


2020–2021 Annual Report

Stanford Institute for Human-Centered Artificial Intelligence

See the progress HAI has made in research, policy, and education during our second academic year.



Contents

- 1 Letter from the Denning Co-Directors
- 2 HAI by the Numbers

- 5 **Research**
- 6 HAI Research Focus
- 7 Grant Programs for AI Research
- 10 Center for Research on Foundation Models
- 11 Stanford Digital Economy Lab
- 12 AI Index
- 13 Research Events / Ethics and Society Review

- 15 **Policy**
- 16 National Research Cloud / Education and Training for Policymakers
- 17 Policy Steering Committees / Policy Resources / Policy Workshops

- 19 **Education**
- 20 Programs for Today's Leaders / AI+Activism+Art
- 21 Programs for Emerging Leaders

- 22 **People**
- 22 Denning Co-Directors
- 23 Associate Directors
- 25 Staff
- 26 Faculty and Fellows
- 27 Junior Fellows Program / Visiting Artist

- 28 Events
- 29 Corporate Members Program
- 30 Research Centers and Partners
- 31 Content / Advisory Council / Support for HAI
- 32 Financials

Letter from the Denning Co-Directors



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

As we write this letter, we can't help but reflect on how the development of the Institute for Human-Centered Artificial Intelligence (HAI) since our founding in 2019 has taken a course no one could have predicted. Despite the challenges of the ongoing pandemic, HAI has grown in size and strength as an organization and campus community, all while we were working remotely. A silver lining of working from home was that it enabled us to entirely remodel the Gates Computer Science Building, providing HAI with its first home. Please come visit!

During the 2020-2021 academic year our milestones included progress with the National Research Cloud, a significant program for EU policymakers, exciting findings from our first six Hoffman-Yee grant recipients, and our first executive education class. We invite you to read about these advancements and more in this, our second annual report.

We sincerely appreciate the support of our growing community of faculty, fellows, and students, as well as scholars and researchers beyond Stanford, around the world. Our supporters, advisors, and corporate members have been generous with their time and resources, while the HAI staff has worked tirelessly to move the organization forward. We simply could not achieve this important work without the contributions of our community.

The landscape of AI has changed markedly in the last three years. The cost of technical AI research has skyrocketed, commensurate with the ever-increasing capabilities and complexity of AI. For example, a new class of large-scale neural networks has the ability to generate text, audio and imagery, transforming the field of AI in ways we are only beginning to understand. This past year HAI responded by establishing a Center for Research on Foundation Models that involves hundreds of Stanford scholars across many different fields, contributing to research and driving important conversations about the risks and opportunities of these large models.

Today, conversations about AI are more charged than ever, spanning tough topics like bias, facial recognition, fake generative data, large models and more. At HAI we believe it's our responsibility to engage in challenging debate, facing questions without easy answers, and learning from failure.

Even after a decade of historic advances, AI's most impactful discoveries have yet to be made. Academic research, by Stanford scholars and others, is key to pushing the frontiers of basic science, both for the sake of discovery and to evaluate the ethical and societal implications that industry developers may not consider, or may not be equipped to address.

As the pandemic winds down, AI is clearly intertwined with the most critical issues of our time: geopolitics, work, climate, healthcare, and more. Courage has never been more important in the field of AI—to deeply examine evidence-based research, to listen to all voices, and to have conviction of reason. We look forward to ongoing collaboration with you in the coming year and welcome your ideas and participation in our journey.

Fei-Fei Li


John Etchemendy

HAI by the Numbers

Although we measure our success by impact, these facts paint a picture of the scope and depth of our work over the 2020-2021 academic year.

108

Stanford faculty funded by HAI



113

blog articles featuring the work of HAI faculty, students, and fellows



All 7

Stanford schools represented by HAI faculty


\$11M

issued in grants



102

projects funded by HAI grants



120,000

views of HAI virtual events






Research

Championing Human-Centered AI

Interdisciplinary work is key to solving society's most complex problems. During the 2020-21 academic year, HAI continued to promote human-centered AI technologies and applications. We are proud to have funded dozens of new faculty across all seven Stanford schools, bringing new voices, perspectives, and talent into the community.

HAI Research Focus

We are committed to ensuring that humanity benefits from AI technology and that the benefits are shared broadly by all. In support of this commitment, our research falls into three key focus areas: Human Impact, Augmenting Human Capabilities, and Intelligence.



Human Impact

To develop equitable and trustworthy technology, we must understand how AI performs in practice and affects humanity. Our multidisciplinary research on this technology's human impact aims to realize this vision.



Augmenting Human Capabilities

AI should communicate and collaborate effectively with people to augment their capabilities and make their work better and more enjoyable. Breakthroughs in this research area will allow great progress in healthcare, education, sustainability, automation, and countless other domains.



