Stanford University
Human-Centered
Artificial Intelligence

Stanford

2019–2020 ANNUAL REPORT

Stanford Institute for Human-Centered Artificial Intelligence

Find out how HAI is advancing AI research, education, policy, and practice to improve the human condition.

Contents

Letter from the Denning Co-Directors

Artificial Intelligence will change lives in profound ways. We feel a collective responsibility to help guide this technology in a positive direction. To this end, in 2019 we founded the Stanford Institute for Human-Centered Artificial Intelligence (HAI) to advance AI research, education, policy, and practice to improve the human condition. We want to thank the many faculty and advisors who joined us in conceiving and creating the institute and helping to define its charter and scope.

Our growing community of researchers, scholars, fellows, students, and partners brings true diversity of thought to the critical question of how AI will affect our society. This is by nature multidisciplinary work that could not be undertaken without collaboration across many different academic departments, educational institutions, and civil society, government, and industry organizations.

As 2020 unfolded the world faced unprecedented challenges and disruption. At HAI we scarcely missed a beat, embracing each new challenge with flexibility and resolve. Just weeks into the first shelter-in-place order, HAI hosted a virtual conference on COVID-19 and AI which was viewed by 85,000 people. In June, we hosted a second virtual conference on COVID-19 + AI: The Road Ahead. In many ways the online format proved liberating, allowing more rapid planning and execution of a wide range of workshops and events.

During the 2019–20 year, we continued our seed grant program and launched our first round of major, Hoffman-Yee grants. Meanwhile, the HAI policy team championed preliminary legislation for a National Research Cloud, taking it from concept through to its passing, on January 1, 2021, as part of the National Defense Authorization Act. To round out the institute's first year of operations, we introduced a range of multidisciplinary education programs focused on helping high-impact decision makers as well as future leaders respond thoughtfully to the social and ethical implications of AI.

Fei-Fei Li and John Etchemendy, Denning Co-Directors of the Stanford Institute for Human-Centered AI

We are sincerely grateful for the support of faculty from all seven schools at Stanford as well as our extended community, our generous donors, advisors, and corporate members, and the dedicated HAI staff. As we look to the future, HAI will continue to build on Stanford's strong culture of innovative research and long tradition of leadership in AI. Together with our ability to convene stakeholders from all sectors, we will continue to work toward a shared vision of AI that truly improves the human condition.

Exploring the full scope, scale, and impact of AI feels more urgent than ever in AY21. We look forward to continuing to work with you this year.

Fei-Fei Li

John Etchemendy

HAI by the Numbers

We measure our success by impact, not numbers. But we'd like to share some numbers we're proud of.

