritus, James F. and Mary Lynn Gibbons Professor and Professor of Electrical Engineering and of Computer Science, Stanford University
See biography on page 13.

Elena Cryst (moderator), Director of Policy and Society, Stanford HAI
See biography on page 20.

Staffer Biographies

Anisa Alazraie, Legislative Assistant, Office of Rep. Ted Lieu (D-CA)

Anisa Alazraie is Congressman Ted Lieu’s legislative assistant covering AI, tech, and IP. She joined Lieu’s office in September 2025, after working for Senator Alex Padilla for three years. Alazraie was born and raised in the Bay Area, where she did small business advocacy before moving to Washington, D.C., to work in Congress.

Shinnola Alexander, Policy Advisor, Office of Rep. Sara Jacobs (D-CA)

Shinnola Alexander is a policy adviser in the office of Congresswoman Sara Jacobs, where she leads policy development and legislative strategy on artificial intelligence, as well as energy, environment, agriculture, transportation, and infrastructure. She previously served as a legislative assistant, covering both domestic and foreign policy. Before Capitol Hill, she held roles in both government and the private sector. She holds a BA in political science from Yale University and an MS in environmental policy and regulation from the London School of Economics and Political Science.

Shaefer Bagwell, Legislative Director, Office of Rep. Suhas Subramanyam (D-VA)

Shaefer Bagwell serves as the legislative director to Rep. Suhas Subramanyam of Virginia’s 10th District. Over 11 years on Capitol Hill, Bagwell has worked for members of the California delegation in both the House and Senate, as well as then-Rep. Mikie Sherrill of New Jersey. His policy work has focused primarily on defense and foreign affairs issues, with a concentration on munitions development, emerging defense technologies, and international security matters. Bagwell holds a BA in political science from Biola University in La Mirada, California. He was born and raised in Sacramento, California, and has lived in Washington, D.C., since 2015. In his free time, he enjoys reading, cooking, learning about wine, and making fun of his dog.



Joel Burke, AI Policy Advisor, Office of Sen. Mike Rounds (R-SD)

Joel Burke is a tech policy expert with experience leading teams and providing strategic advice in contexts as diverse as Silicon Valley, Estonia, and Cameroon. He currently leads the AI portfolio for a senior senator on the Armed Services Committee, supporting the senator’s efforts to advance the use of AI for the American warfighter and to accelerate American science. Prior to his work in D.C., Burke led business development and partnerships for the Republic of Estonia’s e-Residency program, was a partner at Tribe AI, where he helped enterprises adopt AI, and was an early employee at a YC/Andreessen backed startup. In his free time, Burke is passionate about writing and recently authored *Rebooting a Nation* about Estonia’s modernization post re-independence.



Will Burns, Senior Legislative Assistant, Office of Rep. Jay Obernolte (R-CA)

Will Burns is the senior legislative assistant for Congressman Jay Obernolte (CA-23). He is the office lead handling the artificial intelligence (AI), emerging technologies (robotics, quantum, etc.), national security, and foreign affairs policy portfolios. Burns worked closely with Representative Obernolte during his time as chairman of the Bipartisan Artificial Intelligence Task Force, and helped work on Task Force hearings as well as in drafting, negotiating, and publishing the final report. In addition, Burns works closely with relevant AI and technology policy committees, including the House Committee on Science, Space, and Technology (on which Obernolte is chairman of the Subcommittee on Research and Technology) and the Committee on Energy and Commerce. He also engages with Senate Member offices and committee staff (Commerce, Science, and Transportation Committee, Judiciary Committee, etc.), relevant White House and Executive Branch agency offices (including OSTP, NIST, NSF, Commerce Department, Energy Department, and others), and key industry and academic stakeholders on developing, writing, and advancing legislative priorities in the AI and emerging technology policy spaces. Burns is currently working on a significant number of AI policy issues and legislative initiatives, including bills targeting the development and implementation of a national framework for AI that focuses on federal preemption, model regulation, research security, consumer protection, national security and international competitiveness implications, education, and driving innovation and adoption. A native of Washington, D.C., he has worked for Representative Obernolte since October 2021. Prior to his current position, he served as the congressman’s military legislative assistant. Before getting his start on Capitol Hill, Burns received dual undergraduate degrees in international relations and medieval history from the University of St. Andrews in Scotland.