Intelligence

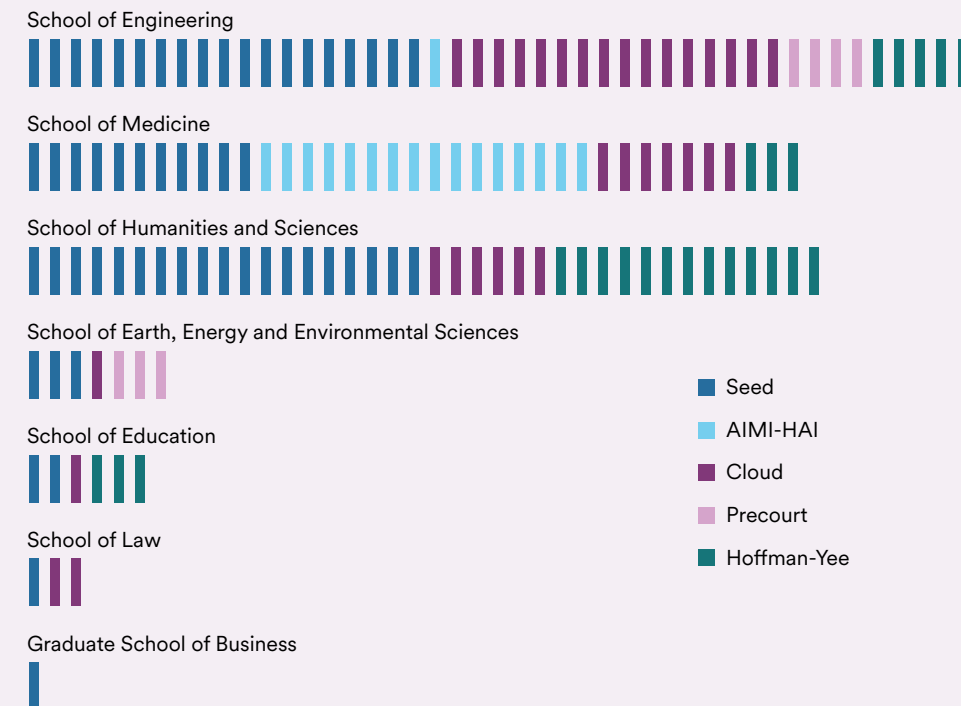
Current AI systems resist explanation or interaction in terms that are comprehensible by humans. Ultimately, we need to develop machine intelligence that understands human language, emotions, intentions, behaviors, and interactions at multiple scales.



Grant Programs for AI Research

HAI strives to foster a culture of AI research in which technological advances are linked inextricably to research about their potential societal impact. HAI builds on the strength of Stanford research by offering multiple [grant programs](#).

108 faculty funded in AY 2020-21



AIMI-HAI Grants for AI in Healthcare and Medicine

HAI contributed \$550,000 toward a first-time grant partnership with the [Stanford Center for Artificial Intelligence in Medicine and Imaging \(AIMI\)](#), providing a total of \$1.14 million to [six teams](#) over the next two years for healthcare AI research with near-term applications.

Precourt Institute for Energy Grants

HAI joined the Precourt Institute for Energy grants and the Bits & Watts initiative in funding \$1.2 million for new projects in energy AI and environmental justice. The first two Precourt Pioneering Projects were selected for their efforts using artificial intelligence and machine learning to make energy systems more sustainable, affordable, resilient, and fair to all socioeconomic groups.

HAI Seed Grants

HAI seed research grants support new, ambitious, and speculative ideas with the objective of getting initial results. These projects range from discrete studies and book-length research to speaker series and system building and evaluation. In AY 2020-21, HAI awarded 25 seed grants totaling approximately \$1.8 million to 54 faculty across all seven schools. Almost half of the funded faculty (24 people) had not been previously funded by HAI.

25 seed grants \$1.8 million
56 faculty

Hoffman-Yee Research Grants

Hoffman-Yee research grants are made possible by a gift from philanthropists Reid Hoffman and Michelle Yee. These grants support multidisciplinary teams that are leveraging AI to focus on real-world problems in health care, education, and society. In September 2021, Stanford HAI held the first Hoffman-Yee Symposium for inaugural grant recipients to present their research results to date and their plans for the future. Four of the six inaugural teams were selected for additional funding. Their work focuses on these goals:

- building robotics to improve human locomotion
- developing artificial tutors to augment teachers in the classroom
- understanding how humans play to improve AI tools
- investigating cultural concepts through time with AI

Each of these four teams will receive an additional \$500,000 to \$1 million for each of the next two years, on top of their original awards of \$500,000.

Featured HAI Seed Grants

Featured Grant in Human Impact

AI Accountability in Practice: What Keeps Companies from Implementing Socially Responsible AI?

This project examines the concrete organizational roadblocks shaping the implementation of FATE (fairness, accountability, transparency, and ethics) values as the technology industry designs and implements AI systems.

Featured Grant in Augmenting Human Capabilities

Using Wearable Electrodermal Activity Sensors to Identify ADHD Biotypes

This study examines the use of longitudinal measurement of electrodermal activity (EDA)—a peripheral index of autonomic arousal—as a biomarker to aid in the objective diagnosis of ADHD and augment the diagnosis of ADHD by identifying potential subtypes, or biotypes.

Featured Grant in Intelligence

Self-Supervised Representation Learning of Screen Data to Facilitate Behavioral and Reinforcement Learning Models

This team is using newly available data and analyses grounded in media psychology to develop visual representations that facilitate semantic labeling of human screen-based behavior and support development of new computational models of human learning and curiosity.

Cloud Credit Grants

Cloud credit grants from Google and Amazon Web Services (AWS) provide the cloud compute required for rigorous AI research. HAI allocates these credits to projects showing promising, novel, or emerging research.


During the 2020-21 academic year, 28 faculty completed 52 projects using \$2 million in AWS cloud credits awarded by HAI in AY 2019-20. HAI also launched the Google cloud credit program, distributing \$2 million to 36 faculty members across 66 projects. Each recipient was eligible for up to \$100,000 in credits for the year.

Google Cloud Credits

36 faculty 66 projects

AWS Cloud Credits

28 faculty 52 projects



Researchers created a computer-simulated playground where arthropod-like agents dubbed “animals” learn and evolve.

Featured Cloud Credit Grants

Featured Grant in Human Impact

Learning in Cooperative and Competitive Games

Cooperative and competitive games have been well-studied in game theory and machine learning literature. However, in real life there are many settings in which the optimal action is sometimes to cooperate and other times to compete. This project aims to discover how humans behave in these games, so we can investigate how to design a new AI-agent that increases cooperativeness and results in higher payoffs for everyone.

Featured Grant in Augmenting Human Capabilities

Advancing the Neuroscientific Discovery of Disease-Specific Phenotypes

This team received AWS credits to develop deep learning and statistical models for neuroscience applications. The work contributed to a better understanding of brain disorders and documented healthy neurodevelopment during adolescence.