96

Stanford departments represented by HAI faculty

240+

Stanford faculty funded by HAI

112

Faculty, students, and fellows whose research is featured on the HAI blog

54

Events

$10.5M

Issued in grants from the HAI launch to the end of 2020

15,334

Attendees across all HAI events

374,000+

Views of HAI event videos on YouTube

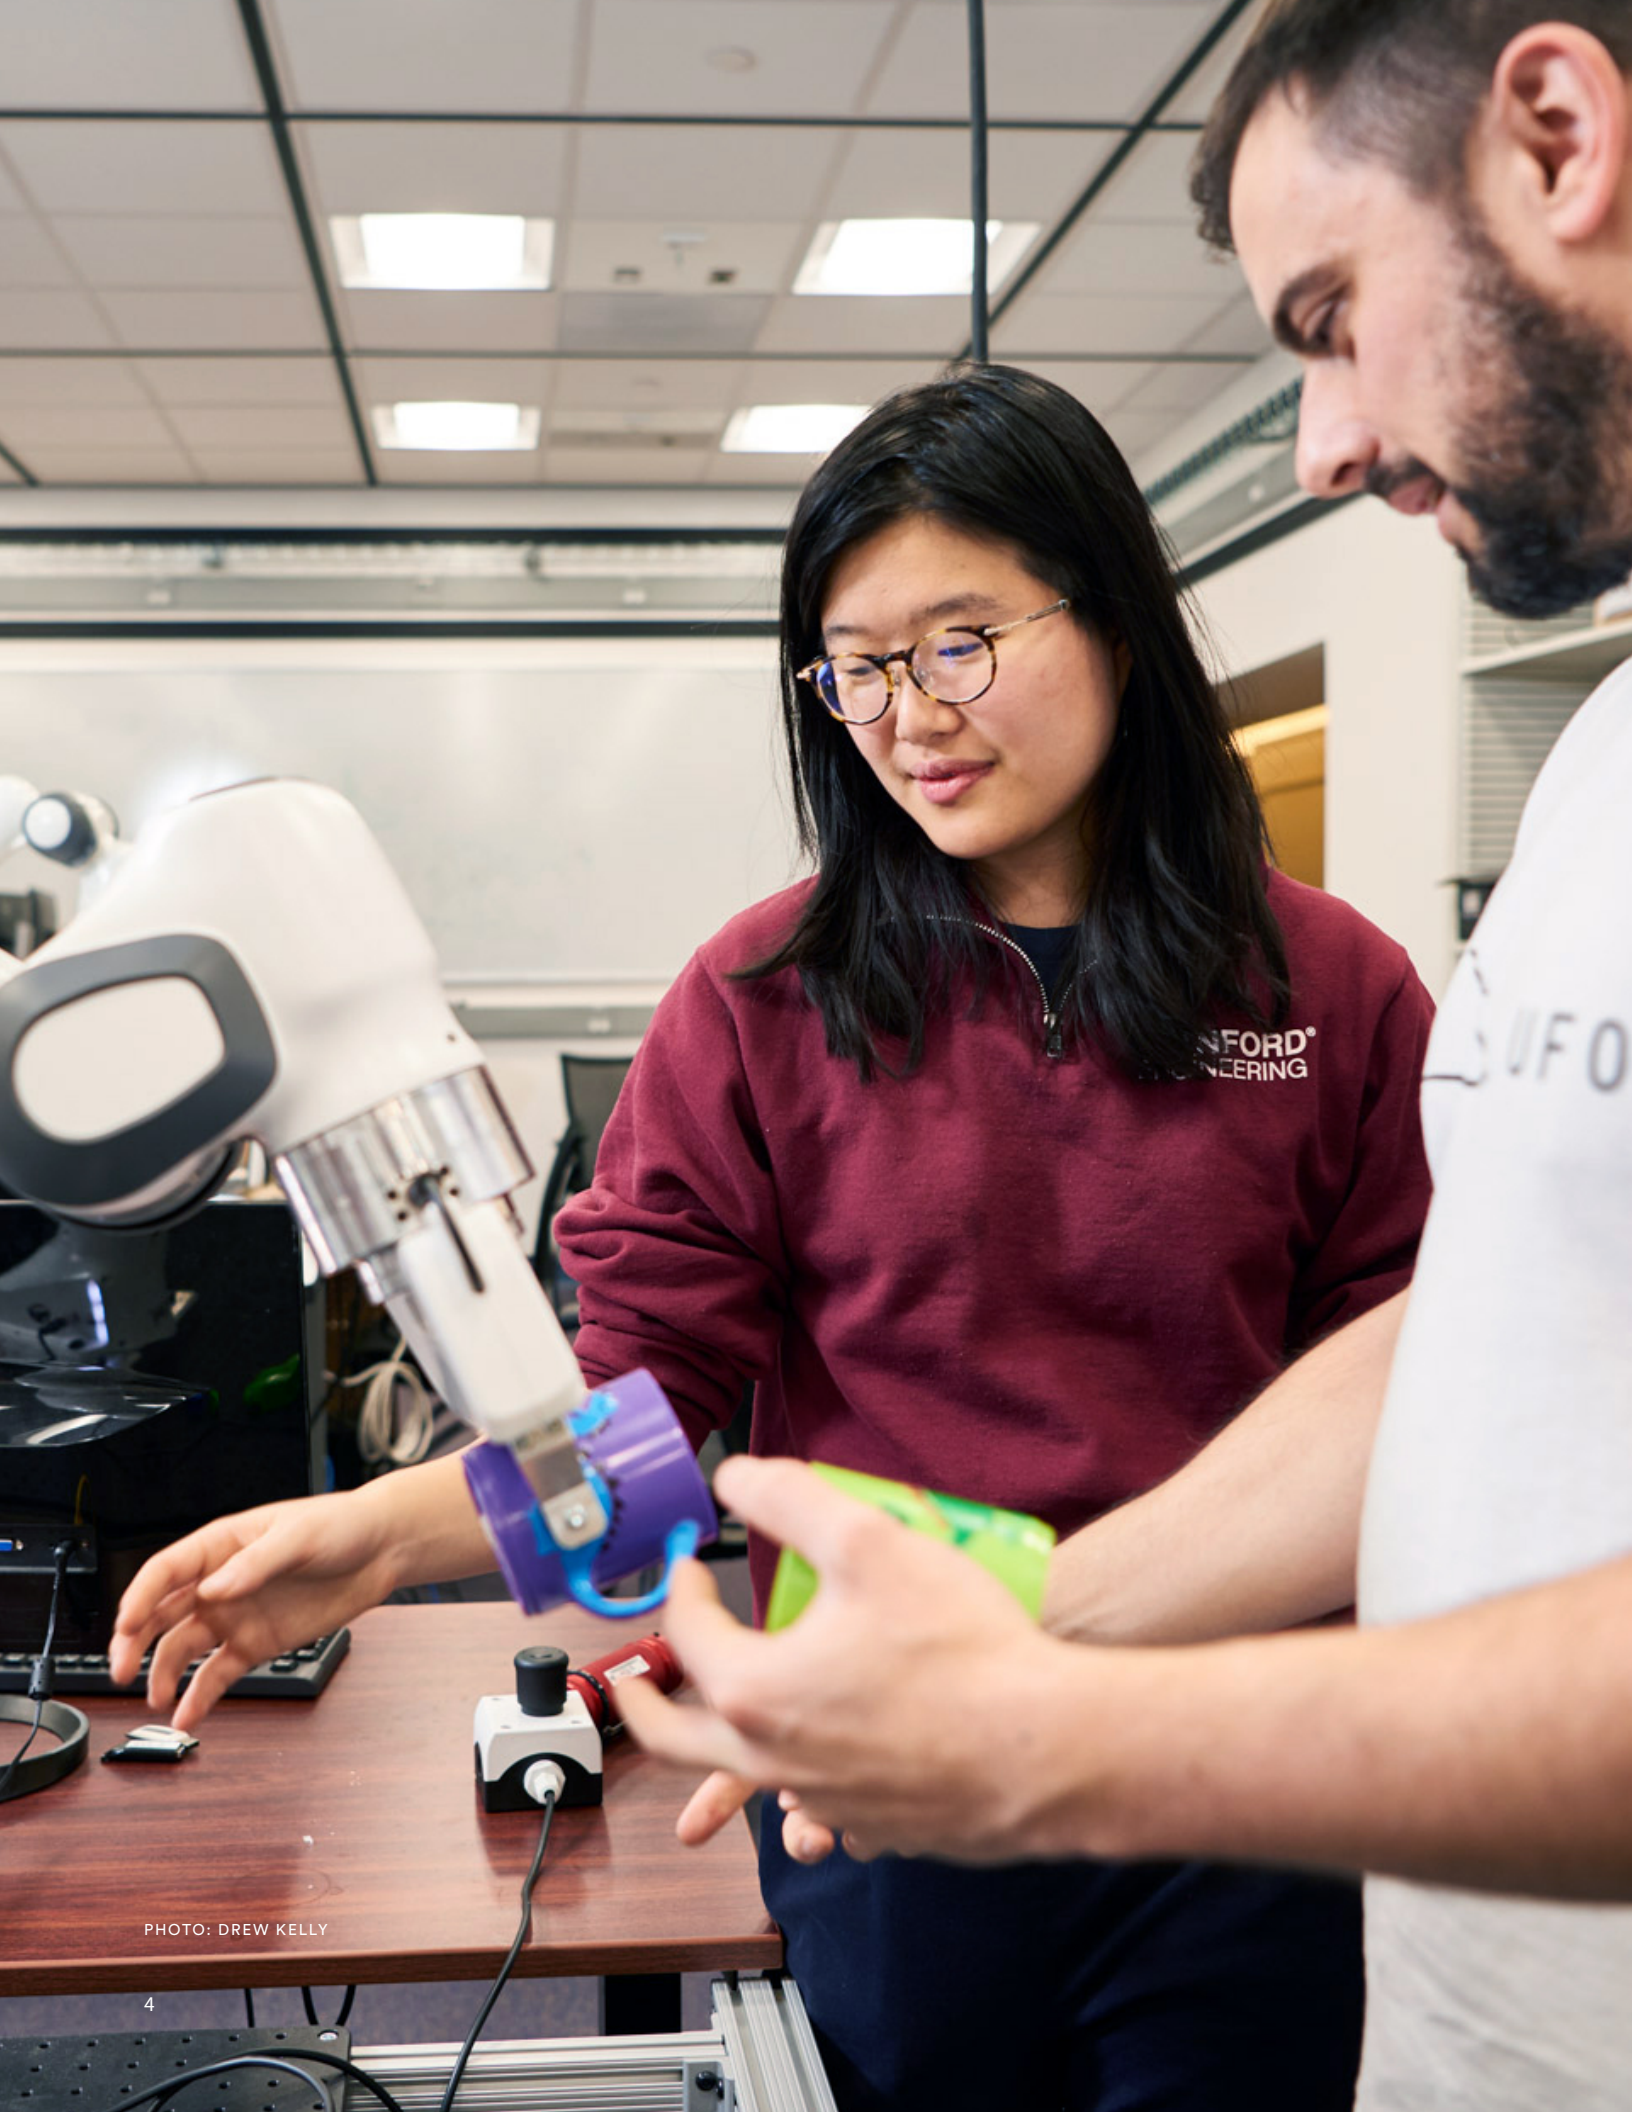

PHOTO: DREW KELLY

Research: Guiding Human-Centered AI

Our vision for the future is led by our commitment to studying, guiding, and developing human-centered AI technologies and applications. We believe AI should be collaborative, augmentative, and enhancing to human productivity and quality of life.

HAI Research Focus

During the 2019–20 academic year, HAI faculty, fellows, and partners engaged in dozens of collaborative research projects spanning many aspects of technology and society. Here are a few highlights, organized around HAI's three main research themes:

Human Impact

Our research in this area aims to understand and guide the global societal impact of AI technologies for the greater good. This includes the impact of AI on economics, society, government, ethics, philosophy, policy, social sciences, and humanities. Consideration for human impact also may include race, ethnicity, and gender studies; interpretable, trustworthy, and fair AI; and the intellectual and conceptual foundations of AI, its history, and its cultural impact.

Augment Human Capabilities

In this research theme we focus on the design and creation of AI technologies that augment humans rather than replace them. For example, AI and human-computer interaction; health, medicine, and wellness; robotics and automation; and education and law.

Intelligence

Our research mission aims to develop novel technologies inspired by the depth and versatility of human intelligence. This includes AI inspired by neuroscience, cognitive science, and psychology; novel unsupervised, semi-supervised, self-supervised, and supervised methods for diverse data types; and knowledge and semantics.

Featured Projects

HUMAN IMPACT

The Carbon Footprint of AI

AI systems use considerable amounts of power and can generate alarmingly high volumes of carbon emissions. Accurate reporting is essential for understanding the potential climate impact of research projects. A team from Stanford, Facebook AI Research, and McGill University devised a framework for tracking real-time energy and carbon usage, along with a leaderboard to motivate engineers to think about the environmental impact of their work. Stanford computer science PhD student Peter Henderson served as lead author of the study, working closely with Stanford professors Dan Jurafsky and Emma Brunskill. By making it easy to measure the carbon footprint of AI, this group hopes to drive more commitments to energy-efficient algorithms.

AUGMENT HUMAN CAPABILITIES

Smarter Hospitals

Advances in machine learning and contactless sensors have given rise to ambient intelligence—physical spaces that are sensitive and responsive to the presence of humans. Fei-Fei Li, HAI Co-director, Arnold Milstein, Professor of Medicine and Director of Stanford's Clinical Excellence Research Center; and Albert Haque, Department of Computer Science graduate student, co-authored a study showing how this technology could improve patient outcomes in healthcare settings. They concluded that thoughtful use of emerging data science and AI technology would enable a better understanding of the complex interplay between physical environments and health-critical human behaviors. Their findings were published in the September 2020 issue of *Nature*.

INTELLIGENCE

Curiosity-Driven Algorithms

Computational neuroscientist and Assistant Professor of Psychology and Computer Science Daniel Yamins and Professor of Human Biology Michael Frank have been collaborating on developing curiosity-driven systems that mimic the way children learn, with the ultimate goal of making more effective diagnoses for a variety of developmental disorders. In July 2020, Yamins published "Active World Model Learning with Progress Curiosity," which provides a foundation for their research into children's curiosity.

Grant Programs for AI Research

HAI strives to foster a culture of AI research in which technological advancements are inextricably linked to research about their potential societal impacts. The three HAI grant programs support novel approaches and interdisciplinary research collaborations.

