Stanford University
Human-Centered
Artificial Intelligence

Stanford
University

Annual Report 2022

Stanford Institute for Human-Centered Artificial Intelligence

HAI built considerable momentum across our research, policy, and education programs during the 2021-22 academic year.

Contents

Research

Policy

Education

This Annual Report reflects the research, activities, and impact of Stanford HAI between September 2021 to August 2022

About HAI

Letter from the Denning Co-Directors

Silicon Valley startups are notorious for the demands they place on early employees—and the more successful they are, the more demanding they can be. For the past three years, Stanford HAI has been an academic startup. We have accomplished a lot since our launch, and we are deeply grateful to our hardworking staff, faculty leaders, and generous donors for getting us where we are today.

In three years, HAI has established itself not only as a leading presence at Stanford, but as one of the leading centers for AI in the world. We've earned a reputation as a catalyst for innovative cross-disciplinary research and have developed a dynamic ecosystem for the study and advocacy of a distinctively human-centered approach to AI.

As we reflect on the past academic year, we're excited to see the Stanford campus active again with faculty, staff and students back in person. We've settled into a new home on the ground floor of the Gates Computer Science Building, where we can host groups of people from across the campus and beyond.

Our foundational programs in research, education and policy are all thriving in the wake of an extremely productive year. Among the highlights from the 2021-2022 academic year:

- **In Research:** We've strengthened our Stanford scholar community through more than 20 in-person and hybrid format workshops, meetings and collaborations, while also collaborating with external organizations including the IEEE on its Planet Positive 2030 campaign.
- **In Education:** More than 7,000 individuals took part in our educational programs, and we helped to organize the largest AI+Health conference in Stanford's history.
- **In Policy:** We engaged with more than 30 U.S. federal agencies, congressional offices, foreign government representatives and delegations this past academic year, which deepened HAI's relationships with policymakers and accelerated their understanding about the impact of AI.

We sincerely appreciate the support and engagement of our advisors, affiliates, scholars and students–at Stanford as well as other universities and organizations worldwide–in helping us achieve these goals.

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

In the broader world of AI, 2021-2022 has been an extraordinary year, one that has convinced even the most diehard skeptics of the importance of AI technology. Gone are any worries about a coming "AI winter" or the sudden deflation of a hype-fueled bubble. It is impossible to ignore such developments as DeepMind's near complete solution of the protein-folding problem, or the steady release of new foundation models for generating text, images and even computer code.

These developments are being addressed, from very different perspectives, by our Center for Research on Foundation Models and our Digital Economy Lab, the former focusing on expanding research, resources and community around emerging AI technology, while the latter assesses its impact on the economy and the future of work.

The breathtaking pace of AI technology only magnifies the importance of an institute devoted to a thoroughly human-centered perspective on this rapidly evolving technology. It is undeniable that AI holds extraordinary promise for improving the human condition, but only if we can successfully guide its future in a responsible direction.

Thank you for being with us on this most important journey.

Fei-Fei Li John Etchemendy

HAI Events

HAI brings its extended community together by hosting a variety of events throughout the year, including workshops, conferences, and seminars. The institute held 81 events during AY 2021-22, reaching a total 19,475 attendees between in-person and virtual platforms.

Hoffman-Yee Symposium

Inaugural recipients of Hoffman-Yee Research Grants presented results from their research to date and plans for the future.

HAI Fall Conference on Policy & AI

Four Radical Proposals for a Better Society

Proposals were presented to a panel of experts from multiple disciplines and backgrounds, who vetted, debated, and judged the merits of each one.

Strengthening the Technical Foundations of U.S. Security

HAI Managing Director for Policy and Society Russell Wald, CSET Senior Fellow Andrew Lohn, and Stanford HAI Postdoctoral Fellow Jeff Ding discussed how a National AI Research Resource (NAIRR) will impact U.S. national security.

HAI Spring Conference

Key Advances in Artificial Intelligence

The HAI Spring Conference explored three key advances in artificial intelligence—accountable AI, foundation models, and embodied AI in virtual and real worlds—as well as what the future of this technology might hold.

Artificial Intelligence and the Economy: Charting a Path for Responsible and Inclusive AI

This symposium brought together leaders from government, industry, civil society, and academia to explore potential opportunities and challenges posed by artificial intelligence and machine learning deployment across different economic sectors.

Advancing Technology for a Sustainable Planet

Hosted together with IEEE and the Stanford Woods Institute for the Environment, this event explored regulatory, policy, and financial frameworks critical to advancing technology which prioritizes people, planet, and purpose-driven progress.

September · November · December · January · February · April · July · August

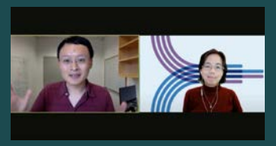

Data-Centric AI Virtual Workshop

This workshop explored challenges and opportunities across the data-for-AI pipeline.

AI+HEALTH

This recorded online conference convened experts and leaders from academia, government, and clinical practice to explore critical and emerging issues related to AI's impact across the spectrum of health, healthcare, and related arenas.

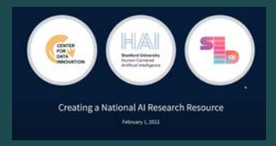

Creating a National AI Research Resource

HAI Managing Director for Policy and Society Russell Wald, Center for Data Innovation Policy Analyst Hodan Omaar, and Founder & President of SeedAI Austin Carson held a discussion on how to design, implement and maintain a National Artificial Intelligence Research Resource (NAIRR).

DEL Spring Workshop

Avoiding the Turing Trap: Workshop on AI & the Future of Work

This workshop brought together DEL's community of faculty, students, fellows, industry leaders, and supporters to share research and discuss groundbreaking ideas related to the future of work and the digital economy.

Stanford HAI Congressional Boot Camp on Artificial Intelligence

This multi-day, bi-partisan educational event for a diverse cohort of 26 congressional staffers unpacked what AI means for international security, the future of work, and healthcare, and included field trips to Stanford labs for interactive experiences.

Research

An Interdisciplinary Approach

HAI's research efforts gained significant momentum in the 2021-22 academic year. The team was excited to return to in-person events and to establish the Ethics and Society Review process as part of HAI's grant programs.

HAI Research Focus

HAI focuses its research initiatives around three main priorities, each with its own vision and mission.

H Human Impact

To develop technology that is equitable, researchers must understand how AI interacts with humans, as well as how it interacts with vital social structures and institutions.

Spotlight

Civics Education for a Just and Sustainable Future

Understanding and changing textbook content is one important lever in the multifaceted process required to redesign history and civics education. Patricia Bromley in the Graduate School of Education has designed a study that uses Natural Language Processing (NLP) to assess textbook content and document depictions of diversity, equity, and sustainability.

A Augment Human Capabilities

HAI believes that breakthroughs in human-centered design methods will further progress in healthcare, education, sustainability, automation, and countless other domains.

Spotlight

Rating Systems and the Future of Algorithmic Worker Evaluation

A team from the Schools of Engineering and Humanities and Sciences is developing a forward-looking digital resumé for workers in online labor platforms. The project brings together statistical methodology, algorithm design, human-computer interaction, and political economy to enable the same career pathways for online workers that society takes for granted in offline work.