Featured Grant in Intelligence

Embodied Intelligence Via Learning and Evolution

Animals display remarkable degrees of embodied intelligence by leveraging their evolved morphologies to learn complex tasks. This project investigates how intelligent behavior is a function of the brain, morphology, and the environment. Since the principles governing relations among these three functions require large-scale in silico experiments, the team has created a new framework to evolve AI agents with diverse morphologies to learn hard locomotion and manipulation tasks in complex environments.

Center for Research on Foundation Models

Stanford is the first university to launch an initiative to examine “foundation models”—an emerging class of AI models, including BERT, DALL-E, and GPT-3—that are trained on broad data using self-supervision at scale and can be adapted to a wide range of downstream tasks.

Foundation models represent a dramatic increase in capability in terms of accuracy, generation quality, and extrapolation to new tasks; however, they also pose clear risks in that they may be used for widespread disinformation, potential exacerbation of historical inequities, and problematic centralization of power.

In the human-centered spirit of HAI, the [Center for Research on Foundation Models \(CRFM\)](#) consists of an interdisciplinary team of 200+ students, faculty, and researchers across 10+ departments or schools, including computer science, economics, political science, law, and medicine. The center’s first report, “[On the Opportunities and Risks of Foundation Models](#),” provided an unprecedented look into the legal ramifications, environmental and economic impact and ethical issues associated with these new models. In August 2021, our launch event, called [Workshop on Foundation Models](#), brought leaders from the broader community to discuss and debate opportunities and risks, technological considerations, industry applications, and potential societal harms.



Percy Liang, Director of the Center for Research on Foundation Models

Percy Liang is an associate professor of computer science at Stanford University and the director of the Center for Research on Foundation Models. His research spans many topics in machine learning and natural language processing, including robustness, interpretability, semantics, and reasoning. Percy’s awards include the Presidential Early Career Award for Scientists and Engineers (2019), IJCAI Computers and Thought Award (2016), an NSF CAREER Award (2016), a Sloan Research Fellowship (2015), and a Microsoft Research Faculty Fellowship (2014).

As the director of CRFM, he is leading Stanford’s efforts to make fundamental advances in the study, development, and deployment of foundation models through an interdisciplinary group of faculty, students, postdocs, and researchers.

Stanford Digital Economy Lab

The [Stanford Digital Economy Lab](#) brings together leading researchers and experts to pursue a deeper understanding of the digital economy.



Erik Brynjolfsson, Director of the Stanford Digital Economy Lab

Jerry Yang and Akiko Yamazaki Professor and Senior Fellow at HAI, 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

Christie Ko, Executive Director



Research

After one full year in operation, the lab currently supports more than 50 projects across four major research areas: AI and the Future of Work, Measuring the Digital Economy, Digital Platforms, and Data-Driven Decision-Making and Management Practices. The lab prioritized the scale-up of the following two topics:

AI and the Future of Work: We are investigating the effects of the current, ongoing phase of automation and augmentation of the workforce; the introduction of AI and robotics on employment, wages, and inequality; and the nature of work in the future.

Measuring the Digital Economy: We are working to create better methods of measuring the health and progress of an increasingly digital economy.

Events and Community

Our November 2020 launch event, AI & The Future of Work, featured an all-star lineup of thought leaders, including James Manyika, Gina Raimondo, and Eric Schmidt.

The lab hosted 15 seminars featuring research from world-renowned experts, including Cornell University Professor Erica Groshen, Columbia University Professor Joseph Stiglitz, and London School of Economics Professor Phillippe Aghion.

Publications

The lab published more than two dozen academic papers by affiliated faculty, fellows, and students. We also launched a [YouTube channel](#) and published a [working paper series](#) to showcase new research and voices.

Research Program

AI Index

In March 2021, HAI published the fourth edition of its AI Index, one of the most comprehensive reports about AI to date. The [AI Index](#) measures and evaluates the rapid rate of AI advancement, enabling leaders and decision makers to take meaningful action to advance AI responsibly and ethically with humans in mind.


The 2021 report significantly expanded the amount of data tracked compared to the prior index; this latest report covered hiring data, corporate investment, global policy and regulation, technical advances and academic trends, diversity in the workforce, and ethical challenges. The AI Index team also worked with a broader set of academic, private, and civil society organizations to gather and calibrate the information. The fourth edition also shows the effects of COVID-19



on AI development from multiple perspectives. Most notably, data demonstrates that the pandemic had little effect on AI hiring and private investment.

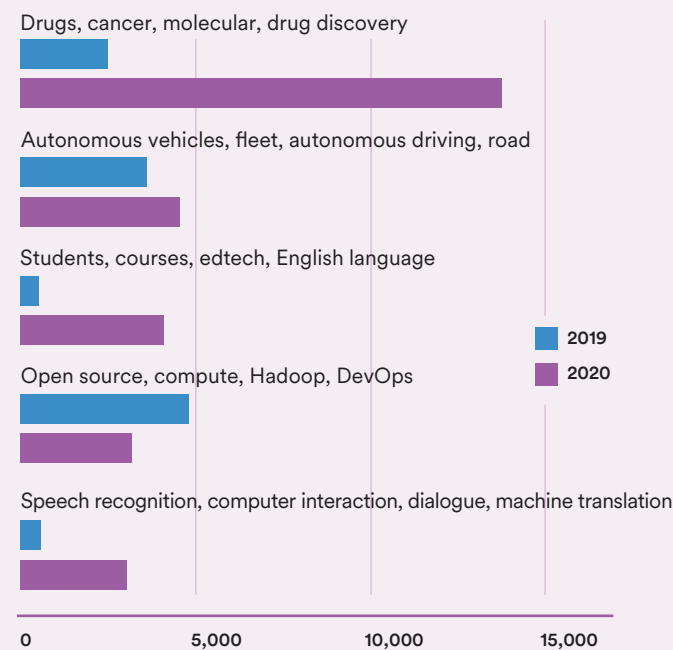
HAI promoted the 2021 report through two public events and five virtual briefings with policymakers and industry leaders, and the institute received considerable positive press coverage in both business and technology media.