244 funded faculty team members

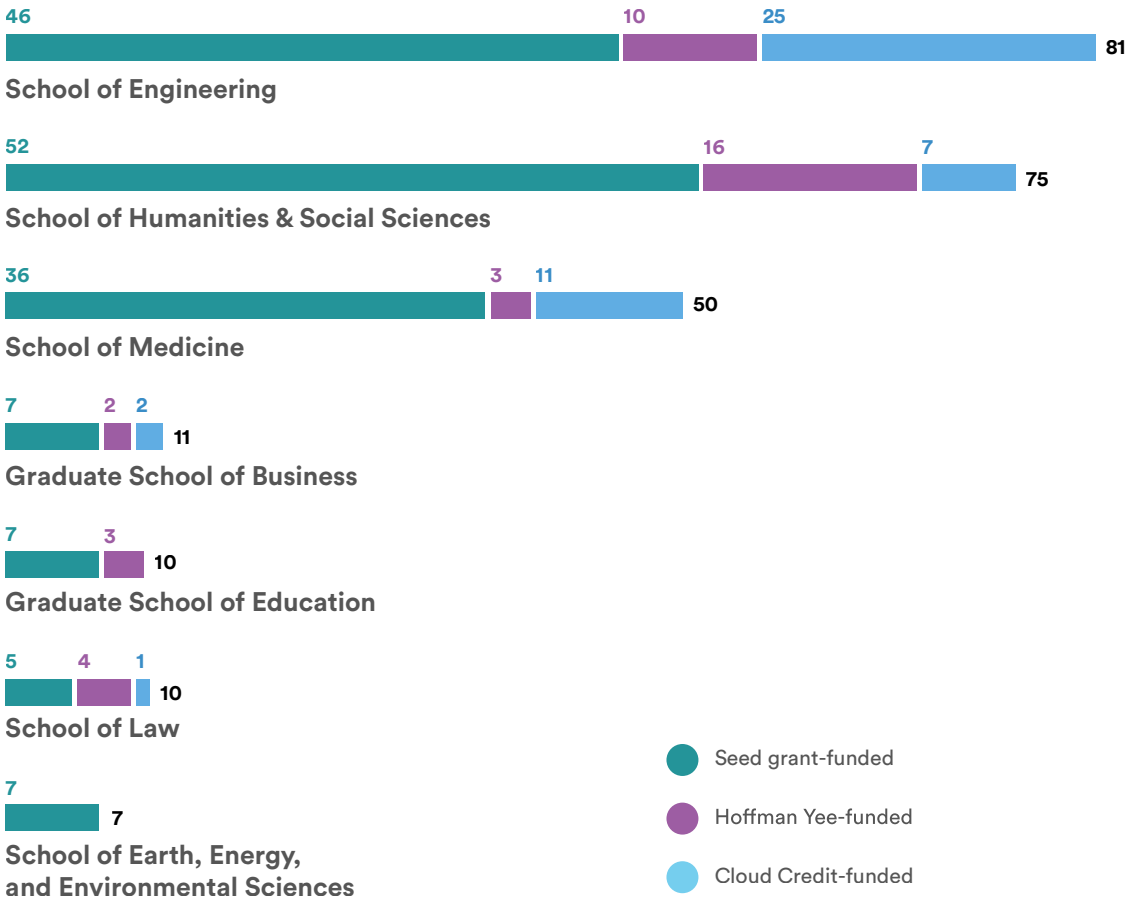

Projects spanning multiple departments

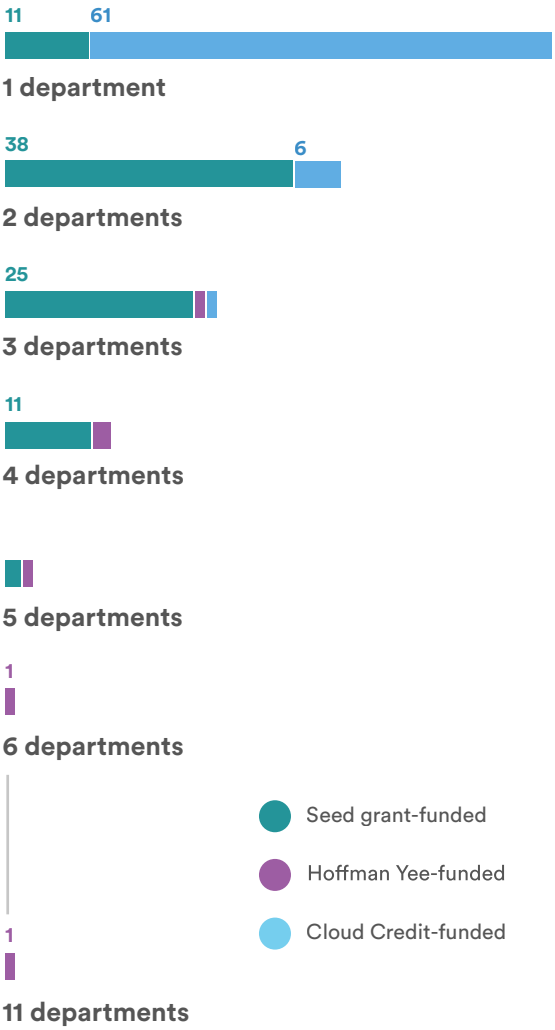

HAI Seed Grants

HAI Seed Grants fund speculative, early-stage research proposals with an objective of getting initial results that are likely to secure further support from internal and external funders. Since 2018, this program has allocated nearly $5.5 million. Topics ranged from AI for Scientific Discovery to Virtual Multisensory Interaction to AI in the Regulatory State.

> “We’re bringing together unprecedented access to data and building algorithms that will allow us to find and start to intervene for victims of human trafficking.”

HAI SEED GRANT RECIPIENT VICTORIA WARD, CLINICAL ASSISTANT PROFESSOR OF PEDIATRICS

160 faculty members

87 projects

Hoffman-Yee Research Program

Hoffman-Yee Research Grants fund ambitious projects that address significant scientific, technical, or societal challenges requiring interdisciplinary collaboration. In its first year, the program awarded $500,000 to six interdisciplinary teams. Successful projects were selected for their boldness, ingenuity, and potential for transformative impact. Teams will be eligible to compete for additional funding. This program was made possible by a gift from philanthropists Reid Hoffman and Michelle Yee.

> “It’s a shame when someone wants to put in the work to learn and make a contribution but no teacher can show up. Our system works toward addressing that gap.”

HOFFMAN-YEE GRANT RECIPIENT CHRIS PIECH, ASSISTANT PROFESSOR OF COMPUTER SCIENCE

35 faculty members

20 projects

Cloud Credit Grants

Cloud Credit Grants from Google and AWS provide the cloud compute required for rigorous AI research. Credits are used to advance promising, novel, or emerging research that requires advanced computational resources provided by the commercial cloud. In the 2019–20 academic year, AWS awarded $2 million in credits to faculty representing 15 departments across six of the seven university schools. A similar program launched in the fall of 2020 for Google cloud credits.

> “The HAI cloud credits gave us the ability to run a variety of projects in parallel, from terrain/object detection to reinforcement learning of locomotion policy to trajectory optimization of musculoskeletal systems.”

CLOUD CREDIT RECIPIENT KAREN LIU, ASSOCIATE PROFESSOR OF COMPUTER SCIENCE

45 faculty members

70 projects

Stanford Digital Economy Lab

The Stanford Digital Economy Lab brings together leading researchers and experts to examine how digital technologies are transforming work, organizations, and the economy.

Professor Erik Brynjolfsson joined Stanford in July 2020 to create the new Digital Economy Lab. Brynjolfsson is a widely cited thought-leader on the effects of information technologies on the economy and business. His new lab brings together researchers and experts to examine the economic implications of AI and other digital technologies—including employment, wages, business organization, productivity, and income inequality.

The lab's research generates insights that can help companies, policymakers, researchers, and workers rise to the challenges created by an era of profound digitization. Current research themes include AI and the future of work, measuring the digital economy, and the economic impacts of COVID-19. Christie Ko joined the Lab as Executive Director in July and manages research funding, seminars, lab meetings with postdocs and graduate students, and interactions with executives, policymakers and sponsors.

In June, Professor Brynjolfsson presented at the HAI COVID + AI the Road Ahead conference on "How COVID has Accelerated Automation in the Workplace." During his 25 years with the Massachusetts Institute of Technology, he published nine books, including *The Second Machine Age*, and wrote more than 100 academic articles.

Erik Brynjolfsson

Jerry Yang and Akiko Yamazaki Professor and Senior Fellow at HAI, Director of the Stanford Digital Economy Lab, Ralph Landau Senior Fellow at the Stanford Institute for Economic Policy Research (SIEPR), and professor, by courtesy, at the Stanford Graduate School of Business and Stanford Department of Economics

STANFORD DIGITAL ECONOMY LAB

Ethics Review Board

The 2020 Hoffman-Yee Research Grant and HAI seed grant processes included proposal reviews by an ethics review board to evaluate the potential benefits and harms to society of the proposed research. All grant proposals included a statement of how the projects would mitigate ethical issues and negative impacts on society. Four professors led this effort: Associate Professor of Computer Science Michael Bernstein, Professor of Political Science and Director of CASBS Margaret Levi, Professor of Philosophy and Dean of Humanities & Sciences Debra Satz, and Professor of Medicine and Biomedical Ethics David Magnus.