I Intelligence

To create a machine-assisted—yet human-centered—world, the AI community must develop the next generation of techniques that overcomes the limitations of current algorithms, expands the class of problems that can be addressed, and complements human cognitive and analytic styles.

Spotlight

Toward Machine Models of Structured Human Reasoning through Explanation-based Meta-learning

Some of the most impressive achievements of the human mind have led to deep scientific understanding and powerful technologies. Two researchers from the Schools of Engineering and Humanities and Sciences seek to combine human and machine learning research to build artificially intelligent systems that could potentially achieve these abilities.

Grant Programs for AI Research

Since its founding, HAI's grant programs have supported 248 faculty members from all seven Stanford schools. Many of these projects span multiple departments, in keeping with the commitment to support interdisciplinary AI research.

How HAI Research Grants Connect Departments

Seed Research Grants

HAI seed research grants are designed to bring faculty together across disciplines to share knowledge and pursue new ideas. Three projects received notable awards to continue their research:

Decision Systems that Reflect Minority Perspectives

In April 2021, Computer Science Professor Michael Bernstein won the CHM Patrick J. McGovern Tech for Humanity Prize for his work toward introducing diversity and reducing bias in AI systems. The project includes a public-awareness website and open-source models for developers.

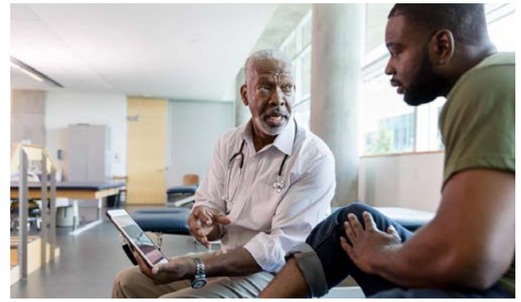

Prototyping an Intelligent Agent for CO_2 Sequestration in Saline Aquifers

This project addresses the urgent need to develop and scale commercial operations for sequestering CO_2 in the earth's subsurface. Professor of Geological Sciences Jef Caers received significant support from an Austrian company that is working toward sustainability.

Equitable Education Access in San Francisco

After winning seed funding in 2020 from HAI for a project project called "School Choice Mechanisms: Improving Diversity and Equity," Irene Lo and Itai Ashlagi from the School of Engineering received follow-on funding from Stanford Impact Labs to develop a method that San Francisco and other districts can use to improve equity in school assignment policies and practices.

90 total research grants since 2018

188 faculty funded in AY 2021-22

Hoffman-Yee Research Grants

Hoffman-Yee research grants are made possible by a gift from philanthropists Reid Hoffman and Michelle Yee. In August 2022, HAI announced its second cohort of Hoffman-Yee Grant recipients. Six research teams were each awarded a total of $2.75 million to solve some of the most challenging problems in the field of AI. Each of the proposals went through a rigorous review with faculty, as well as an Ethics and Society review.

Tuning Our Algorithmic Amplifiers: Encoding Societal Values into Social Media Algorithms

Main PI: Michael Bernstein

Co-PIs: Angele Christin, Jeffrey Hancock, Tatsunori Hashimoto, Nathaniel Persily, Jeanne Tsai, Johan Ugander

Imagine a future in which the AI algorithms that underpin social media platforms help us achieve societal goals rather than undermining them. A team of social science, engineering, and policy researchers is studying how to encode societal values into social media AIs.

MARPLE: Explaining What Happened through Multi-modal Simulation

Main PI: Tobias Gerstenberg

Co-PIs: Chelsea Finn, Noah Goodman, Thomas Icard, Robert MacCoun, Jiajun Wu

This project aims to bridge the gap between vision and language models by creating a way for AI systems to integrate different sources of evidence into a causal model of the world. The resulting framework will seek to mimic the human ability to infer what happened.

Foundation Models: Integrating Technical Advances, Social Responsibility, and Applications

Main PI: Percy Liang

Co-PIs: Russ B. Altman, Jeannette Bohg, Akshay Chaudhari, Chelsea Finn, Tatsunori Hashimoto, Dan E. Ho, Fei-Fei Li, Tengyu Ma, Christopher Manning, Christopher Re, Rob Reich, Dorsa Sadigh, Matei Zaharia

Foundation models represent a paradigm shift in AI. But developers need to improve their technical capabilities, while being sensitive to social responsibility and finding more real-world applications. A diverse team with expertise in machine learning, law, political science, biomedicine, vision, and robotics is on the case.

EAE Scores: A Framework for Explainable, Actionable and Equitable Risk Scores for Healthcare Decisions

Main PI: Carlos Ernesto Guestrin

Co-PIs: Carissa Carter, Emily Fox, Ramesh Johari, David Maahs, Priya Prahalad, Sherri Rose, David Scheinker

Researchers in medicine and engineering are collaborating to develop new AI algorithms and open-source tools that will transform the patient risk assessment process. Their ultimate goal is to improve clinical outcomes for Type 1 Diabetes, epilepsy, and opioid overdose risk.

Matching Newcomers to Places: Leveraging Human-Centered AI to Improve Immigrant Integration

Main PI: Jens Hainmueller

Co-PIs: Avidit Acharya, Yonatan Gur, Tomas Jimenez, Dominik Rothenhaeusler

Exactly where immigrants settle within a host country shapes the experience they have, as well as their contributions to the local economy and society. The goal of this project is to develop data-driven matching tools for location decision-makers—both governments and immigrants themselves.

Dendritic Computation for Knowledge Systems

Co-PIs: Kwabena Boahen, Scott W Linderman, H.-S. Philip Wong, Matei Zaharia

This team is working on a way to rein in AI's unsustainable energy, carbon, and monetary costs, distribute its productivity gains equitably, transform its users' experience, and restore their privacy. Their aim is to move away from learning with synapses (2-D) to learning with dendrites (3-D).

Cloud Credit Program

HAI saw record demand for cloud credit grants during AY 2021-22. This program provides advanced computational resources that researchers need to conduct rigorous AI projects. HAI allocates these credits to studies showing promising, novel, or emerging ideas.

In the second year of the Google Cloud Credit Grants, HAI awarded a total of $2.5 million in GCP credits to 48 faculty across 59 projects. The institute also launched the Microsoft Azure Cloud Credit Grants and awarded $330,000 in Azure credits to 7 faculty across 11 projects.

Featured Cloud Grants

Intelligence—Google

"Large-scale Evaluation of Visual Intuitive Physics," Dan Yamins, Assistant Professor of Psychology and of Computer Science

An artificial system that could mimic human-like intuitive physical prediction abilities, given real-world visual input in complex environments, would be useful for applications in computer vision, robotics, reinforcement learning, and other areas. To help advance research in this area of AI, this team plans to create a large-scale evaluation benchmark that captures a variety of aspects of intuitive physics. Afterward, they will evaluate a large set of state-of-the-art models on this new benchmark and provide a comprehensive comparison to human performance in a unified way.

Human Impact—Azure

"Foundation Models Homogenize Outcomes," Percy Liang, Director, Center for Research on Foundation Models, and Associate Professor of Computer Science

This study seeks to bring attention to a new class of potential harms that may arise with the increasing use of algorithmic systems—ones that are especially problematic at scale. The research looks at whether foundation models tend to homogenize outcomes, producing results that are more similar across different tasks or applications. The team aims to provide a scientific understanding of the principles that dictate how foundation models homogenize outcomes, as well as evaluate possible interventions to correct for homogenization.