Brandon Gindt, National Security Advisor, Office of Sen. Ted Cruz (R-TX)

Brandon Gindt serves as Senator Ted Cruz’s national security advisor. Prior to joining Senator Cruz’s office, he worked in the Directorate of Operations at the Defense Intelligence Agency (DIA), where he worked on Middle East and strategic competition issues. During his tenure at DIA, he deployed to Afghanistan with a special operations joint task force and managed crisis operations on the Joint Staff. While on the Joint Staff, he served as the executive briefer to the Chairman of the Joint Chiefs of Staff. Later, Gindt served as the Israel and Lebanon country director in the Office of the Under Secretary of Defense for Policy. He received his undergraduate and graduate degrees from Texas A&M University.



Angelica Espinoza, Senior Legislative Assistant, Office of Rep. Lauren Underwood (D-IL)

Angelica Espinoza serves as the legislative director to Congresswoman Lauren Underwood, where she leads the appropriations, foreign affairs, defense, homeland security, and cybersecurity portfolio. She previously served as shared committee staff on the House Appropriations Committee under Ranking Member Lauren Underwood on the Homeland Security Subcommittee, where she led work on emergency management, border security, immigration, transportation and maritime security, and broader national security issues. Prior to joining Congresswoman Underwood’s office, Espinoza served as a legislative fellow for the late Congressman Raúl Grijalva of Arizona’s 7th district. She earned an MS in homeland security from San Diego State University in 2022 and a BA in international security and conflict resolution in 2020.



Michael Harris, Energy and Environment Policy Advisor, Office of Sen. Adam Schiff (D-CA)

Michael Harris is Senator Adam Schiff’s energy and environment policy advisor. Previously, Harris spent three years as a political appointee in the Biden-Harris administration, where he advised two members of the president’s Cabinet on congressional engagement and legislative strategy. At the U.S. Environmental Protection Agency, he managed congressional oversight for the agency, and he led House affairs for the U.S. Department of Energy. Prior to his service in the executive branch, Harris directed Congresswoman Debbie Wasserman Schultz’s work on the House Appropriations Committee and helped her secure funding to fight climate change and restore ecosystems. Before his time in the House, he advocated for the protection of public lands at Defenders of Wildlife. He began his environmental policy career in the office of Senate Democratic Leader Harry Reid. Harris received his BA in public policy studies from Duke University and his MPP from the Trachtenberg School.



Eric Henshall, Deputy Chief of Staff, Office of Rep. Sam Liccardo (D-CA)

Eric Henshall is the deputy chief of staff for Congressman Sam Liccardo, who represents part of Silicon Valley in Congress. Henshall has been in his current role since Representative Liccardo took office in January 2025 and serves as his policy adviser on technology, telecommunications, science, education, transportation, and appropriations. Previously, Henshall spent eight years in the office of former Congresswoman Anna Eshoo, most recently as legislative director. He earned a bachelor's degree in government and economics from Georgetown University and is originally from the San Francisco Bay Area.



Sydney Hess, Legislative Assistant, Office of Sen. Mark Kelly (D-AZ)

Sydney Hess serves as a legislative assistant for Senator Mark Kelly, where her work spans a range of domestic policy issues, including space and science, banking, public safety, homeland security and immigration, and appropriations. She supports Senator Kelly's work on artificial intelligence, including the release of his AI for America road map, which focuses on ensuring that the benefits of AI are broadly shared through investments in workers, infrastructure, and responsible innovation. Hess has worked for Senator Kelly for four years. Before joining his office, she worked at a nonprofit focused on higher education and student leadership. An Arizona native, Hess earned her bachelor's degree in philosophy, politics, economics, and law from the University of Arizona.



Talia Katz, Policy Advisor, Office of Sen. Dave McCormick (R-PA)

Talia Katz Fortgang is a policy adviser for U.S. Senator Dave McCormick (R-PA), where she oversees the commerce, science, and transportation portfolio.



Leah Li, Legislative Assistant, Office of Sen. Brian Schatz (D-HI)

Leah Li is a legislative assistant to Senator Brian Schatz of Hawaii, covering a legislative portfolio that includes technology, telecommunications, consumer protection, trade, and government modernization. Previously, she was senior adviser to Congressman Derek Kilmer of Washington. Li received a BA from the University of Washington, Seattle, and a master of international public policy from the Johns Hopkins School of Advanced International Studies.



Geornailer Marcial Acaba, Health Legislative Assistant, Office of Rep. Michael Rulli (R-OH)

Geornailer Marcial serves as Congressman Rulli's Health legislative assistant, handling issues before the House Committee on Energy and Commerce Health Subcommittee. Prior to joining Congressman Rulli's office, Marcial served as Congressman Randy Feenstra's legislative correspondent and worked as a legislative fellow in the House Committee on Natural Resources. In 2020, Marcial graduated from Florida International University with a BA in political science and a minor in international relations.