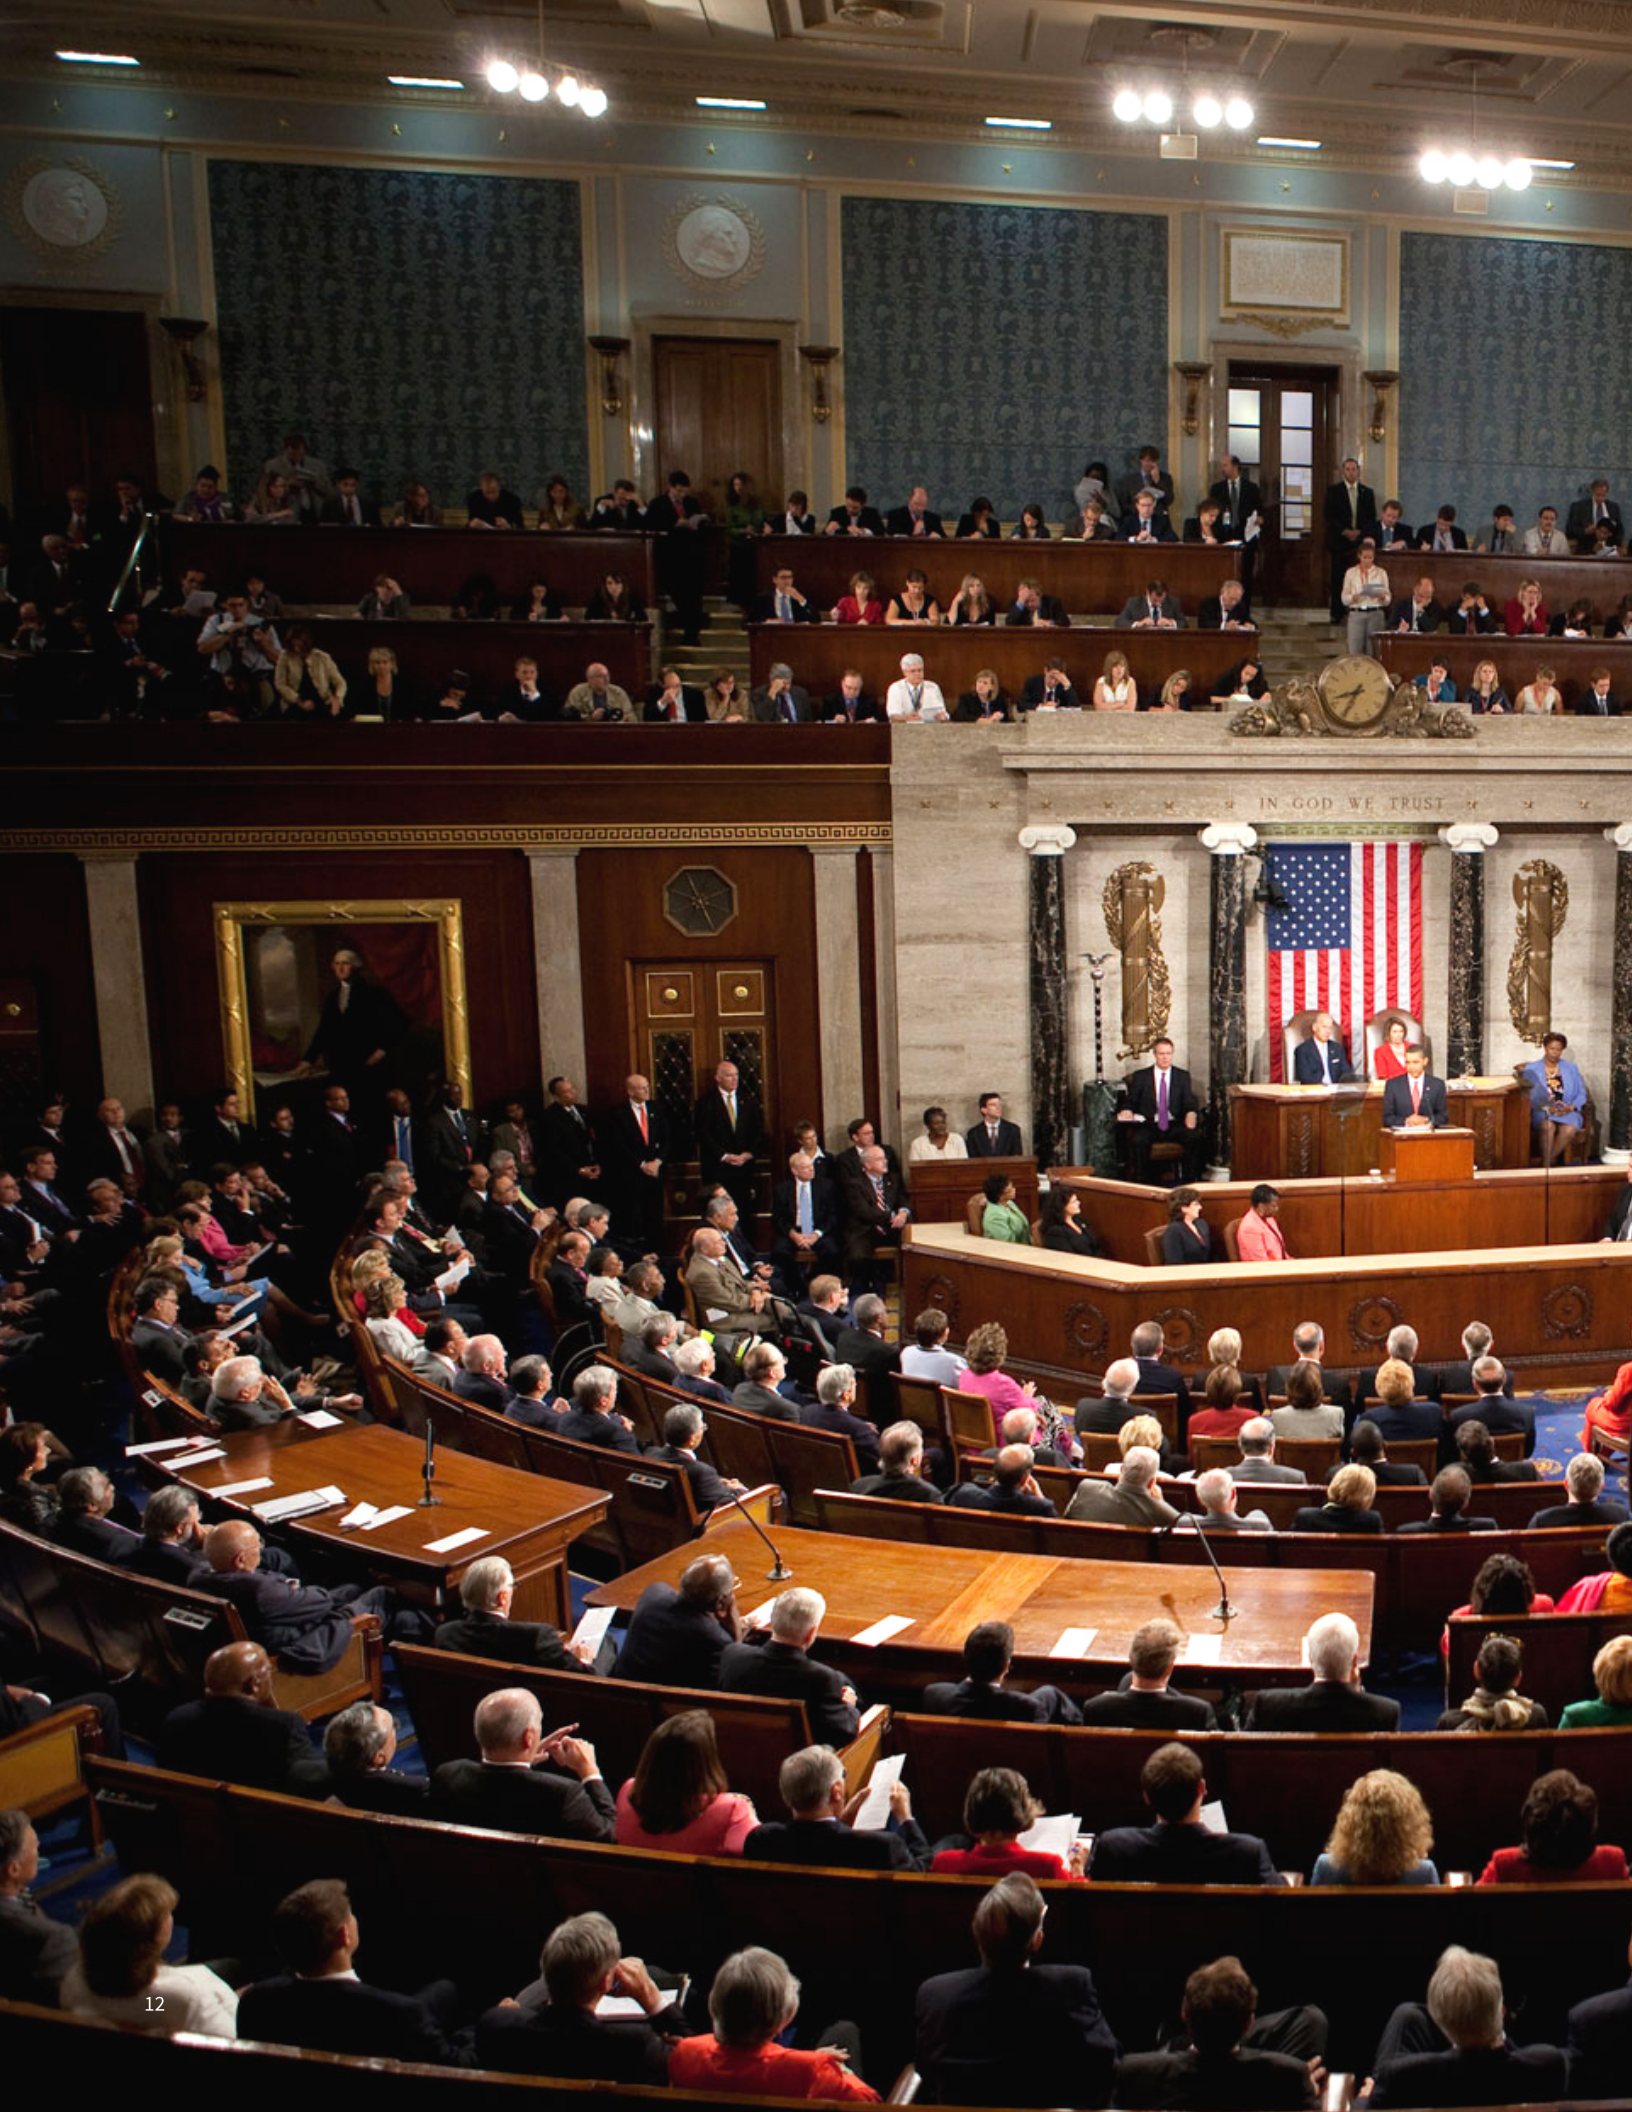
IN GOD WE TRUST

A Hub for Policy Impact

Since our founding, HAI has played a pivotal role in convening national and global discussions related to artificial intelligence policy. We have engaged directly with local, state, federal, and international government officials and civil society organizations, produced cutting-edge, policy-relevant multidisciplinary research, and provided policymakers with tools for effective decision making. As a result of these efforts, HAI has become a natural hub for AI policy.