Augment Human Capabilities—Google

"Online Surgical Analytics through Computer Vision via Multi-task Learning," Serena Yeung, Assistant Professor of Biomedical Data Science

Due to complexity and lack of suitable data training sets, AI has not been widely developed for open surgical procedures. This project aims to overcome existing limitations by curating the largest dataset of open surgical videos to date. Using this dataset, the team has developed a multi-task AI model that is capable of understanding surgical behaviors, hands, and tools across both space and time.

248 faculty members supported through HAI research grants to date

HAI Research Center

Center for Research on Foundation Models

Center for Research on Foundation Models (CRFM) celebrated its one-year anniversary in August 2022 and now counts more than 300 students, faculty, and postdocs working on relevant research projects.

This center was founded to drive fundamental advances in the responsible study, development, and deployment of foundation models. Foundation models are resource-intensive AI systems, such as GPT-3, Codex, StableDiffusion, Imagen, and AlphaFold, which are trained on broad data at scale. Foundation models provide valuable representations for many downstream applications; however, the technology is advancing at a rapid pace and there are unknown risks that must be considered.

Stanford established CRFM to focus on three pillars:

Research

The center conducts interdisciplinary research that lays groundwork for how foundation models should be built to make them more efficient, robust, interpretable, multimodal, and just.

Resources

The center builds and releases foundation models, code, and tools to improve the ecosystem.

Community

CRFM engages with universities, companies, policymakers, and civil society to develop professional norms that govern the responsible development and deployment of foundation models.

Research

During the past academic year, CRFM scholars authored dozens of papers, including:

The Time Is Now to Develop Community Norms for the Release of Foundation Models

In this paper, the authors propose setting up a review board to develop community norms and encourage coordination on release of foundation models for research access.

FlashAttention: Fast and Memory-Efficient Exact Attention with IO-Awareness

This innovation enables significant gains in speed for training foundation models and has been integrated into major frameworks, as well as adopted by many industry teams.

Large Language Models Can Be Strong Differentially Private Learners

Contrary to conventional wisdom that DP optimization fails at learning high-dimensional models, empirical results from this study revealed that private learning with pre-trained language models doesn't tend to suffer from dimension-dependent performance degradation.

Meet CoAuthor, an Experiment in Human-AI Collaborative Writing

A team of researchers studied how humans and AI can write together by designing large interaction datasets. This paper won an honorable mention at the 2022 ACM CHI Conference on Human Factors in Computing Systems.

Borrowing from the Law to Filter Training Data for Foundation Models

Using "Pile of Law," a dataset of legal materials, Stanford researchers explore filtering private or toxic content from training data for foundation models.

Notably, CRFM was awarded a 2022 Hoffman-Yee Research Grant to continue this important work.

Resources

To make progress on building concrete artifacts and tools, CRFM has hired its first few engineers. This is a critical step for building the infrastructure and expertise in academia to grapple with the scale and complexity of foundation models.

Community

To build community at Stanford, CRFM holds weekly meetings that bring together scholars from different departments, including computer science, law, political science, and biomedicine. To build norms and standards across the foundation model ecosystem, the center has put forth a proposal (featured as an op-ed in Protocol) on norms for the release of and access to foundation models. CRFM researchers engaged with policymakers at the HAI Congressional AI Bootcamp and are beginning to write a series of policy briefs together with HAI's policy team on foundation models, the evolving ecosystem, and norms/standards/benchmarks for ensuring the impact of this technology is in the interest of public good.

In collaboration with HAI's Industry Programs team, CRFM presented its work at key community engagement events with Google, IBM, and Wells Fargo.

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.

HAI Research Center

Stanford Digital Economy Lab

Over the last academic year, the Digital Economy Lab has been focused on conducting research and hosting major events. The lab currently has more than 50 research projects in its portfolio. Those projects fall into four areas of research:

Measuring the Digital Economy

Creating better methods of measuring the health of an increasingly digital economy

AI & The Future of Work

Understanding the future of the workforce in a rapidly changing global economy

Adoption of Advanced Technologies and Management Practices

Measuring, assessing, and predicting how data-driven decision-making and management practices are impacting the global economy and workforce

Digital Platforms and Society

Exploring how digital technologies can transform platforms and social media infrastructure to benefit society

Research Highlights

National Employment Report

The lab collaborated with the ADP Research Institute on a new version of the National Employment Report. This work represents an important step toward properly measuring labor and the economy using fine-grained and high-frequency data.

Digital Resilience: How Work-from-Home Feasibility Affects Firm Performance

In this exploration of remote-work, researchers provided a first look at how work-from-home practices improved resilience to a major, unanticipated social and economic shock. The team also completed a second wave of its survey with Gallup, which measures how Covid-19 and remote work affected organizations and workers.

Community Engagement and Public Outreach

To gather high-level feedback on its research areas, the lab hosted two advisory group meetings, which pointed the organization toward some key directions for growth and the expansion of its work. To build community at Stanford, the lab hosted two Stanford faculty roundtable discussions as well as a summer community BBQ in collaboration with HAI, the Golub Capital Social Impact Lab, and Stanford's RegLab.

The Digital Economy Lab is building community beyond Stanford, too. Two major hybrid events were hosted on campus, and a biweekly hybrid seminar during the academic school year featured top thinkers and academics from around the world.

Efforts like the *DigDig*, a semimonthly newsletter that captures both important Digital Economy Lab and HAI content, provide a way of reaching audiences outside the academic community. Finally, in social media, the lab's audience continues to grow primarily on Twitter, where it introduced a weekly update on important news from the digital economy.

The Turing Trap: The Promise & Peril of Human-like Artificial Intelligence

AI that complements or augments humans has a lot of economic potential, but right now the incentives for entrepreneurs, engineers, and governments are misaligned in favor of substituting human labor. In January 2022, Digital Economy Lab Director Erik Brynjolfsson published an essay outlining the social and economic dangers of disproportionately focusing on AI that automates human tasks. The essay explains that engineers have strong incentives to focus on automation, but for AI to create economic value in the long term, ideas that lead to augmentation are essential. It was widely covered in the media and has continued to gain traction and seed new research ideas.

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

AI Index

Measuring trends in artificial intelligence.

HAI and the AI Index Steering Committee published the 2022 AI Index report in March 2022. The latest edition included data from a broad set of academic, private, and nonprofit organizations as well as more self-collected data and original analysis than any previous edition. The authors included a variety of new data points to cover technical performance in reinforcement learning, the pricing of robotic arms, and artificial intelligence legislation.

The AI Index received positive press coverage, both around the time of the launch and throughout the rest of the academic year: highlights included pieces from publishers like *IEEE Spectrum, Fast Company* and *Morning Brew*. A number of prominent publications cited information contained in the report—*TechCrunch, Fortune, Forbes* and the *Harvard Business Review* among them.