Number of Paper Titles Mentioning Ethics Keywords at AI Conferences, 2000-19



Prates et al., 2018; adapted from the AI Index 2021 Annual Report

Global Private Investment in AI by Focus Area, 2019 vs 2020



Source: CapIQ, Crunchbase, and NetBase Quid, 2020; adapted from the AI Index 2021 Annual Report

Research Program

Research Events

Each spring and fall, the HAI research team organizes weekly seminars and frequent workshops to bring the AI community together around common interests. In AY 2020-21, HAI hosted 27 seminars featuring scholars from Stanford and beyond across a wide range of topics, from collaborative robots to ethical AI, and healthcare AI applications to embodied intelligence. This year's seminars included DeepLearning founder Andrew Ng's discussion around bridging [AI's proof-of-concept to production gap](#), and HAI Visiting Artist Rashaad Newsome's innovative creation "[BEING: The Digital Griot](#)." On average, about 300 attendees joined for each virtual live event.

In October 2020, researchers from OpenAI, HAI, and other universities convened to discuss open research questions surrounding GPT-3, the largest publicly disclosed, dense language model at the time. Participants represented a variety of disciplines including computer science, linguistics, philosophy, political science, communications, and cyber policy. The new [Center for Research on Foundation Models](#), organized by HAI and made up of more than 120 faculty and students, emerged as a result of these discussions.



HAI research events bring together the AI community of scholars for discussion and collaboration.

Ethics and Society Review

For the first time at Stanford, a new program piloted with support from HAI requires AI researchers to evaluate the potential societal implications of their proposed research before receiving funding from HAI. The Ethics and Society Review (ESR) requires researchers seeking funding to consider how their proposals might pose negative ethical and societal risks, to come up with methods to lessen those risks, and, if needed, to collaborate with an interdisciplinary faculty panel to ensure those concerns are addressed before funding is received.



The Ethics and Society Review creates a path for researchers to consider the wider ramifications of their work before embarking on research.

The program has reviewed more than 70 grants to date and is in the process of expanding beyond HAI to other funding programs on campus. This work was published in the [Proceedings of the National Academy of Sciences](#) and after its initial funding by HAI, the program has received additional funding from the Public Interest Technology University Network; Stanford's Ethics, Science, and Technology Hub; HAI; and the National Science Foundation. The Patrick J. McGovern Foundation also has approved the ESR for funding through a gift to the Center for Advanced Study in the Behavioral Sciences (CASBS).



Policy

Convening Global Discussions

Stanford HAI plays a leading role in producing critical scholarship on AI governance and appropriate uses, as well as serving as a convener of national and global discussions. By engaging AI leaders across all sectors and training policymakers around the world on these issues, HAI leaders have positioned the institute to be a natural hub for policy impact in human-centered AI. In AY 2020-21, HAI published 10 policy briefs, established three policy steering committees, and held several pivotal events for legislators and policy experts.

National Research Cloud

On January 1, 2021, the U.S. Congress authorized the National AI Research Resource Task Force Act as part of the National Defense Authorization Act for Fiscal Year 2021. The Biden administration established the task force to study and provide an implementation pathway to create world-class computational resources and robust government datasets for researchers across the country in the form of a National Research Cloud (NRC). HAI Denning Co-Director Fei-Fei Li was named to the task force, and a Stanford Law School policy lab composed of graduate students in computer science, business, engineering, and law completed six months of investigation into the issue. The resulting white paper, [Building a National AI Research Resource](#), was released in October 2021.



Policy Steering Committees

In early 2021, HAI established three faculty-led policy committees on subjects that represent unique opportunities for HAI:

- The Committee on Healthcare Policy, chaired by [Curt Langlotz](#) from Stanford School of Medicine
- The Committee on International Security, chaired by [Amy Zegart](#) from the Freeman Spogli Institute (FSI) and Hoover Institution
- The Committee on Technology Governance, chaired by [David Engstrom](#) from Stanford Law School

Policy Workshops

HAI hosted three policy workshops in AY 2020-21:

- Opportunities and Challenges for Civil Society
- Assessing and Managing Risks with the Use of Artificial Intelligence
- Breakthroughs and Barriers for Industry

Education and Training for Policymakers

European-North American Intensive Policy Lab

In May, HAI convened this lab together with Brookings and the European University Institute. The event engaged 28 current and former senior policymakers, as well as 24 prominent speakers. During the multiday event, the group engaged in critical conversations around the economics of AI, international security and cooperation, and regulating emerging technologies.

Congressional AI Caucus Briefing

HAI hosted a briefing on various AI policy subjects for members of the Congressional Artificial Intelligence Caucus.

Congressional Testimony

In September 2020, Susan Athey, Economics of Technology Professor at the Stanford Graduate School of Business and HAI associate director, provided [congressional testimony](#) to the House Budget Committee on the impact of AI on economic recovery and the future of work.

National Research Cloud Feedback Session

The co-authors of the National Research Cloud white paper hosted a feedback session on the NRC executive summary for approximately 50 participants from other universities, government, industry, and civil society organizations.

Policy Resources

HAI publishes numerous policy briefs, explainers, and white papers throughout the year to help lawmakers understand issues related to AI. Highlights from AY 2020-21 include these publications:

- [AI's Promise and Peril for the U.S. Government](#), by Dan Ho
- [The Moment of Reckoning: AI and the Future of U.S. Intelligence](#), by Amy Zegart
- [European Regulation on AI](#), by Marietje Schaake
- [Improving AI Software for Healthcare Diagnostics](#), by David Larson, Curt Langlotz, and Daniel Rubin





Education

Programs for Current and Future Leaders

HAI creates multidisciplinary programs for both Stanford affiliates and the general public. Our programs help people pursue cutting-edge research, create useful and responsible AI, and respond thoughtfully to the societal and ethical implications of global AI. In AY 2020-21, we held multiple events for each audience.