National Research Cloud

Artificial intelligence, one of the most consequential technologies of the 21st century, is a direct outgrowth of federally funded university research, further advanced by exceptional R&D in the private sector. But the research prowess that's powered decades of growth and prosperity is at risk.

Today's AI requires vast amounts of compute, data, and expertise to train and deploy the massive machine learning models powering the most advanced research. And access to sufficient compute power and the high-quality datasets needed for research are increasingly out of reach for most colleges and universities. The United States has been and remains a leader in AI research. But declining government investment in basic and foundational research, combined with lack of access to computational resources and large datasets, threatens America's position on the global stage.

Over the past year, HAI's co-directors and policy team played a leading role in crafting legislation to make world-class computational resources and robust government-held datasets available for researchers across the country. This effort involved securing support from 22 leading computer science colleges and universities and the National Security Commission on Artificial Intelligence. In addition, Senator Rob Portman (R-OH) and Congresswoman Anna Eshoo (D-CA 18th District) both supported the initiative by sponsoring legislation.

The effort to shape a bicameral, bipartisan bill was more than a year in the making before reaching a major milestone when it passed as part of the National Defense Authorization Act on January 1, 2021. The new law provides for a federal task force to develop and propose recommendations for a National Research Cloud. Although it is only the beginning of what will be a lengthy process, this initiative represents a critical first step toward democratizing AI technology.

Featured Policy Activities

June 2019—HAI Co-Directors John Etchemendy and Fei-Fei Li in Washington D.C.

Visit to Washington, D.C.

In June 2019, HAI Co-directors John Etchemendy and Fei-Fei Li met with senior policymakers in Washington, D.C. to discuss the growing importance of AI and why government leaders need to be prepared to play a role in shaping its future. The agenda was planned with the goal of building key relationships and included briefings with the National Institutes of Health, National Institute of Standards and Technology, House AI Caucus, Executive Office of the President, the Departments of Defense and Veterans Affairs, DARPA, and a key meeting with Senator Rob Portman (R-OH), Chair of the Senate AI Caucus, whose participation would prove invaluable to garnering support for the National Research Cloud bill.

Government by Algorithm

AI use is widespread across the U.S. federal government; however, few if any agencies are applying the technology in ways that rival the sophistication of the private sector. This uneven adoption curve has raised questions around transparency, fairness, ethics, and efficiency. In 2019, the Administrative Conference of the United States commissioned a team of HAI-affiliated lawyers and computer scientists to examine AI uses among federal agencies. Published in February 2020, the resulting report, Government by Algorithm: Artificial Intelligence in Federal Administrative Agencies, is the most comprehensive study of the subject ever conducted in the United States. It revealed deep concerns about growing government use of AI tools and suggested how AI could be applied to make the federal government work better, more fairly, and at lower cost.

California Future of Work Commission

In August 2019, California Governor Gavin Newsom announced the Future of Work Commission to study, understand, analyze, and make recommendations regarding the impact of technology on work and the best way to ensure prosperity for all. HAI Co-director Fei-Fei Li and Associate Director Susan Athey joined other leaders representing technology, labor, business, education, and venture capital in accepting the invitation to serve as members. Two HAI Advisory Council members co-chair the commission: James Manyika, Chairman and Director of the McKinsey Global Institute, and Mary Kay Henry, International President of the Service Employees International Union (SEIU).

Policy Workshops

HAI regularly organizes invitation-only workshops that convene diverse stakeholders on urgent and timely topics. As an example—Facial Recognition Technology raises profound questions about the role of technology in society. In May 2020, HAI convened a half-day workshop to discuss the question of operational accuracy as the technology is considered for use in new domains. Leading computer scientists, legal scholars, and representatives from industry, government, and civil society participated, and the results of their discussions were published as a white paper, "Evaluating Facial Recognition Technology: A Protocol for Performance Assessment in New Domains."

As another example, in July 2020, HAI convened a workshop for EU and U.S. regulators on "Assessing and Managing Risks with the Use of Artificial Intelligence." This was the first in a series of workshops with outside stakeholders to build a research agenda in the AI regulation space. The resulting agenda is expected to be completed in 2021.

Education for the Leaders of Today and Tomorrow

Working with faculty across the university, the HAI education team creates multidisciplinary programs for high-impact decision makers and emerging leaders. Our portfolio of courses helps people pursue cutting-edge research, create useful and responsible AI applications, and respond thoughtfully to the societal and ethical implications of global AI.

PHOTO: DREW KELLY

FEATURED PROGRAMS FOR

High-Impact Decision Makers

Today's leaders need to understand the current state of AI technology and how it is likely to affect society as rapid advancements continue.

Cyber and Artificial Intelligence Boot Camp

In August 2019, HAI partnered with the Hoover Institution and the Freeman Spogli Institute to host AI training for policymakers. HAI researchers presented congressional staff with a conceptual framework to help them anticipate and manage converging technology and policy issues.

Expert Roundtable

In partnership with the Stanford Program on Geopolitics, Technology and Governance, HAI invited 14 major media outlets to a roundtable discussion in October 2020, during which Stanford experts addressed technology questions facing the next administration. Journalists from *The New York Times*, Reuters, *The Economist, The Financial Times*, and *Wired* joined the conversation.

Executive Education

In 2019 and 2020, HAI co-sponsored two executive education courses with the Stanford Graduate School of Business: "Harnessing AI and Big Data: Analysis to Action" and "Harnessing AI for Breakthrough Innovation and Strategic Impact." Designed in 2019 by Associate Professor of Aeronautics and Astronautics Mykel Kochenderfer and Professor of Economics Paul Oyer, the new courses drew participants from 20 industry sectors.

Professional Development

During the 2019–20 academic year, HAI issued 1,200 professional credentials to participants of its AI education programs.

FEATURED PROGRAMS FOR

Emerging and Future Leaders

Tomorrow's leaders will face unprecedented challenges as AI becomes pervasive in society. These programs are designed to help future decision makers embrace the opportunity to shape how technology influences our society.

Embedded EthiCS Program

HAI partnered with the McCoy Family Center for Ethics in Society to support the Embedded EthiCS program in the Stanford Computer Science Department. This new initiative seeks to integrate ethical thinking into all aspects of computing.

Public Interest Technology University Network

HAI joined other departments across the Stanford community in supporting the Public Interest Technology initiative, a group that's committed to developing a new generation of civic-minded technologists and bridging the gap between technology and public service.

AI4ALL

This U.S.-based nonprofit works to increase inclusion in AI education, research, development, and policy. In the summer of 2020, HAI teamed up with the Stanford AI Lab, AI4ALL, and Stanford Pre-Collegiate Studies to offer the first online program affiliated with AI4ALL at Stanford. High school students from groups that have been historically under-represented in AI fields completed the three-week class with Stanford AI researchers.

HAI Undergraduate Concentration

In partnership with the Stanford Symbolic Systems program, HAI introduced the Human-Centered Artificial Intelligence undergraduate concentration.

The People of Stanford HAI

Our community thrives on the energy and commitment of the many individuals who share a desire to responsibly guide the future of AI. From the leadership team and advisory council to faculty and fellows, who are supported by a dedicated group of staff members, all are essential to ensuring that we make continual progress toward our mission.

Denning Co-Directors

John Etchemendy

Provost Emeritus and Patrick Suppes Family Professor in the School of Humanities and Sciences, Stanford University; Denning Co-Director, Stanford Institute for Human-Centered Artificial Intelligence

In addition to co-directing the institute, John taught "Philosophy of Artificial Intelligence" in the winter quarter and played an instrumental role in advancing the National Research Cloud initiative.

Fei-Fei Li

Sequoia Professor, Computer Science Department; Denning Co-Director, Stanford Institute for Human-Centered Artificial Intelligence

Among the many roles she served for HAI in AY20, Fei-Fei published "Illuminating the dark spaces of healthcare with ambient intelligence" in *Nature.* She also served as a member of the Future of Work Commission convened by California Governor Gavin Newsom and helped formulate the concept for a National Research Cloud.

The co-directors of HAI share a vision of artificial intelligence serving the collective needs of humanity. It is this goal that guides the work of HAI.

Associate Directors

The associate directors of HAI represent the diverse fields and expertise that help enrich the multidisciplinary approach of the institute.