HAI organized a total of 17 policy, public, and academic events to promote the findings in the 2022 report. Some of the most notable public events were:

- Virtual briefings at the HAI weekly seminar and a virtual presentation and fireside chat with a Chinese translation partner.
- Policy events included a trip to Washington, DC to brief policymakers and tech policy experts, briefings with Canadian Members of Parliament, and a general reception with the British Consulate General in San Francisco.
- For the academic audience, HAI presented to the World Economic Forum's AI Team, the Centre for Data Ethics and Innovation, and the Center for AI Governance.

The findings in the 2022 report also were used to inform policy-making in Canada and the European Union. Canadian parliamentarians consulted with the AI Index team on the current state of facial recognition, and data from the AI Index made it into the official report that the parliamentary committee released on its recommended regulation of facial recognition technology in Canada. Meanwhile, members of the European Commission consulted with the AI Index team to better understand the European Union's relative position in the international AI ecosystem.

Chess Software Engines: Elo Score

Progress in reinforcement learning, an important AI skill capability, can be captured by the performance of the world's top chess software engines. Computers surpassed human performance in chess a long time ago, and since then have not stopped improving.
Source: Swedish Computer Chess Association, 2021. Chart: 2022 AI Index Report

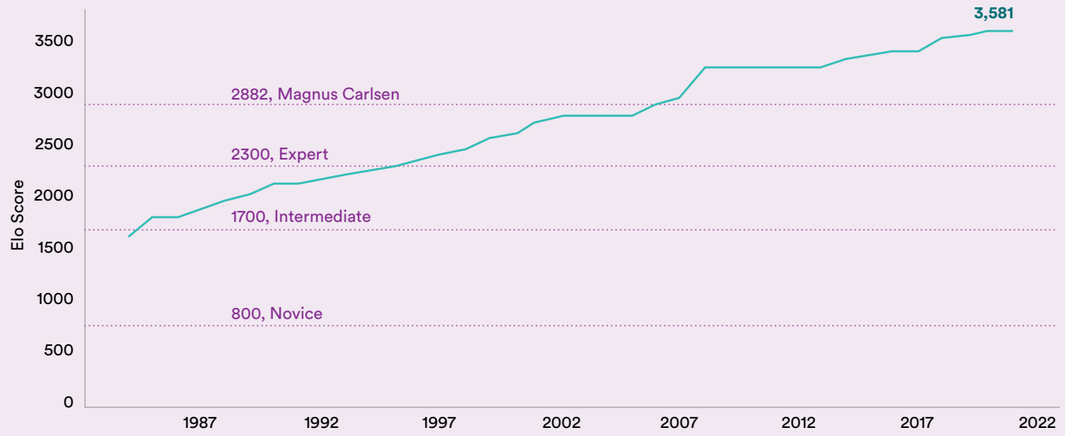

Research Events

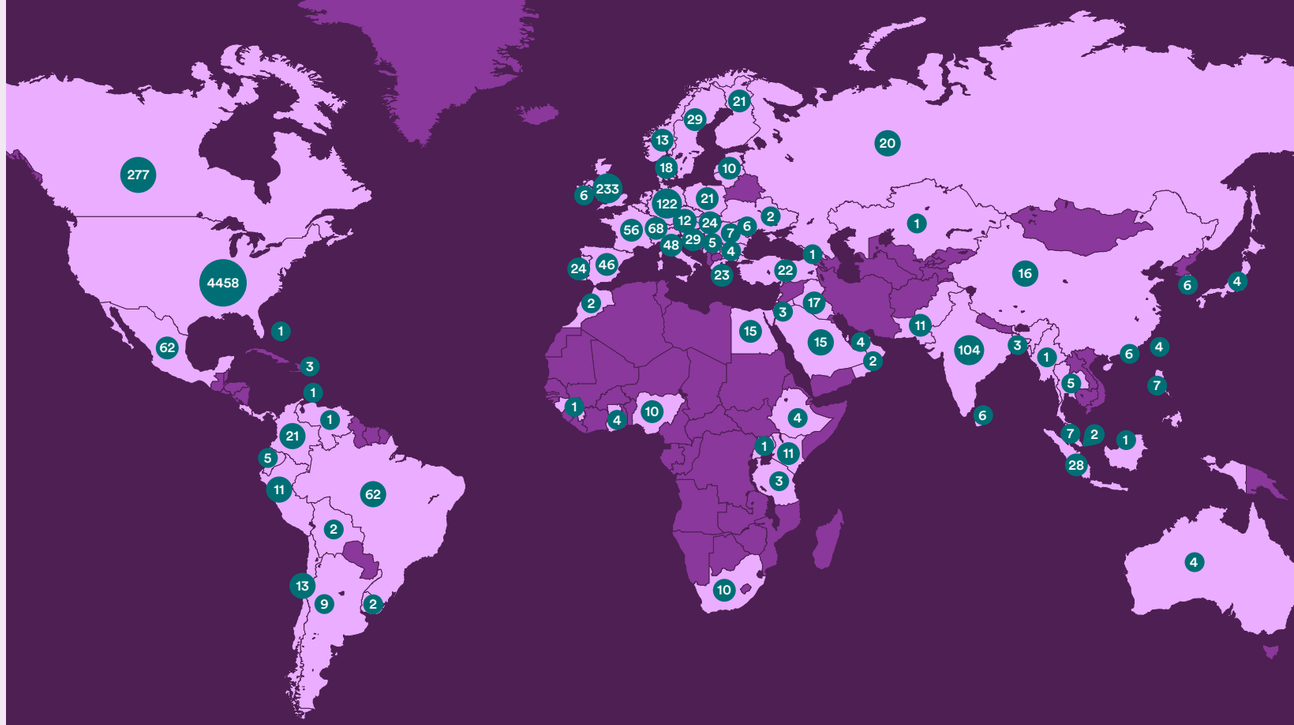

Weekly seminars

HAI's research team hosted 19 weekly seminars from Stanford and non-Stanford researchers across a wide range of topics, from "Decolonial AI" to "Embodied Intelligence." Seminars for the academic year were held as virtual live events averaging about 170 attendees. In all, the seminars reached more than 3,300 people from 90 countries.

Workshops

In November 2021, HAI partnered with ETH Zurich AI Center to host the Data-Centric AI Virtual Workshop. Sessions were hosted in both the U.S. and Europe. The recorded version of this event is one of HAI's most-viewed videos on YouTube. At the end of the workshop, HAI Research launched the Data-Centric AI Benchmarking Competition.

In February 2022, the HAI Collaboration Workshop provided a great example of community building for more than 40 Stanford faculty. Designed as an in-person workshop to support faculty in forming interdisciplinary research collaborations, the three-hour event included a combination of interactive activities and discussion.

In July 2022, HAI collaborated with the IEEE Standards Association and Stanford Woods Institute for the Environment to host Advancing Technology for a Sustainable Planet, a two-day event as part of IEEE's Planet Positive 2030 campaign. Panel sessions explored regulatory, policy, and financial frameworks critical to advancing technology that prioritizes people, planet, and purpose-driven progress. After the keynotes and panels, attendees were invited to contribute to a compendium document, "Strong Sustainability by Design," which seeks to address climate change issues with recommendations and practical approaches.