Programs for Today's Leaders

High-impact decision makers need to stay abreast of rapid advancements in AI technology and applications so they can advocate for appropriate human-centered development and policies.

Harnessing Design Thinking, Ethics, and the Power of AI for Social Innovation

Together with the Stanford Social Innovation Review (SSIR), HAI hosted a [two-part webinar](#) for executives in social innovation nonprofit and civil society organizations. Part one focused on Integrating Design Thinking in AI for Social Innovation, and part two covered Ethical and Societal Implications of AI.

The Value of Data and AI: Strategies for Senior Leadership

This [course](#) for executives, launched in January 2021, helps company leaders understand how to leverage their data and gain a competitive advantage.

AI+Activism+Art

In the spring quarter of 2021, [Michele Elam](#), William Robertson Coe Professor of Humanities and HAI associate director, introduced a pilot class, [AI+Activism+Art](#). This course puts arts and humanities right at the center of conversations about the impact of AI on humans and society at large.

Students enrolled from a wide range of fields, including computer science, art, English, comparative studies in race and ethnicity, African and African American studies, symbolic systems, and science, technology and society. Throughout the quarter, they examined AI from the perspective of feminist, gender, and sexuality studies; decolonialism, and race studies. Michele featured lectures by visiting artists and technologists, classroom discussion, and a final project of either a scholarly paper or a creative interdisciplinary project.



I came into Stanford as a humanities major and always wondered how to connect the arts and computer science. This course reassured me that I could bridge them and that there was a lot to be pursued at this intersection.



Isabelle Lee, a rising senior studying computer science

Programs for Emerging Leaders

Future leaders will face immense challenges as AI continues to shape society. Our programs for this audience seek to build awareness and community around key issues.



Stanford AI4ALL

This [summer camp](#) program for high school students from underserved populations provides hands-on learning and a supportive peer community. In June-July 2021, attendees worked closely with AI practitioners and went deep into new ideas that are defining the AI landscape.

Technology and the 2020 Election Project

In Fall 2020 this was offered as a Stanford course, a Continuing Studies course, a Coursera course, and a public event, engaging more than 8,000 students across all platforms.

Art+Tech Monthly Salon

Launched for the Stanford community through the Office of the Vice President for the Arts, this [program](#) focuses on building connections at the intersection of art and technology.

Summer Faculty Speaker Series Highlights

- AI for Good: The Promise and the Peril, with [Jeremy Weinstein](#)
- Artificial Intelligence in Space Exploration, with [Marco Pavone](#)
- Artificial Intelligence, Activism, and Art, with [Michele Elam](#)

People

Passionate people are essential to the work and progress of our organization. From our leaders and advisors to faculty, fellows, and staff, we all share a commitment to the mission of guiding human-centered AI.

Denning Co-Directors

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.



John Etchemendy

Provost, Emeritus, and Patrick Suppes Family Professor in the School of Humanities and Sciences

In AY 2020-21, [John Etchemendy](#) provided leadership across the HAI community, including as a member of the AI Index Steering Committee. He taught “Philosophy of Artificial Intelligence” to undergraduate students and was elected to the American Academy of Arts and Sciences.



Fei-Fei Li

Sequoia Professor, Computer Science Department

Among her many contributions as co-director of HAI, [Fei-Fei Li](#) was named to the National Artificial Intelligence Research Resource Task Force and participated in HAI’s most viewed event of the year, Healthcare’s AI Future: A Conversation with Fei-Fei Li & Andrew Ng. She was elected to the National Academy of Medicine and to the American Academy of Arts and Sciences.

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

Despite a sabbatical in 2021, [Russ Altman](#) helped lead the HAI Spring Conference on Intelligence Augmentation: AI Empowering People to Solve Global Challenges. He was also the faculty director for the AI100 project, leading the publication of its second report, and contributed to the report on “Opportunities and Risks of Foundation Models.”



Susan Athey

Economics of Technology Professor, Senior Fellow at the Stanford Institute for Economic Policy Research and Professor, by courtesy, of Economics

[Susan Athey](#) was named president-elect of the American Economic Association and continued to lead the Golub Capital Social Impact Lab, which deployed an AI-based recommendation system for an educational technology application, leading to a 30 percent improvement in user engagement. She taught a course in which engineers and business students collaborated to use AI and experimentation in projects for social impact organizations. Susan also testified for the U.S. House Budget Committee on Artificial Intelligence & the Workforce. Her papers on artificial intelligence and causal inference were published in *AIStats*, *ICML*, *PNAS*, *JASA*, *Econometrica*, and the *Journal of Econometrics*.



Michele Elam

William Robertson Coe Professor of Humanities, Department of English, Center for Comparative Studies in Race & Ethnicity, African & African American Studies

[Michele Elam](#) helped launch the HAI VPA Office Arts-Tech Salons and served on the Hoffman-Yee, Junior Fellows, Post-Doc and Education Committees. Michele was invited to speak at the Critical Borders: Radical (Re)Visions of AI conference, hosted by the Leverhulme Centre for the Future of Intelligence.

In the spring Michele piloted a new interdisciplinary course, “AI+Activism+Art.” She authored articles related to AI, art, and race in *Daedalus: Publication of the Academy of Arts and Sciences*, the *American Literature* journal, and *Feminist AI*. She presented in the HAI Summer Faculty Speaker Series on AI, Activism, and Art, as well as the capstone event with Ruha Benjamin at the 2021 Tech & Racial Equity Conference: Anti-Racist Technologies for a Just Future.



Surya Ganguli

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

Surya Ganguli was a co-lead of the HAI Fall Conference on Triangulating Intelligence: Melding Neuroscience, Psychology, and AI.