Kaitlin McNiffe, Professional Policy Staffer, House Republican Study Committee

Katie McNiffe is a professional policy staffer for the Republican Study Committee (RSC), the oldest and largest conservative caucus in the House of Representatives. She works to advance conservative policy priorities across a broad portfolio, including energy, environment, transportation, infrastructure, natural resources, telecommunications, and interstate commerce. Prior to joining the RSC, McNiffe spent four years with her hometown representative, Congressman Richard Hudson, where she developed expertise in legislative strategy, constituent engagement, and coalition building. A proud alumna of the University of North Carolina at Chapel Hill, she earned degrees in political science and philosophy while actively engaging in extracurricular pursuits that laid the foundation for her career in public service and politics. McNiffe lives in Falls Church, Virginia, with her husband and their dog.



Charlie Mleck, Legislative Assistant, Office of Rep. Bob Latta (R-OH)

Charlie Mleck is a legislative assistant to Congressman Bob Latta (R-OH-05), where he leads the office's work on artificial intelligence and emerging technology policy, along with commerce, manufacturing, and trade issues. Before joining Congressman Latta's staff, Mleck served as an AI policy fellow with the House Energy and Commerce Committee, working with the majority staff of the Commerce, Manufacturing, and Trade Subcommittee. A native of Monclova, Ohio (just outside of Congressman Latta's district), Mleck first interned in Latta's office as an undergraduate. He holds a BA in economics from Washington and Lee University.



Daniel Park, Policy Counsel, Office of Sen. Josh Hawley (R-MO)
Daniel Park is policy counsel to Senator Josh Hawley of Missouri and covers tech-related issues, including AI. He comes to AI policy from a background in law and finance. Park is interested in helping shape the future of AI governance and is especially interested in AI alignment and control.



Ben Schwartz, Legislative Assistant, Office of Rep. Josh Gottheimer (D-NJ)
Ben Schwartz is a legislative assistant for Representative Gottheimer with a focus on AI and technology more broadly. Previously, he served as an advisor for Families Over Billionaires, the Democrats’ Campaign HQ for the 2025 Tax Fight. Prior to that, he served as a special assistant in the Biden administration’s Office of Management and Budget. Schwartz is a graduate of the University of Rochester and a native of Bethesda, Maryland.



Michael Stadelmaier, Legislative Assistant, Office of Rep. Claudia Tenney (R-NY)
Michael is a legislative assistant and district liaison for Rep. Claudia Tenney. Previous to this role, he served as district director. His current policy portfolio is the Science, Space, & Technology Committee work, energy, infrastructure, telecom and broadband. He has a BA and an MPA from UTSA and previously worked for the U.S. Senate and New York State Assembly.



Dan Taylor, Legislative Director and Senior Counsel, Sen. Maria Cantwell (D-WA)
Dan serves as Legislative Director and Senior Counsel for Sen. Maria Cantwell (D-WA). He has served in personal and committee offices of senior members of the Senate and House of Representatives in addition to working in the executive and judicial branches. A native Washingtonian, Dan is a graduate of Claremont McKenna College and George Washington University Law School.



Elya Taichman, Legislative Director, Office of Rep. Eugene Vindman (D-VA)
Elya “Eli” Taichman is legislative director for Representative Eugene Vindman (D-VA), a position he assumed in January 2025. He previously served as a senior adviser in the Bureau of Legislative Affairs at the State Department. Before that, he was legislative director for Congresswoman Lori Trahan (D-MA) and worked in the office of then-Congresswoman Michelle Lujan Grisham (D-NM). Taichman holds a law degree from Temple University Beasley School of Law and an undergraduate degree from the University of Pennsylvania. He lives in Washington, D.C., with his wife and son. He is originally from Lower Merion, Pennsylvania. His hobbies include piano, history, and cooking.



Marcus Towns, Director of Member Services, House Appropriations Committee
Marcus Towns serves as the director of member services for the House Appropriations Committee.