Policy

Engaging AI Leaders and Policymakers

Stanford HAI's policy team studies AI governance and appropriate uses, with a commitment to bringing AI leaders and policymakers together to learn about and discuss strategies for guiding human-centered AI. Activities include education and training for policymakers, as well as policy initiatives, publications, and events.

Education and Training for Policymakers

Promote National AI Research Resource (NAIRR)

In an effort to better educate policymakers and combat false narratives on the NAIRR, the policy team partnered with organizations on two events, one with Georgetown CSET and another with the Center for Data Innovation and SeedAI.

Artificial Intelligence and the Economy: Charting a Path for Responsible and Inclusive AI

In April 2022, HAI co-hosted a one-day symposium with the Department of Commerce, National Institute of Standards and Technology, and FinRegLab at the Commerce Department in Washington, D.C. The symposium focused on responsible AI integration in the economy, general risk, financial services, and healthcare.

HAI Congressional Staff Boot Camp

In August 2022, HAI welcomed a diverse cohort of 26 congressional staffers from both the House and Senate for a three-day educational event. Sessions covered what AI means for international security, the future of work, and healthcare. The event also included fireside chats with leading AI experts, networking dinners, and field trips to Stanford labs.

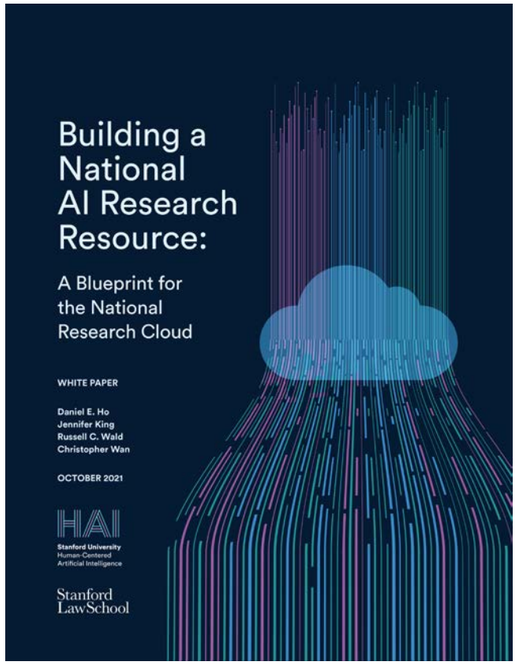

On April 27, former HAI Associate Director Susan Athey joined Deputy Secretary of Commerce Don Graves and Acting Comptroller of the Currency Michael J. Hsu in a fireside chat during a symposium co-hosted by HAI, the U.S. Dept. of Commerce, NIST, and FinRegLab.

From August 8-11, a bipartisan cohort of 26 congressional staffers from the House and the Senate joined our first Congressional Bootcamp on AI. They participated in sessions which explored what AI means for international security, future of work, healthcare, and more.

Policy Briefings in Washington, D.C.

The HAI policy team led Stanford faculty members to hold briefings for policymakers in D.C. on AI-related research at Stanford.

U.S. Federal Agencies

White House office of Science and Technology (OSTP)

Department of State (Bureau of Cyberspace and Digital Policy, Office of the Science and Technology Adviser to the U.S. Secretary of State, Center for Analytics)

General Accountability Office (GAO)

National Security Council (NSC)

Central Intelligence Agency (CIA)

National Science Foundation (NSF)

Think Tanks

Special Competitive Studies Project (SCSP)

Center for a New American Security (CNAS)

Georgetown Center for Security and Emerging Technology (CSET)

Center for Strategic and International Studies (CSIS)

Congressional Offices and Committees

Congressional Black Caucus

Congressional and Senate AI Caucuses

Senate and House Intelligence Committee

Senate Committee on Commerce, Science, and Transportation

House Select Committee on Economic Disparity and Fairness in Growth

House Committee on Science, Space, and Technology

Office of Senator Martin Heinrich

Office of Senator Mark Warner

Office of Senator Cory Booker

Office of Senator Alex Padilla

Office of Senator Rob Portman

Office of Senator Mike Rounds

Office of Representative Anna Eshoo

Office of Representative Ted Lieu

Office of Representative Jerry McNerney

Office of Representative Jay Obernolte

U.S. and E.U. Policymakers Visit Stanford

Throughout the academic year, HAI hosted a number of visits from high-level representatives to engage and inform government officials on AI policy opportunities and challenges.

Meetings included Senator Mark Warner, Representative Jerry McNerney, Representative Anna Eshoo, then-UK Chancellor of the Exchequer Rishi Sunak, eight members of the E.U. parliament with MEP Victor Negrescu, an Austrian government delegation, a Norwegian government delegation, an Estonian government delegation, and the Taipei Economic and Cultural Office (TECO).

These visits served to enhance the visibility of cutting-edge Stanford AI research, promote direct dialogues between the federal government and academia, and identify opportunities and gaps for collaboration in AI R&D, governance initiatives, and development of AI policy norms.

On May 25, HAI faculty affiliates and directors met with members of the European Parliament at Stanford University to discuss the Digital Revolution, privacy, surveillance, and other AI regulatory issues.

Policy Initiatives

Emerging Technology Policy Writing Competition

HAI worked together with the Digital Economy Lab and the Stanford Institute for Economic Policy Research (SIEPR) to seek submissions of policy briefs from Stanford graduate students. The competition aimed to cultivate succinct, effective, and evidence-based policy writing that can help policymakers map out a human-centric approach to responsible development and deployment of AI technologies.

AI Audit Challenge

In May 2022, HAI's policy team coordinated with Marietje Schaake, Rob Reich, and Jack Clark to launch an algorithmic auditing competition. HAI invited participants to design and create applied tools to assess whether deployed AI systems exhibit bias or other harm. With a grant from the Rockefeller Foundation, the competition had a total of $71,000 to offer the winning submissions. The application process was still underway at the end of AY 2021-22 with awards planned to be announced in early 2023.

Multilateral AI Research Institute (MAIRI)

HAI published this white paper to map a blueprint for a research institute that would provide a way to accelerate AI research and development collaboration among like-minded countries. The authors make the case for how an institute like this could help to establish better AI research standards worldwide for the responsible and inclusive development of AI.

Policy Publications

HAI's policy team translates cutting-edge, multidisciplinary AI research into digestible formats for policymakers. The team also produces original AI-related policy research to equip policymakers with tools to understand and govern this technology.

Policy Briefs

During the last academic year, we published five policy briefs:

Risks of AI Race Detection in the Medical System

Using Satellite Imagery to Understand and Promote Sustainable Development

Toward Stronger FDA Approval Standards for AI Medical Devices

A New Direction for Machine Learning in Criminal Law

Promoting Algorithmic Fairness in Clinical Risk Prediction

Policy Explainers

HAI's policy explainers summarize key points of AI-relevant legislation or other legal documents to help decision-makers find the information they need to make effective decisions. These three explainers were published in AY 2021-22:

EU AI Act

AI-related Provisions from the National Defense Authorization Act 2022

What The CHIPS and Science Act means for Artificial Intelligence

Federal Requests for Information

Federal Requests for Information (RFIs) synthesize academic research and perspectives to provide the federal government with necessary tools and knowledge to develop AI-related governance initiatives. In response to the White House Office of Science and Technology Policy, HAI prepared the following:

AI Bill of Rights for public and private sector use of biometric technologies

Recommendations on updating the National AI R&D Strategic Plan

Feedback to the National AI Research Resource Interim Report

Recommendations to the National Institute of Standards and Technology

To assist the National Institute of Standards and Technology, HAI prepared recommendations on how standards can be applied to a more productive digital economy.