Daniel E. Ho

William Benjamin Scott and Luna M. Scott Professor of Law; Professor of Political Science; Senior Fellow, SIEPR; Faculty Fellow, CASBS; Faculty Director, Stanford RegLab

Daniel E. Ho co-authored “Building a National AI Research Resource: A Blueprint for the National Research Cloud.” He co-led research on AI to improve COVID-19 testing and contact tracing with the Santa Clara County Public Health Department, published in part in the *Proceedings of the National Academy of Sciences*.

Daniel was appointed to the Administrative Conference of the United States and was awarded the Carole Hafner Award for Best Paper at the International Conference on Artificial Intelligence and Law.



James Landay

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

James Landay continued to lead the Hoffman-Yee and Seed grant process for HAI, from calls for proposals through reviews and decisions. James was also the co-lead of the HAI Spring Conference on Intelligence Augmentation: AI Empowering People to Solve Global Challenges, and served on the Junior Fellows, Post-Doc and Education Committees.



Curtis P. Langlotz

Professor of Radiology and Biomedical Informatics; Director, Center for Artificial Intelligence in Medicine & Imaging

Interim Associate Director Curtis P. Langlotz chaired the HAI Steering Committee on Healthcare Policy and helped to establish and lead the Medical Imaging and Data Resource Center, a national repository to support the development of machine learning methods for the diagnosis and treatment of COVID-19. He authored a white paper proposing regulatory frameworks for the development and evaluation of AI in medical imaging, and he also collaborated on the first multi-institutional prospective clinical trial of a radiology AI algorithm.



Christopher Manning

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

Christopher Manning co-organized the HAI Fall Conference on Triangulating Intelligence: Melding Neuroscience, Psychology and AI. He co-authored a paper on emergent linguistic structure in artificial neural networks trained by self-supervision in the *Proceedings of the National Academy of Sciences of the United States of America*, as well as a paper in *Computational Linguistics* on universal dependencies (for language). Christopher contributed to a workshop and report on “Opportunities and Risks of Foundation Models” and released a 2021 update of the widely used, freely available online course CS224N Natural Language Processing with Deep Learning.



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 Reich published two important books in AY 2020-21: *Digital Technology and Democratic Theory*, edited with Lucy Bernholz and Hélène Landemore, and *System Error: Where Big Tech Went Wrong and How We Can Reboot*, with Mehran Sahami and Jeremy M. Weinstein. He joined the Center for Research on Foundation Models and contributed to the report and workshop on “Opportunities and Risks of Foundation Models.” In addition to these accomplishments, Rob continued to guide development of the Embedded EthiCS initiative.



Michael Sellitto, Deputy Director

Michael Sellitto led the institute through a substantial growth phase, more than doubling its size during the pandemic and significantly expanding its internal and external reach. He is also a co-author of the annual AI Index report and co-chairs an experts working group on national AI policy implementation at the Organization for Economic Cooperation and Development (OECD).

Staff

The staff of HAI serves to further the mission of the institute, working closely with faculty to design and implement programs and initiatives with global impact.

Krysten Hommel

Senior Associate Director, Strategic Initiatives

Amita Kumar

Director of Administration and Operations

Panos Madamopoulos-Moraris

Managing Director of Industry Programs and Partnerships

Vanessa Parli

Interim Director of Research

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 250 affiliated Stanford faculty from across all seven schools and more than 90 departments.

Featured Fellows



Sarah Bana, Stanford Digital Economy Lab Postdoctoral Fellow

[Sarah Bana](#) is passionate about the opportunities that unstructured data create to answer questions about the future of work. Her main projects use online job postings and syllabi to characterize valuable tasks and skills. These insights serve to provide improved strategic advice for firms and workers, along with improved policy prescriptions for governments.



Jennifer King, HAI Privacy and Data Policy Fellow

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. In AY 2020-21, she contributed to the HAI-sponsored National Research Cloud report, as well as three major areas of research: AI/algorithmic regulation, dark patterns, and data governance.

Junior Fellows Program

The HAI [Junior Fellows Program](#) is a prestigious program for exceptionally qualified 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. In AY 2020-21, our first Junior Fellow appointee [Johannes Eichstaedt](#) was joined by our second Junior Fellow, Hariharan Subramonyam.



Hariharan Subramonyam, HAI Junior Faculty Fellow

[Hari Subramonyam](#) is jointly appointed as the Ram and Vijay Shriram HAI Junior Faculty Fellow, and an assistant professor (research) in the Graduate School of Education. He is also a member of the Human-Computer Interaction Group at Stanford. Hari's research focuses on augmenting critical human tasks (such as learning, creativity, and sensemaking) with AI by incorporating principles from cognitive psychology. He also investigates support tools for multidisciplinary teams to co-design human-centered AI experiences.

VPA-HAI Visiting Artist

Rashaad Newsome

[Rashaad Newsome's](#) work blends collage, sculpture, film, photography, music, computer programming, software engineering, community organizing, and performance to create a new field that celebrates Black contributions to the art canon and creates inclusive forms of culture and media.

Working with Stanford students, he has created a text-generating machine learning model that allows the new generation of his primary work, *Being: The Digital Griot*, to develop their own poetry. The model was trained on the work of Oakland-based Black Queer poet Dazié Grego-Sykes, as well as bell hooks and Audre Lorde.

In addition, Rashaad surveyed members of the Black community at Stanford to discover what culturally specific sounds they find soothing. Using the material from this survey, he worked with composer Robert Aiki Aubrey Lowe to create a Black ASMR soundscape.