Russ Altman

Kenneth Fong Professor and Professor of Bioengineering, of Genetics, of Medicine (General Medical Discipline), of Biomedical Data Science, and, by courtesy, of Computer Science

Russ led HAI's efforts for the April and June virtual conferences on COVID-19 and AI. He also published "Geographic Distribution of US Cohorts Used to Train Deep Learning Algorithms" research in *JAMA. Forbes* magazine highlighted his COVID-19 research in "What A Stanford Researcher's Fight Against Covid-19 Can Tell Us About The Future Of Drug Discovery."

Susan Athey

Economics of Technology Professor, Stanford Graduate School of Business; Senior Fellow at The Stanford Institute for Economic Policy Research; Member, Institute for Computational and Mathematical Engineering; and Professor, by courtesy, of Economics

Susan published her research on "Computational social science: Obstacles and opportunities" in *Science*. She also served as a member of the California Governor's Council of Economic Advisors and provided congressional testimony to the House Budget committee on the impact of AI on economic recovery and the future of work.

Michele Elam

William Robertson Coe Professor of Humanities, Department of English, Center for Comparative Studies in Race & Ethnicity, Race & Technology Affiliate at the Center for Comparative Studies in Race & Ethnicity

Michele led HAI's diversity, equity and inclusion efforts and created the HAI Visiting Artist program, a collaboration with the Office of the Vice President of the Arts. She also moderated the HAI premiere of "Coded Bias," with Fei-Fei Li and filmmaker Shalini Kantayya. Michele helped lead the April 1 conference on COVID 19 + AI with Russ Altman and Rob Reich.

Surya Ganguli

Associate Professor of Applied Physics, and by courtesy, of Neurobiology, of Electrical Engineering, and of Computer Science

Surya published three important papers in AY20: "Fundamental bounds on the fidelity of sensory cortical coding" in *Nature*, "Discovering Precise Temporal Patterns in Large-Scale Neural Recordings through Robust and Interpretable Time Warping" in *Neuron*, and "A deep learning framework for neuroscience" in *Nature Neuroscience*.

Daniel E. Ho

William Benjamin Scott and Luna M. Scott Professor of Law, Professor of Political Science, Senior Fellow at the Stanford Institute for Economic Policy Research

Daniel published a report on "Government by Algorithm: Artificial Intelligence in Federal Administrative Agencies" as well as "Algorithmic Accountability in the Administrative State" in the *Yale Journal of Regulation*. Daniel also led HAI's Facial Recognition Technology workshop and resulting white paper, "Evaluating Facial Recognition Technology: A Protocol for Performance Assessment in New Domains."

James Landay

Anand Rajaraman and Venky Harinarayan Professor in the School of Engineering, Professor of Computer Science

James led the Hoffman-Yee and Seed grant process for HAI, from calls for proposals through reviews and decisions. He published "QuizBot: A Dialogue-based Adaptive Learning System for Factual Knowledge," in Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19). In the spring quarter, James co-taught CS 335: "Fair, Accountable, and Transparent (FAT) Deep Learning."

Christopher Manning

Thomas M. Siebel Professor in Machine Learning, Professor of Linguistics and of Computer Science; Director, Stanford Artificial Intelligence Laboratory

Christopher's article on "Emergent linguistic structure in artificial neural networks trained by self-supervision" was published in *PNAS (Proceedings of the National Academy of Sciences)*. He also mentored the team that won second place in the Alexa Prize Socialbot Grand Challenge 3.

Rob Reich

Professor of Political Science, Faculty Director of the McCoy Family Center for Ethics in Society, and the Marc and Laura Andreessen Faculty Co-Director of the Stanford Center on Philanthropy and Civil Society

Rob led HAI's efforts to create an embedded ethics initiative that will insert ethics modules into core courses in the computer science major. He published "Teaching Computer Ethics: A Deeply Multidisciplinary Approach" in Proceedings of the 51st ACM Technical Symposium on Computer Science Education (SIGCSE'20). Rob also led HAI efforts for the April and June virtual conferences on COVID-19 and AI.

Staff

The staff of HAI is the lifeblood of the institute, collaborating with faculty to design and implement programs and initiatives with global impact.

Michael Sellitto, Deputy Director

As the founding Deputy Director of HAI, Michael made fundamental contributions concerning all aspects of designing, building, and launching the institute, including those relating to mission, strategy, operations, and budget. He also spearheaded work on HAI's policy and educational projects, both developing and executing key initiatives spanning the institute and the university itself.

While Michael led much of the work of HAI staff, Amita Kumar, Director of Administration, oversaw finance, operations, and administration. In Spring 2020, HAI bolstered its staff with seasoned professionals across key areas in support of HAI's research, education, and policy mission. The additions to the team included the following individuals.

Deep Ganguli
Director of Research

Krysten Hommel
Senior Associate Director, Strategic Initiatives

Panos Madamopoulos-Moraris
Director of Partnerships

Stacy Peña
Director of Communications

John Robichaux
Director of Education

Russell Wald
Director of Policy

Faculty and Fellows

HAI aims to appoint and support promising researchers working at intersections often overlooked by traditional academic departments as well as outstanding researchers pursuing core disciplinary topics. In addition to our fellows, we have more than 175 affiliated Stanford faculty from across all seven schools and more than 90 departments. Marietje Schaake and James Zou are featured as representative profiles of our fellows and affiliated faculty.

Johannes Eichstaedt, HAI Junior Faculty Fellow

The HAI Junior Fellows program invites early-career scholars to conduct innovative AI research as Assistant Professors in a unique, supportive, and interdisciplinary environment with unparalleled opportunities for impact in research, policy, and education. Johannes Eichstaedt is HAI's first Junior Fellow appointee.

Johannes is a computational social scientist who is jointly appointed as Ram and Vijay Shriram HAI Faculty Fellow and Assistant Professor (Research) of Psychology. During the 2019–20 academic year, he formed a team to study the impact of COVID-19 on mental health, after early data suggested that depression and anxiety rates had doubled in the months after the virus outbreak. Johannes also received a four-year NIH grant to study the use of social media to detect causal patterns in the mental health of the United States. He co-authored 12 articles, three of which used social media to track COVID symptoms and adherence to public health guidelines, and was featured several times in *The New York Times* for his expertise on the pandemic's impact on mental health.

Marietje Schaake, HAI Fellow

Marietje is a Dutch politician who served as a Member of the European Parliament (2009-2019). She is the HAI International Policy Fellow and International Policy Director of the Cyber Policy Center at Stanford. In AY20, she taught two courses: "AI and the Rule of Law, a Global Perspective" and "Technology and the 2020 Election"—and contributed a chapter to *Which Side of History*, edited by Jim Steyer. Her work has appeared in numerous publications including *Foreign Affairs, MIT Tech Review*, and Brookings Tech Stream.

James Zou, HAI Faculty Affiliate

James is an Assistant Professor of Biomedical Data Science and, by courtesy, of Computer Science and Electrical Engineering at Stanford. He also serves as the faculty director of the university-wide AI for Health program. He works on making machine learning more reliable, human-compatible, and statistically rigorous. In AY20, his team published a paper in *Nature* on the first AI algorithm to diagnose heart diseases from cardiac ultrasound videos. The group also published an AI algorithm that can generate molecular details from cancer histology images in *Nature Biomedical Engineering*. Learn more about James' work in this HAI blog post: When Algorithms Compete, Who Wins?

HAI Commitment to Action

The HAI mission assumes that AI can be used to improve the human condition. This assumption holds only insofar as we are cognizant of issues of accessibility, diversity, fairness, and justice. Accordingly, in early June, we strengthened our values statement to make explicit our commitment to a more inclusive institute.

We formed a commitee to assess our current efforts and support the development of best practices. We are commited to an ongoing assessment of our attempts to promote equity at all levels and to support justice and fairness in our research.

Fellows

HAI Fellowship programs create opportunities for scholars at diverse stages of their careers to explore topics, conduct research, and collaborate across disciplines related to AI technologies, applications, or impact. In the 2019–20 academic year, HAI sponsored 15 fellows, including several fellowships offered in partnership with other Stanford centers. These joint fellowship programs allow the institute to support researchers in ways that expand our community and foster engagement of scholars working on diverse issues related to human-centered AI.