Faculty Engagement

Healthcare Policy Steering Committee

A committee led by HAI Associate Director Curt Langlotz met three times during the academic year to build the community and discuss research agendas. The committee also funded three policy-related healthcare research grants.

International Security Steering Committee

Led by Hoover Institution Senior Fellow and Professor of Political Science (by courtesy), Amy Zegart, the committee commissioned a study on DOD use cases in AI at the non-combat level and supported participation in the newly launched Stanford Emerging Tech Review. It also held various engagements across campus.

Faculty-originated Policy Research

HAI's policy team is working with Professor of Law and Health Policy Michelle Mello on a comprehensive study of liability related to AI use in healthcare, funded by the Healthcare Policy Steering Committee.

Education

Educating the Leaders of Today and Tomorrow

HAI's growing portfolio of programs helps high-impact decision-makers and emerging leaders to pursue cutting-edge research, create useful and responsible AI, and respond thoughtfully to the societal and ethical implications of global AI.

Reaching Faculty at Stanford and Beyond

During AY 2021-22, HAI enrolled more than 7,000 people across its education programs, including 50 faculty from Stanford and other institutions who were trained in matters related to human-centered AI.

Special Initiatives

Stanford AI+Health Conference

The inaugural Stanford AI+Health Conference, held December 8-9, 2021, was the largest AI+Health Conference in Stanford's history in terms of audience size and faculty participation, with more than 150 Stanford faculty presenting. During this online event, Stanford HAI, the Artificial Intelligence in Medicine & Imaging (AIMI) center, and the Center for Continuing Medical Education (CME) convened experts and leaders from academia, industry, government, and clinical practice to explore emerging issues related to AI's impact on healthcare.

Stanford Art + Tech Salon

Stanford Art + Tech Salon is a quarterly get-together for affiliated artists, researchers, scholars, and students who work at the intersection of arts and tech to share work and engage in discussions. Events are produced through a partnership with Stanford Arts. AY 2021-2022 events included a spotlight on Jeremy Bailenson and Cyan DeVeaux's work creating classes and new worlds in VR; Lark Alder and Khuyen Le on Queer Technology and Speculative Design; and Ge Wang and Kunwoo Kim on Music, Computing and Design.

In the summer of 2022, 30+ high school students participated in a three-week AI4ALL program where they learned about AI through lectures, mentorship, and hands-on research projects.

For Current Leaders

Government Education

HAI now offers six distinct programs targeting more than 100 government executives and policymakers. In addition to the Congressional Boot Camp, four additional training programs provided opportunities for executives across federal government bodies in the U.S. and internationally to learn about the latest AI trends and issues.

Executive Education

5,000 executives and professionals enrolled in HAI's Executive and Professional Education offerings during the academic year. Highlights included the launch of an invitation-only CEO Roundtable Series and a new custom Executive Education program that gives organizations the option to choose among 100 AI topics. In December 2021, Journalist Day at HAI focused on AI Imperatives for 2022. The education team also provided faculty training on "Teaching AI and Ethics: Lessons from Stanford's Embedded EthiCS Team" for the Stanford Educational Partnership for Internationalizing Curriculum (EPIC) Community College Faculty Fellowship.

For Future Leaders

A group of high school students attended a three-week live virtual program called Stanford AI4ALL to get immersed in AI through a combination of lectures, hands-on research projects, and mentoring activities. HAI also offered two courses in Stanford's Pre-collegiate Summer Institutes: Artificial Intelligence and Fundamentals of Augmented Reality.

In the spring, HAI's "Visions of the Future" online course reached 500 global students, covering the future impact of AI, as well as AI and the metaverse, healthcare, and blockchain.

Denning Co-Directors

HAI's co-directors lead the institute much like an academic startup, making sure that all of the institute's work reflects a passion for the safe and responsible development of AI technology.

John Etchemendy

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

In AY 2021-22, John served as chair of the Hoffman-Yee Grant Selection Committee and was named to the AI Index Steering Committee. He also taught "Philosophy of Artificial Intelligence" to undergraduate students. Notably, John received the K. Jon Barwise Prize from the American Philosophical Association, honoring "significant and sustained contributions to the intersection of philosophy and computation."

Fei-Fei Li

Sequoia Professor, Computer Science Department

Fei-Fei co-led the HAI Spring Conference on Key Advances in Artificial Intelligence and spoke at the Stanford HAI Congressional Boot Camp on Artificial Intelligence in AY 2021-22. She was named to the The Center for Strategic and International Studies (CSIS) AI Council, and she published several papers: "Embodied intelligence via learning and evolution" in Nature Communications; "Advances, challenges and opportunities in creating data for trustworthy AI" in *Nature Machine Intelligence*; and "Searching for Computer Vision North Stars" (co-author) in *Dædalus*. Fei-Fei received the Thomas S. Huang Memorial Prize which honors scientists who are recognized as examples in research, teaching/mentoring and service to the computer vision community.

Associate Directors

HAI's associate directors bring a diverse mix of backgrounds and experience to the institute. They are essential to the execution of HAI's multidisciplinary approach to studying AI.

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 co-hosted the HAI Fall Conference with James Landay and co-chaired the search for HAI Junior Fellows with John Etchemendy, which resulted in the hiring of Shriti Raj, HAI's newest Junior Faculty Fellow (July 2023). Russ also served as Faculty Director for AI100, spoke at the School of Engineering event on Principled Engineering: AI and Drug Development, and worked with Curt Langlotz to strengthen ties between HAI and the School of Medicine.

Susan Athey

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

Susan published "Uncovering interpretable potential confounders in electronic medical records" in MedRxiv and spoke at the HAI joint event with IEEE, "Advancing Technology for a Sustainable Future." In July, she moved on from HAI to accept a position as chief economist of the antitrust division at the U.S. Department of Justice.

Michele Elam

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

Michele presented on AI & the Future of the Creative Industry at the Stanford HAI Congressional Boot Camp on Artificial Intelligence, and with Russell Wald, briefed congressional staffers on art, equity, and generative AI. She authored four peer-reviewed articles placed in top journals in the humanities as well as "The Stanford HAI Bill of Rights," co-led with Rob Reich. She taught two interdisciplinary classes on human-centered AI that examine race, gender, disability & equity, and gave keynote talks on human-centered AI at the University of Chicago, Ohio State, UCLA, Princeton, and Cambridge (UK). Her HAI committee work included Hoffman-Yee and Junior Fellows selection, and faculty lead for the HAI Visiting Artist selection committee. Michele is the faculty lead for the Arts & Tech Salon and represented HAI on the Race & Tech Initiative (co-sponsored by HAI and Center for Comparative Studies in Race & Ethnicity).

Surya Ganguli

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

Surya published "Embodied intelligence via learning and evolution" in *Nature Communications* and co-led an HAI seminar on this work. He also published "Emergent reliability in sensory cortical coding and inter-area communication" in Nature and gave a lecture on Applied Physics/Physics Colloquium.