Fellows

HAI Fellowship Programs are designed to support researchers in exploring topics related to AI technologies. In the 2020–21 academic year, HAI sponsored 16 fellows, including several fellowships offered in partnership with other Stanford centers, including the following:

- 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)

2020–21 HAI Fellows

Elizabeth Adams, HAI-CCSRE Practitioner Fellow

Renata Avila, HAI-CCSRE Practitioner Fellow

Sarah Bana, S-DEL Postdoctoral Fellow

Kathleen Creel, Embedded EthICS Fellow

Jeffrey Ding, HAI-CISAC Predoctoral Fellow

Samir Doshi, HAI-CCSRE Practitioner Fellow

Jazmia Henry, HAI-CCSRE Practitioner Fellow

Todd Karhu, HAI-EIS Postdoctoral Fellow

Geoff Keeling, HAI-EIS Postdoctoral Fellow

Jennifer King, HAI Privacy and Data Policy Fellow

Jae Joon Lee, S-DEL Postdoctoral Fellow

J. Frank Li, S-DEL Postdoctoral Fellow

Hong Qu, HAI-CCSRE Practitioner Fellow

Marietje Schaake, HAI International Policy Fellow

Michelle Spektor, HAI-CISAC Predoctoral Fellow

Xiupeng Wang, S-DEL Postdoctoral Fellow

Stanford University

Events

Events are an integral part of the HAI mission. HAI hosted weekly seminars and several conferences in AY 2020–21, all designed to bring our community together in new ways that lead to insight and action:

- **September 2020**
Coded Bias Screening and Panel Discussion, co-sponsored by Stanford Arts
- **October 2020**
HAI Fall Conference on Triangulating Intelligence: Melding Neuroscience, Psychology, and AI
- **March 2021**
HAI Spring Conference on Intelligence Augmentation: AI Powering People to Solve Global Challenges
- **April 2021**
Healthcare's AI Future: A Conversation with Fei-Fei Li & Andrew Ng, HAI's most viewed event of the academic year
- **May 2021**
2021 Tech & Racial Equity Conference: Anti-Racist Technologies for a Just Future, co-sponsored with the Stanford Center for Comparative Studies in Race & Ethnicity (CCSRE), Digital Civil Society Lab (DCSL) at the Stanford Center on Philanthropy and Civil Society, and the Stanford Program in African & African American Studies, with support from the Public Interest Technology University Network

8 workshops 27 weekly seminars

Corporate Members Program

HAI recognizes that industry is at the forefront of AI, and we are committed to working closely with business leaders seeking to harness the opportunities posed by AI while understanding and mitigating its risks. To help establish these bridges, HAI published industry briefs on healthcare and life sciences, the future of work post-COVID, and financial services. HAI also launched the next phase of the Founding Members Program for industry stakeholders. It now includes a research fund that can be allocated according to members' priorities, new pathways of engaging with and recruiting students, and new partnerships with the Stanford Digital Economy Lab and the Stanford AI Lab. HAI founding members have been using the program's new structure to seed, support, and collaborate on research projects across Stanford, as well as to curate tailored professional education programs for their companies.

In addition to these efforts, HAI launched a Corporate Affiliate Program for companies interested in AI research, policy, and practice to interact with HAI. The inaugural focus of this program is on AI and financial services.



HAI offers corporate members access to Stanford researchers on the leading edge of AI, spanning multiple disciplines.



The strength of the HAI Corporate Program lies in the breadth of opportunity to engage with a wide range of faculty and departments. Our partnership helps to inform the strategic direction across IBM's portfolio of fundamental research. The expansion of the program's 'wallet' feature, along with the white-glove treatment shown by the HAI team, will allow us to increase research collaborations and generate even more measurable value.



Jeffrey J. Welser, COO IBM Research,
VP Exploratory Science & University Collaboration



The HAI Founding Members Program allows for customizable, project-based collaborations that engage participating Googlers with new topics, methods, and people in the domain of human-centered AI. We look forward to expanding on these initiatives—Google recognizes that HAI's interdisciplinary perspective is key to shaping an inclusive AI future.



Jeff Dean, Google Senior Fellow and SVP of
Google Research



The HAI Corporate Program facilitates an effective engagement among Stanford University, the larger industry trends, and a broad cross-section of our leadership and employees across the world. The innovative program structure has helped us engage with some interesting challenges specific to the financial services industry and bring new ways of thinking about the future.



Chintan Mehta, CIO, Head of Digital Technology
and Innovation, Wells Fargo

Research Centers and Partners

Close collaboration with labs, centers, and institutes across and beyond the Stanford campus helps HAI achieve progress toward 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 for 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 Research on Foundation Models (CRFM)

The Center for Research on Foundation Models is an interdisciplinary initiative born out of Stanford HAI that aims to make fundamental advances in the study, development, and deployment of foundation models for AI.

Center for the Study of Language and Information (CSLI)

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 (DAWN)

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. DAWN is a five-year research project to democratize AI by making it dramatically easier to build AI-powered applications.

Digital Economy Lab (S-DEL)

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

Golub Capital Social Impact Lab

The Golub Capital Social Impact Lab uses digital technology and social science research to improve the effectiveness of leading social sector organizations. Based out of Stanford GSB, the lab is a research initiative of affiliated academics and staff, as well as researchers and students, who are passionate about conducting research that guides and improves the process of innovation.

Open Virtual Assistant Lab (OVAL)

OVAL is creating an ecosystem founded on open virtual assistant technology that (1) democratizes AI for linguistic user interfaces, (2) creates an open and non-proprietary web, and (3) promotes sharing with individual data ownership.

Regulation, Evaluation, and Governance Lab (RegLab)

Stanford's RegLab partners with government agencies to design and evaluate programs, policies, and technologies that modernize governance.

Stanford Artificial Intelligence Laboratory (SAIL)

SAIL has been a center of excellence for artificial intelligence research, teaching, theory, and practice since its founding in 1962.

Stanford Partners

Center for AI in Medical Imaging

Center for Comparative Studies in Race & Ethnicity

Center for Ethics in Society

Center for International Security and Cooperation

Cyber Policy Center

Digital Civil Society Lab

Institute for Diversity in the Arts

John S. Knight Journalism Fellowships

Office of the Vice President for the Arts

Stanford Institute for Economic Policy Research

External Partners

AI4All

Content

On the HAI Blog, we highlight cutting-edge Stanford AI research, faculty thought leadership, and key insights from HAI scholars. In AY 2020-21, we posted 113 blog articles and 60 videos, featuring 160 scholars, and reached more than 74,000 social media accounts across Twitter, Facebook, LinkedIn, and Instagram. Our news and blog articles brought in more than 360,000 views to the website, representing a third of all HAI website traffic.