To date, HAI joint fellowships have been established with these Stanford centers:

- Center for Comparative Studies in Race & Ethnicity (CCSRE)
- Center for International Security and Cooperation (CISAC)
- McCoy Family Center for Ethics in Society (EIS)
- Graduate School of Business (GSB)
- John S. Knight Journalism Fellowships (JSK)

2019–20 HAI Fellows

LISTED IN ALPHABETICAL ORDER BY TITLE

John Bauer
AI-Engineering Fellow

Bongjun Ko
AI-Engineering Fellow

Elizabeth Adams
HAI-CCSRE Practitioner Fellow

Renata Avila
HAI-CCSRE Practitioner Fellow

Samir Doshi
HAI-CCSRE Practitioner Fellow

Hong Qu
HAI-CCSRE Practitioner Fellow

Shazeeda Ahmed
HAI-CISAC Predoctoral Fellow

Shunryu Colin Garvey
HAI-CISAC Postdoctoral Fellow

Todd Karhu
HAI-EIS Postdoctoral Fellow

Kate Vredenburgh
HAI-EIS Postdoctoral Fellow

Lisa Simon
HAI-GSB Postdoctoral Fellow

John Markoff
Journalist-In-Residence

Garance Burke
HAI-JSK Fellow

Pamela Chen
HAI-JSK Fellow

Marietje Schaake
HAI International Policy Fellow; International Policy Director, Cyber Policy Center, Stanford University

HAI Distinguished Fellows

LISTED IN ALPHABETICAL ORDER BY NAME

Yoshua Bengio
Scientific Director of Mila, Quebec's Artificial Intelligence Institute; Professor, University of Montreal

Rodney Brooks
Panasonic Professor of Robotics (Emeritus), MIT

Jeff Dean
Google Senior Fellow, Senior Vice President of Google AI

Daniel Dennett
University Professor and Director, Center for Cognitive Studies, Tufts University

Susan Dumais
Technical Fellow & Managing Director, Microsoft Research

Edward Feigenbaun
Kumagai Professor in the School of Engineering, Emeritus, Stanford University

Barbara Grosz
Higgins Professor of Natural Sciences, Harvard University

Demis Hassabis
Co-Founder, DeepMind

Geoffrey Hinton
Professor, University of Toronto

Eric Horvitz
Chief Scientific Officer, Microsoft

James Manyika
Chairman and Director, McKinsey Global Institute; Senior Partner, McKinsey & Company

John Markoff
Journalist, HAI Affiliate

Helen Nissenbaum
Professor, Information Science, Director, Digital Life Initiative, Cornell Tech

Judea Pearl
Professor of Computer Science and Director of the Cognitive Systems Lab, UCLA

Stuart Russell
Professor of Computer Science, UC Berkeley

Mustafa Suleyman
Vice President, AI Policy at Google, Co-Founder, DeepMind

Hal Varian
Chief Economist, Google; Emeritus Professor, UC Berkeley

Terry Winograd
Professor of Computer Science, Emeritus, Stanford University

Advisory Council

The volunteer Advisory Council was established to provide HAI leadership with well-informed expert advice on external trends and developments related to the institute's mission and programs. The Advisory Council is chaired by Reid Hoffman, with vice chairs Steve Denning and Srinija Srinivasan.

Jim Breyer
Kenneth I. Chenault
Mariano-Florentino Cuéllar
Steve Denning
John Hennessy
Mary Kay Henry
Mellody Hobson
Reid Hoffman
Eric Horvitz
Bob King
Susan Liautaud
James Manyika
Marissa Mayer
Samuel J. Palmisano
Condoleezza Rice
Heidi Roizen
Eric Schmidt
Mike Schroepfer
Stephen Schwarzman
Kevin Scott
Neil Shen
Ram Shriram
Vishal Sikka
Srinija Srinivasan
Jerry Yang
Songyee Yoon

VPA-HAI Visiting Artist Program

In partnership with the Office of the Vice President for the Arts (VPA) and under the leadership of Associate Director Michele Elam, HAI established a visiting artist residency to augment human creative expression and experience. This national program focuses on artists working at the intersections of art and artificial intelligence, with particular emphasis on HAI's three research focus areas: human impact, human capability, and inspired intelligence.

The HAI program advances a key goal of Stanford's developing Arts Incubator: to cultivate creative research and demonstrate how artists can work with scholars, scientists, engineers, and entrepreneurs to address important societal challenges. HAI envisions its visiting artist program evolving into a university-community, public-private partnership in collaboration with other institutes for the arts.

Stephanie Dinkins: 2019 HAI Artist in Residence

HAI Visiting Artist Stephanie Dinkins is a transmedia artist who creates platforms for dialogue about AI technology as it intersects race, gender, aging, and our future histories. She is particularly driven to work with communities of color to co-create more inclusive, fair, and ethical AI ecosystems. Her work employs lens-based practices, emerging technologies and community engagement to confront questions of bias in AI, consciousness, data sovereignty, and social equity.

Dinkins became the first HAI artist-in-residence in October 2019 as part of a pilot collaboration with the Sundance Institute's New Frontier Labs. While at Stanford, she further developed an experimental, voice-interactive AI memoir and sculpture titled "Not the Only One," which is based on oral histories gathered from her family. It's a project that sits at the intersection of art, artificial intelligence, and social justice practices.

Research Centers and Partners

The work of HAI is only made possible through collaboration with labs, centers, and institutes across and beyond the Stanford campus, all of which are critical partners in achieving our mission.

HAI Centers and Labs

AI100

The One Hundred Year Study on Artificial Intelligence, or AI100, is a 100-year effort to study and anticipate how the effects of artificial intelligence will ripple through every aspect of how people work, live, and play.

AI Index

The AI Index is an effort to track, collate, distill and visualize data relating to artificial intelligence. It aspires to be a comprehensive resource of data and analysis for policymakers, researchers, executives, journalists, and the general public to develop intuitions about the complex field of AI.

Center on AI Safety

The mission of the Stanford Center for AI Safety is to develop rigorous techniques for building safe and trustworthy AI systems and establishing confidence in their behavior and robustness, thereby facilitating their successful adoption in society.

Center for the Study of Language and Information

The Center for the Study of Language and Information (CSLI) serves Stanford faculty and students who are engaged in research involving computational, logical, and stochastic modeling of cognitive functions and processes.

Data Analytics for What's Next

Despite incredible recent advances in machine learning, building machine learning applications remains prohibitively time-consuming and expensive for all but the best-trained, best-funded engineering organizations. Data Analytics for What's Next, or DAWN, is a five-year research project to democratize AI by making it dramatically easier to build AI-powered applications.

The Digital Economy Lab

The Stanford Digital Economy Lab (S-DEL) at HAI is an interdisciplinary research group studying how digital technologies are transforming work, organizations, and the economy. An engine for research and education, S-DEL brings together an unprecedented group of stakeholders to analyze data, run experiments, develop theories, and provide actionable insights.

Open Virtual Assistant Lab

Open Virtual Assistant Lab, or OVAL, is creating an ecosystem founded on open virtual assistant technology that democratizes AI for linguistic user interfaces, creates an open and non-proprietary web, and promotes sharing with individual data ownership.

Regulation, Evaluation, and Governance Lab

The Regulation, Evaluation, and Governance Lab (RegLab) partners with government agencies to design and evaluate programs, policies, and technologies that modernize governance.

Stanford Artificial Intelligence Laboratory

The Stanford Artificial Intelligence Laboratory (SAIL) has been a center of excellence for AI research, teaching, theory, and practice since its founding in 1962.

Stanford Partners

Center for AI in Medical Imaging

AIMI was established in 2018 with the primary mission to solve clinically important problems in medicine using AI.

Drawing on Stanford's interdisciplinary expertise in clinical medical imaging, bioinformatics, statistics, electrical engineering, and computer science, AIMI supports the development, evaluation and dissemination of new AI methods applied across the medical imaging life cycle.

Center for Comparative Studies in Race & Ethnicity

The CCSRE Race and Technology Initiative links research, teaching, and practice to advance racial justice in the analysis, production, and deployment of new technologies.

Center for Ethics in Society

The McCoy Family Center for Ethics in Society is committed to bringing ethical reflection to bear on important social problems through research, teaching, and community engagement. Drawing on the established strengths of Stanford's interdisciplinary faculty, the Center develops initiatives with ethical dimensions that relate to pressing public problems.

Center for International Security and Cooperation

The Center for International Security and Cooperation tackles the most critical security issues in the world today. Founded in 1983, CISAC has built on its research strengths to better understand an increasingly complex international environment. It is part of Stanford's Freeman Spogli Institute for International Studies.

Cyber Policy Center

The Cyber Policy Center at the Freeman Spogli Institute for International Studies is Stanford University's premier center for the interdisciplinary study of issues at the nexus of technology, governance and public policy.

Digital Civil Society Lab

The Digital Civil Society Lab at Stanford PACS (Center on Philanthropy and Civil Society) engages researchers, practitioners and policymakers across four domains: Technology, Organizations, Policy and Norms.

John S. Knight Journalism Fellowships

The JSK Journalism Fellowships supports diverse journalists from around the world who are deeply engaged in exploring solutions to journalism's biggest problems. It focuses on accelerating change in the journalism industry to improve the access to information people need to create and sustain democratic communities.

Office of the Vice President of the Arts

The Vice Presidency for the Arts at Stanford University was established in February 2017 to elevate the arts in the university's priorities and lead strategic planning for the university's arts goals.

Stanford Institute for Economic Policy Research

The Stanford Institute for Economic Policy Research (SIEPR) is a research organization committed to scholarship that helps address the real-world challenges facing governments and businesses in the United States and around the world. SIEPR's goal is to raise living standards and improve the quality of life by making economic policy more effective at all levels.

External Partners

AI4ALL is a nonprofit working to increase diversity and inclusion in artificial intelligence. it creates pipelines for underrepresented talent through education and mentorship programs in the U.S. and Canada that give high school students early exposure to AI for social good. AI4ALL's vision is for AI to be developed by a broad group of thinkers and doers advancing AI for humanity's benefit. AI4ALL and HAI co-sponsor a summer internship program for AI4ALL alumni, in cooperation with CSLI.

Corporate Members Program

HAI actively engages with companies that share our mission to advance AI research, education, policy, and practice to improve the human condition. To date, HAI has worked with industry stakeholders, including our Founding Corporate Members to foster dialogue and identify research and educational opportunities. In FY20, Google and Wells Fargo joined AWS and IBM as Founding Corporate Members. These efforts include collaborations with HAI partner research centers on topics including the future of work and AI applications in healthcare. We also distributed $2 million in AWS cloud credits acrosss campus and began issuing Google cloud credits.