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 was appointed to the National AI Advisory Commission to advise the White House on AI policy, as well as to the Committee on National Statistics at the National Academies. He served as co-lead for the HAI Fall Conference on Policy & AI: Four Radical Proposals for a Better Society. Daniel spoke about Public Sector AI at the Artificial Intelligence and the Economy event in Washington, D.C., which was co-hosted with the Department of Commerce and FinRegLab. He also spoke at the Stanford HAI Congressional Boot Camp on Artificial Intelligence. Additionally, he taught "Antidiscrimination Law and Algorithmic Bias," and co-led a HAI seminar on "A Blueprint for the National Research Cloud." His publications included "Advances, challenges and opportunities in creating data for trustworthy AI" in *Nature Machine Intelligence*.

James Landay

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

James co-hosted the HAI Fall Conference with Russ Altman on "AI in the Loop: Humans in Charge." He was named Vice Director and Faculty Director of Research for HAI, and he continued to lead HAI's grant programs and to co-lead the Hoffman-Yee Symposium. James gave a keynote on "'AI for Good' Isn't Good Enough: A Call for Human-Centered AI" at MobiCom 2022 in Sydney, The Australian National University, Brown University, and Northwestern University. In November 2022, he published "End-User Audits: A System Empowering Communities to Lead Large-Scale Investigations of Harmful Algorithmic Behavior," along with Michelle S. Lam, Mitchell L. Gordon, Danaë Metaxa, Jeffrey T. Hancock, and Michael S. Bernstein, in Proc. ACM Hum.-Comput. Interact. 6, CSCW2, Article 512.

Curtis P. Langlotz

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

Curtis chaired the HAI Steering Committee on Healthcare Policy, which initiated studies of tort liability and medical AI and dangers of health AI dual use. He published "Implementation of Clinical Artificial Intelligence in Radiology: Who Decides and How?" which proposed a risk model and governance structure for the implementation of AI in health imaging. He also led the AIMI-HAI partnership grant program, which provided over $1 million in new and ambitious ideas that reimagine AI in healthcare, using real clinical datasets, with near-term clinical applications. Curtis also helped organize the Stanford AI+Health conference, which explored critical and emerging issues related to AI's impact across the spectrum of healthcare.

Christopher Manning

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

Chris co-led the HAI Spring Conference on Key Advances in Artificial Intelligence. He also leads the HAI cloud credit grant programs. He published papers including "Human language understanding and reasoning" in Dædalus and "Fast model editing at scale" and "GreaseLM: Graph REASoning Enhanced Language Models" at ICLR 2022.

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 published a paperback version of *System Error: Where Big Tech Went Wrong and How We Can Reboot,* with Mehran Sahami and Jeremy M. Weinstein, with Brazilian, Korean, and Japanese translations also released in 2022. He spoke at the HAI Spring Conference on Key Advances in Artificial Intelligence and co-authored a HAI blog post, "The Time Is Now to Develop Community Norms for the Release of Foundation Models." He continued to guide development of the Embedded EthiCS initiative.

Staff

HAI depends on an exceptionally capable staff to realize its mission and run programs that have wide appeal and global impact.

Krysten Hommel
Senior Associate Director, Strategic Initiatives (through May 2022)

Amita Kumar
Director of Administration

Panos Madamopoulos-Moraris
Managing Director of Industry Programs and Partnerships

Maura McGinnity
Senior Director of Development, Strategic Initiatives

Vanessa Parli
Associate Director of Research Programs

Stacy Peña
Director of Communications

John Robichaux
Director of Education

Russell Wald
Managing Director for Policy & Society

Michael Sellitto

Deputy Director (through June 2022)

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. In June 2022, Mike joined AI startup Anthropic as Head of Geopolitics and Security. 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).

Faculty and Fellows

HAI currently has 250 affiliated Stanford faculty from across all seven schools. Each year, the institute aims to support the work of researchers who focus on the intersection of multiple disciplines. In AY 2021-22, HAI launched two new graduate fellowships: The HAI Graduate Fellowship Program, led by HAI Junior Faculty Fellow and Assistant Professor in the Psychology Department, Johannes Eichstaedt, and the Technology and Racial Equity Graduate Fellowship offered in partnership with the Center for Comparative Studies in Race and Ethnicity (CCSRE).

Featured Fellows

Kathleen Creel

Embedded EthiCS Fellow

In her second year of a two-year fellowship, Kathleen Creel seeks to understand ethical questions posed by emerging technologies with an interdisciplinary approach. She teaches in and creates course material for a program designed to embed ethical considerations into the computer science curriculum. She's also pursuing her own research, collaborating with interdisciplinary scholars from HAI's Center for Research on Foundation Models (CRFM) to study the negative impact of algorithmic monoculture on individuals.

Jazmia Henry

HAI-CCSRE Practitioner Fellow

Machine learning specialist Jazmia Henry focuses on finding a way to responsibly incorporate African American Vernacular English (AAVE) into natural language processing models. As a fellow at HAI and the Center for Comparative Studies in Race and Ethnicity, she has created an open-source corpora of more than 141,000 AAVE words to help researchers and builders design models that are both inclusive and less susceptible to bias.

Peter Norvig

Distinguished Education Fellow

AI expert Peter Norvig has joined Stanford HAI as a Distinguished Education Fellow to help develop a human-focused AI curriculum and broaden access to education. A longtime innovator in the field, he co-wrote *Artificial Intelligence: A Modern Approach*, an introductory textbook used by 1,500 universities worldwide. He has taught hundreds of thousands of students through his courses on the online education platform Udacity.

Junior Faculty Fellows

Hariharan Subramonyam

HAI Junior Faculty Fellow

Hari Subramonyam is an Assistant Professor at the Graduate School of Education and Computer Science (by courtesy). He is a Ram and Vijay Shriram Faculty Fellow at Stanford HAI. Hari's research looks at enabling multidisciplinary teams to design and develop ethical, responsible, and human-centered experiences with AI. A specific application area of his research is AI in education, where he draws from cognitive psychology principles to augment human learning, creativity, and visual sensemaking.

Johannes Eichstaedt

HAI Junior Faculty Fellow

Johannes is a computational social scientist jointly appointed as a Ram and Vijay Shriram HAI Faculty Fellow and Assistant Professor (Research) of Psychology. He analyzes large social text (Twitter, Facebook, Reddit, etc.) with NLP and AI to understand and measure psychological processes for individuals and populations—such as building early warning systems for mental health. Currently, Johannes is focused on leveraging generative language models (like ChatGPT) to deliver high-fidelity therapy and positive interventions (like gratitude exercises).