Advisory Council

The HAI Advisory Council consists of leaders in science, civil society, and business who provide HAI with expert advice on issues relevant to the institute's mission and programs. Reid Hoffman serves as chair of the council, along with vice chairs Steve Denning and Srinija Srinivasan.

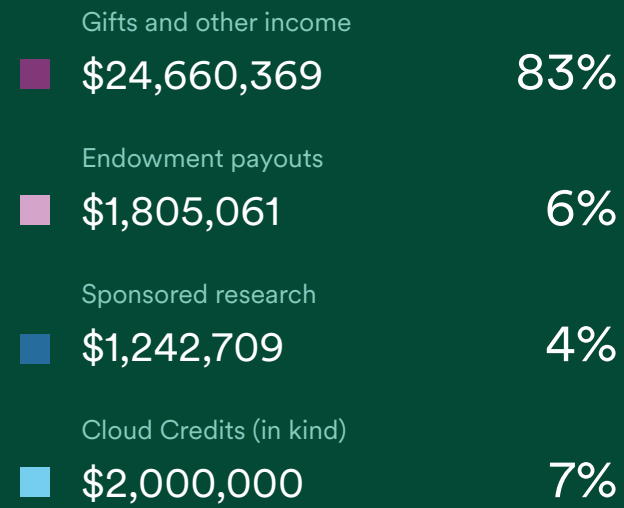
Support for HAI

Stanford HAI gratefully recognizes the generosity of those who provide support and made our work possible. Contributions made prior to August 31, 2021 are listed on the HAI website and will be updated annually, per our fundraising policy.




Financials

Sources of Income¹



Total \$29,708,140

¹ Income does not include increases in endowment principal



HAI Research Support

Seed Grants	\$2,684,144
Hoffman-Yee Grants	\$2,973,974
Total	\$5,658,118

From 2018-2020, HAI issued

\$8.5m

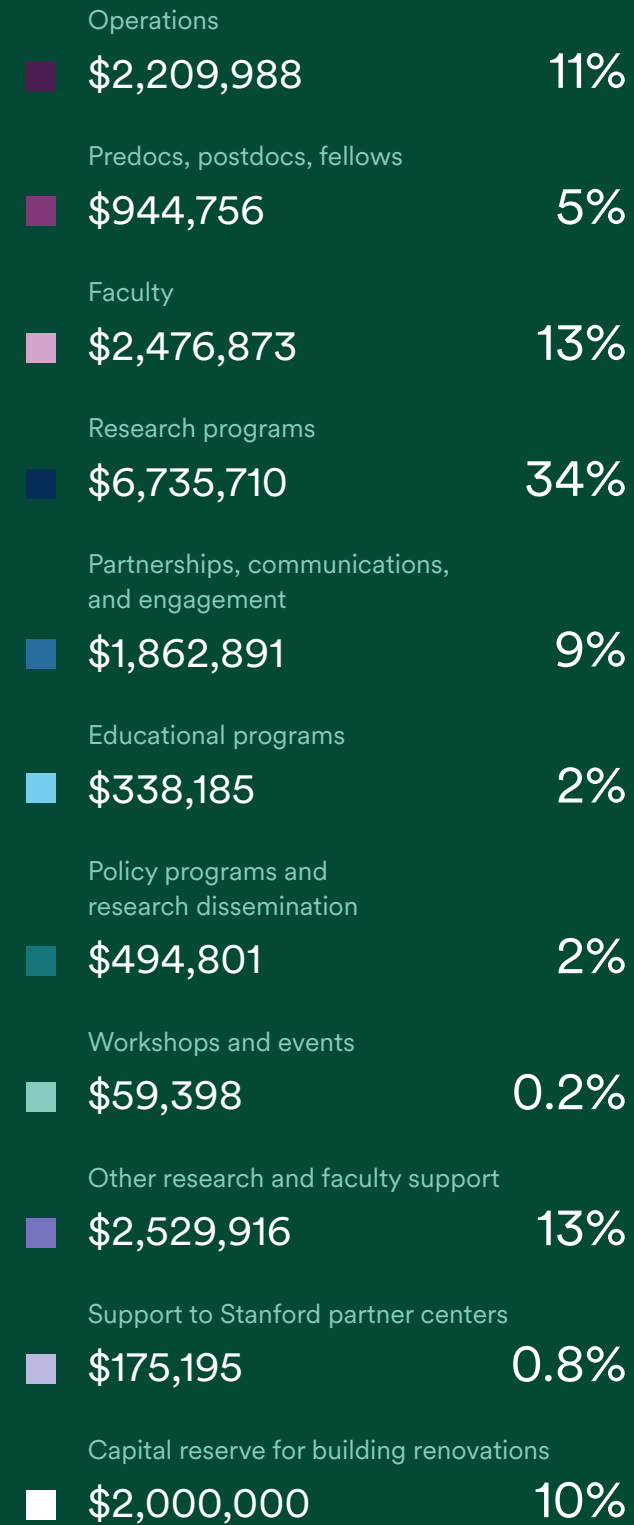
in cash grants.

Cumulatively to date, these projects have subsequently attracted


\$6.2m


in external funding.

Expenditures



Total \$19,827,717





Thank You

Thank you for taking the time to learn about our ongoing work across the AI landscape. We appreciate your support and encourage you to contact us with any questions.

353 Jane Stanford Way
Stanford University
Stanford, CA 94305

Visit us

hai-institute@stanford.edu

Reach out

hai.stanford.edu

Subscribe to our

[newsletter](#)

Follow us