Clint Trocchio, Chief Clerk, Senate Appropriations Committee
Clint Trocchio is chief clerk of the U.S. Senate Committee on Appropriations, where he helps bridge policy, technology, and evidence-based decision-making. A retired Coast Guard aviator, flight instructor, budget examiner, and congressional liaison, he has spent more than 25 years translating missions into resources and resources into measurable outcomes. On Capitol Hill, Trocchio built the committee’s first analytics and digital innovation function, helping modernize how Congress manages appropriations, earmarks, and oversight. His work focuses on helping public servants become public solvers through better data, better tools, and practical applications of AI. Trocchio holds degrees from the U.S. Coast Guard Academy, Spring Hill College, and George Washington University, and recently completed a Hoover Institution Veteran Fellowship at Stanford University, where he developed the committee’s Applied Policy Lab framework.



Zach Vega, Professional Staff Member, Senate Republican Policy Committee

Zach Vega is a professional staff member for the Senate Republican Policy Committee covering artificial intelligence, technology, science, and telecommunications issues. Much of his work focuses on the intersection of technology, policy, and its impact on other sectors, and he publishes weekly updates for fellow Senate staff. His recent focus has been on the proliferation of AI-related infrastructure in federal scientific research. Before joining the committee, Vega worked for Senator Pat Toomey. He is from Silver Spring, Maryland.



Paul Zepp, Senior Counsel, House Select Committee on the CCP

Paul Zepp is senior counsel to the House Select Committee on the Strategic Competition between the United States and the Chinese Communist Party, where he focuses on emerging technology policy. His work examines the national security, governance, and strategic implications of advanced technologies, with a particular emphasis on the convergence of AI and QIST. He brings experience in military service, law, and public sector policy and is especially interested in strengthening trusted cooperation among partners working on critical and emerging technologies. Zepp received his JD from Yale Law School and has completed professional studies in cybersecurity, quantum computing, and artificial intelligence.

Special thank-you to the Stanford HAI staff who contributed to the boot camp:



Annie Benisch, Event Planner, Stanford HAI

Annie Benisch is an event planner and designer who is passionate about crafting memorable experiences that last long after the event is over. Whether virtual, in-person, a conference, a trade show, or an employee event, Benisch consistently creates meaningful and impactful connections that enhance HAI’s bottom line. With a keen eye for detail and a commitment to excellence, she brings brands to life through innovative event strategies and engaging marketing campaigns. Her unique approach to event planning and design has helped numerous companies strengthen their relationships with clients, employees, and stakeholders.



Dana Heinemann, Policy & Society Coordinator, Stanford HAI

Dana Heinemann supports initiatives at the intersection of AI technology, policy, and society by coordinating programs and events. Prior to joining Stanford, Heinemann worked at two healthcare startups, where she managed human resource operations and developed a passion for mission-driven innovation. She holds a degree in project management and brings a collaborative, detail-oriented approach to foster engagement across academia, policy, and industry.



Shana Lynch, Head of Content and Associate Director of Communications, Stanford HAI

Shana Lynch oversees HAI content and channels, ensuring HAI’s cutting-edge research and faculty thought leadership reach audiences ranging from policymakers to industry professionals to the general public. Before joining HAI, she was a senior editor at Stanford Graduate School of Business and managing editor at the *Silicon Valley Business Journal*.



Jeanina Matias, Digital Media Manager, Stanford HAI

Jeanina Matias’ role involves digital community-building through social media channels, blogs, and the web and overseeing brand assets and the voice of the content on digital channels. Prior to joining HAI, she worked with local and global nonprofits in the United States and in the Philippines in various capacities involving digital media, marketing, news writing, graphic design, and events. She majored in journalism and graduated with a BS in development communication from the University of the Philippines Los Banos. In her spare time, she volunteers as a photographer and a graphic designer for *Mahalaya*, an SF-based community newspaper, which centers the experiences and narratives of Filipinos in and beyond the Bay Area.



Madeleine Wright, Communications Specialist, Stanford HAI

Madeleine Wright manages key communications-related projects, assists with marketing efforts, and works closely with the events team. She holds an MS in marketing from Santa Clara University and a BA in economics and communications studies from Gonzaga University.



Maddie Ziff, Event Manager, Stanford HAI

Maddie Ziff leads the events team at HAI and has been managing academic events for almost 10 years, with most of that time at Stanford. Before joining HAI in January 2025, Ziff worked as a conference and events manager at Stanford Graduate School of Business, where she planned large conferences and enjoyed working directly with student groups. Prior to her time at the GSB, Ziff was an associate manager of events with Stanford’s Office of Special Events and Protocol and assisted with planning and executing large scale events such as Stanford’s annual commencement ceremony. Ziff received her BA in business and psychology from Brandeis University.

This page is intentionally left blank.