Content

HAI's content team published 97 blog posts, 9 announcement stories, and 61 videos, accumulating 552,000 pageviews and 169,700 views on YouTube. The team increased social media following by 42% during the academic year and set up a syndication partnership with VentureBeat. Content highlights from the year included:

- A healthcare AI series covering explainability, generalizability, privacy and fairness concerns, and whether AI for health is even useful
- A fall conference series examining crucial ideas around data intermediaries, UBI, middleware, and AI audits
- The People of HAI series, highlighting the personalities and ambitions of the people affiliated with HAI
- A video series on this year's Hoffman-Yee Grant recipients and their research goals

2021–22 HAI Fellows

Diana Acosta Navas
Embedded EthiCS Fellow

Sarah Bana
DEL Postdoctoral Fellow

Kathleen Creel
Embedded EthiCS Fellow

Ruyu Chen
DEL Postdoctoral Fellow

Jeffrey Ding
HAI-CISAC Postdoctoral Fellow

Vael Gates
HAI-CISAC Postdoctoral Fellow

Jazmia Henry
HAI-CCSRE Practitioner Fellow

Todd Karhu
HAI-EIS Postdoctoral Fellow

Jennifer King
HAI Privacy and Data Policy Fellow

Jae Joon Lee
DEL Postdoctoral Fellow

J. Frank Li
DEL Postdoctoral Fellow

Evelyn Mei
HAI-CCSRE Practitioner Fellow

Collier Navaroli
HAI-CCSRE Practitioner Fellow

Rashaad Newsome
Visiting Artist

Peter Norvig
HAI Distinguished Education Fellow

Marietje Schaake
HAI International Policy Fellow

Faye-Marie Vassell
STEM Education, Equity, and Inclusion Fellow

Research Centers and Partners

HAI's work depends on close collaboration with many labs, centers, and institutes across the Stanford campus and beyond.

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

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

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.

Digital Economy Lab

The Stanford Digital Economy Lab at HAI 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

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

Black in AI Research (BlackAIR)

Center for AI in Medicine & Imaging (AIMI)

Center for Comparative Studies in Race and Ethnicity (CCSRE)

McCoy Family Center for Ethics in Society

Center for International Security and Cooperation (CISAC)

Cyber Policy Center

Digital Civil Society Lab (DCSL)

Hoover Institution

Institute for Diversity in the Arts

John S. Knight Journalism Fellowships (JSK)

Office of the Vice President for the Arts

Stanford Aging and Ethnogeriatrics (SAGE) Research Center

Stanford Institute for Economic Policy Research (SIEPR)

Wu Tsai Neurosciences Institute

External Partners

AI4All

Corporate Members Program

For its community of corporate members, HAI published two corporate industry briefs on Robotics & AI and Education & AI. The Founding Members Program leveraged research tokens, Fellow-Mentor-Advisor relationships, and visiting scholars to facilitate research collaborations and distribute almost $1.5 million to faculty and programs across the university. Funded projects are addressing foundation models, AI trust, social computing, conversational AI, and human-computer/human-robot interaction. In addition, HAI developed a tailored AI Seminar Series for corporate members.

The HAI Corporate Affiliate Program launched a second program area on Consumer Goods, Retail & AI. This effort aims to provide an environment for manufacturers, suppliers, distributors, retail chains and consumer goods and retail corporations to engage with Stanford at the nexus of their most fundamental AI challenges. The HAI Corporate Program also includes a general program area for companies with a more horizontal interest in AI.

"The collaboration with HAI has evolved into deeper engagement over the course of the past year. Our researchers interact regularly with faculty and graduate students on research topics of mutual interest, sharing research outcomes through joint publications and open-source code contributions. This is a natural transition into the kind of deep joint collaboration that we envisioned would happen as HAI grew out of the founding stages into steady-state operation.

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

"Google recognizes that HAI's interdisciplinary perspective is key to shaping an inclusive AI future. 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 as we continue our work with HAI.

Jeff Dean, Google Senior Fellow and SVP of Google Research

"HAI continues to be a critical engagement for Wells Fargo's innovation plans, specifically in the AI/ML space. Its cross-discipline approach and resulting access to deep subject matter expertise approaching the problems from different perspectives is the only way we will solve hard problems of explainability and bias for what will be an increasingly complex ecosystem of AI.

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

Funding

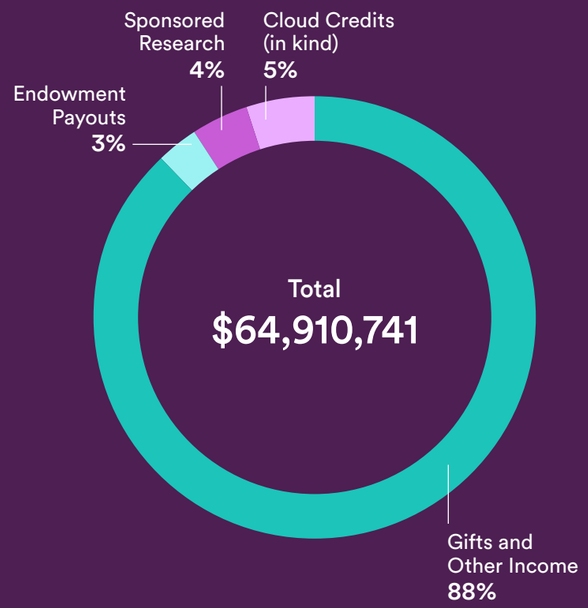

In FY22, HAI issued **$11.1M** in research grants.

Sources of Income[1]

Gifts and Other Income	$56,842,049	88%
Endowment Payouts	$2,280,196	3%
Sponsored Research	$2,616,371	4%
Cloud Credits (in kind)	$3,172,125	5%
Total		$64,910,741

HAI Research Support

Cloud Credit Grants	$3,172,125
Hoffman-Yee Grants	$6,000,000
Seed Grants	$1,960,248
Total	$11,132,373

1 Income does not include increases in endowment principal

Financials

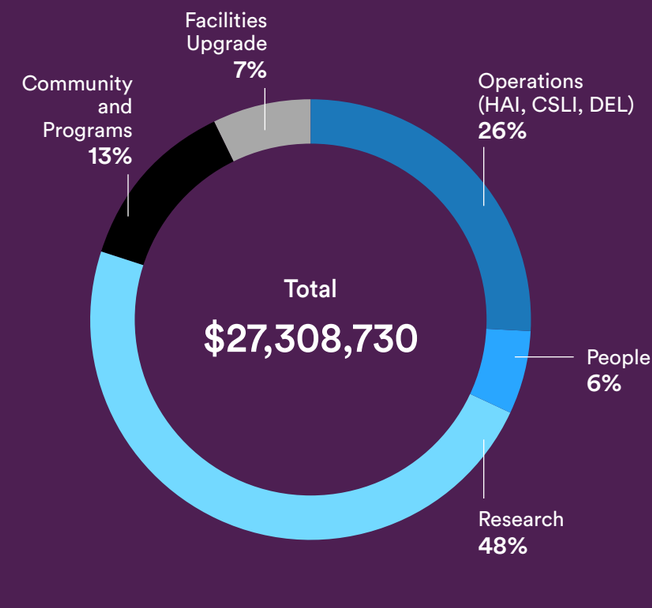

Operations	$7,563,459	26%
People	$1,861,892	6%
Research	$14,001,153	48%
Community and Programs	$3,882,226	13%
Facilities Upgrade	$2,000,000	7%
Total		$27,308,730

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 is grateful for the generosity of those who provide support to make its work possible. The institute welcomes gifts made online.

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.stanford.edu

Reach out
hai-institute@stanford.edu

Subscribe to our
newsletter

Follow us

Stanford University
Human-Centered
Artificial Intelligence

Stanford
University