Philanthropic Support

Stanford HAI gratefully recognizes the generosity of those who provided founding support and make our work possible. Contributions made prior to August 31, 2020, are listed on the HAI website and will be updated annually, per our fundraising policy.

Events

Events are a cornerstone of HAI's offerings to the community, from Directors' Conversations to weekly seminars to frequent workshops and large conferences. Our events program commenced with the 2019 HAI Symposium, the institute's launch event held in March.

Our 2019 Fall Conference convened experts and leaders from academia, industry, civil society, and government for two days to explore critical and emerging issues related to the topics of AI Ethics, Policy, and Governance. Speakers included Joy Buolamwini, Founder, Algorithmic Justice League; Michael Kratsios, Chief Technology Officer of the United State; Reid Hoffman, Co-founder, LinkedIn and Partner, Greylock Partners; and Eric Schmidt, Technical Advisor, Alphabet Inc.

During the onset of the pandemic, HAI held two virtual conferences on COVID-19 and AI, one on April 1 as information was emerging, and another on June 1, looking at the road ahead. Over the 2019–20 academic year we held workshops on such topics such as AI and Environmental Intelligence, AI and Ethics, AI and Cognition, AI and International Security, HAI AI and Labor Markets, Regulation and Artificial Intelligence, and Facial Recognition Technology. We completed the events program with seminars covering diverse topics including The Future of Work, Machine Learning, AI Governance, and Natural Language Processing.

85,000
viewers of COVID-19 virtual events

14
workshops

35
virtual seminars

Financials

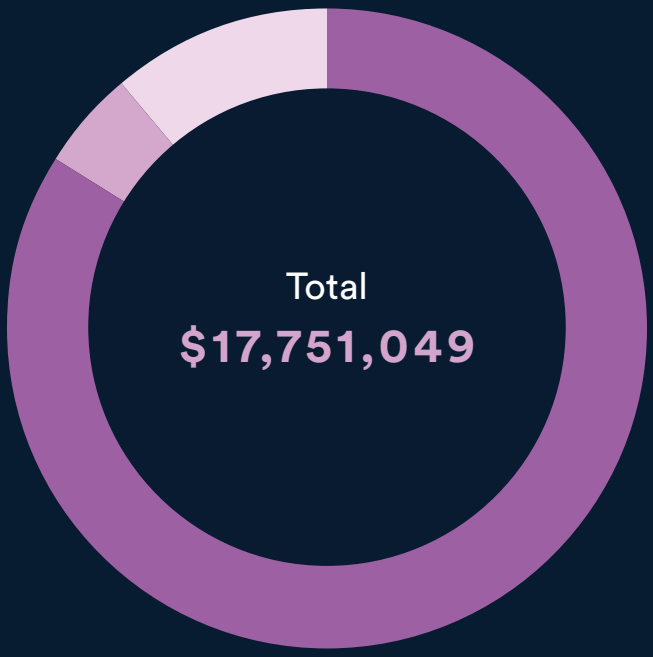

Sources of Income[1]

GIFTS	$14,920,001	84%
ENDOWMENT PAYOUTS	$831,048	5%
CLOUD CREDITS (IN KIND)	$2,000,000	11%

HAI Research Support

2018 SEED GRANTS	$1.2m
2019 SEED GRANTS	$1.8m
2020 SEED GRANTS[2]	$2.5m
2020 HOFFMAN-YEE GRANTS[2]	$3m
2020 AWS CLOUD CREDITS	$2m
Total	$10.5m

1 Income does not include increases in endowment principal, but includes endowment payout

2 Disbursed in FY2021

3 In FY20, $3 million in Hoffman-Yee research grants were awarded but were not disbursed until FY21

Expenditures

OPERATIONS
$2,213,728 — 22%

POSTDOCS AND FELLOWS
$942,590 — 10%

FACULTY
$516,534 — 5%

RESEARCH PROGRAMS[3]
$2,430,753 — 25%

PARTNERSHIPS, COMMUNICATIONS, AND ENGAGEMENT
$2,004,509 — 20%

EDUCATIONAL PROGRAMS
$107,071 — 1%

POLICY PROGRAMS AND RESEARCH DISSEMINATION
$153,212 — 2%

WORKSHOPS AND EVENTS
$762,261 — 8%

OTHER RESEARCH SUPPORT AND FACULTY INCENTIVES
$473,732 — 5%

SUPPORT TO STANFORD PARTNER CENTERS
$225,624 — 2%

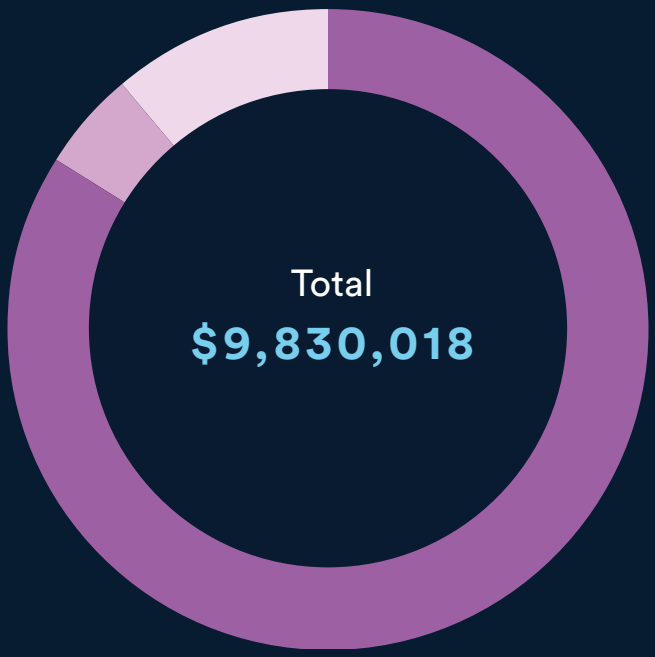

Thank you for taking the time to learn about the work of Stanford HAI. We appreciate your support and encourage you to contact us with any questions.

Cordura Hall
210 Panama Street
Stanford University
Stanford, CA 94305

HAI-INSTITUTE@STANFORD.EDU

HAI.STANFORD.EDU

Follow us on Twitter, LinkedIn, Facebook, Instagram, and YouTube:

@STANFORDHAI

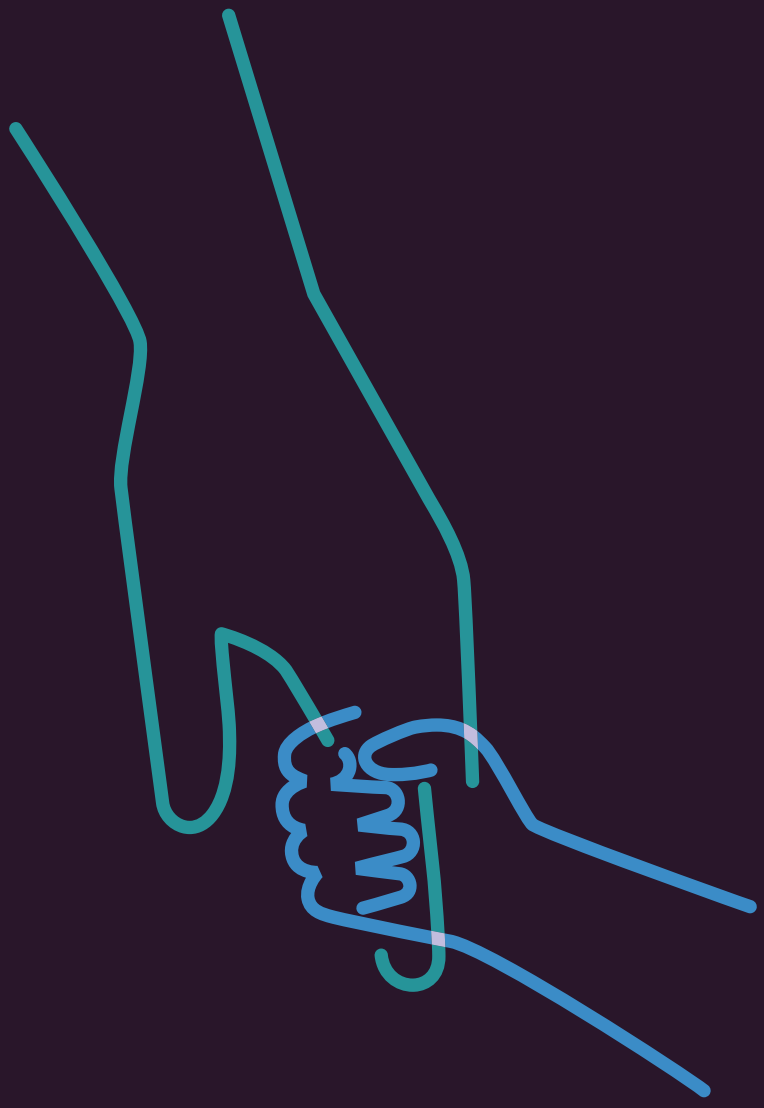

Stanford University
Human-Centered
Artificial Intelligence

Stanford