HAI | **Stanford University
Human-Centered
Artificial Intelligence**
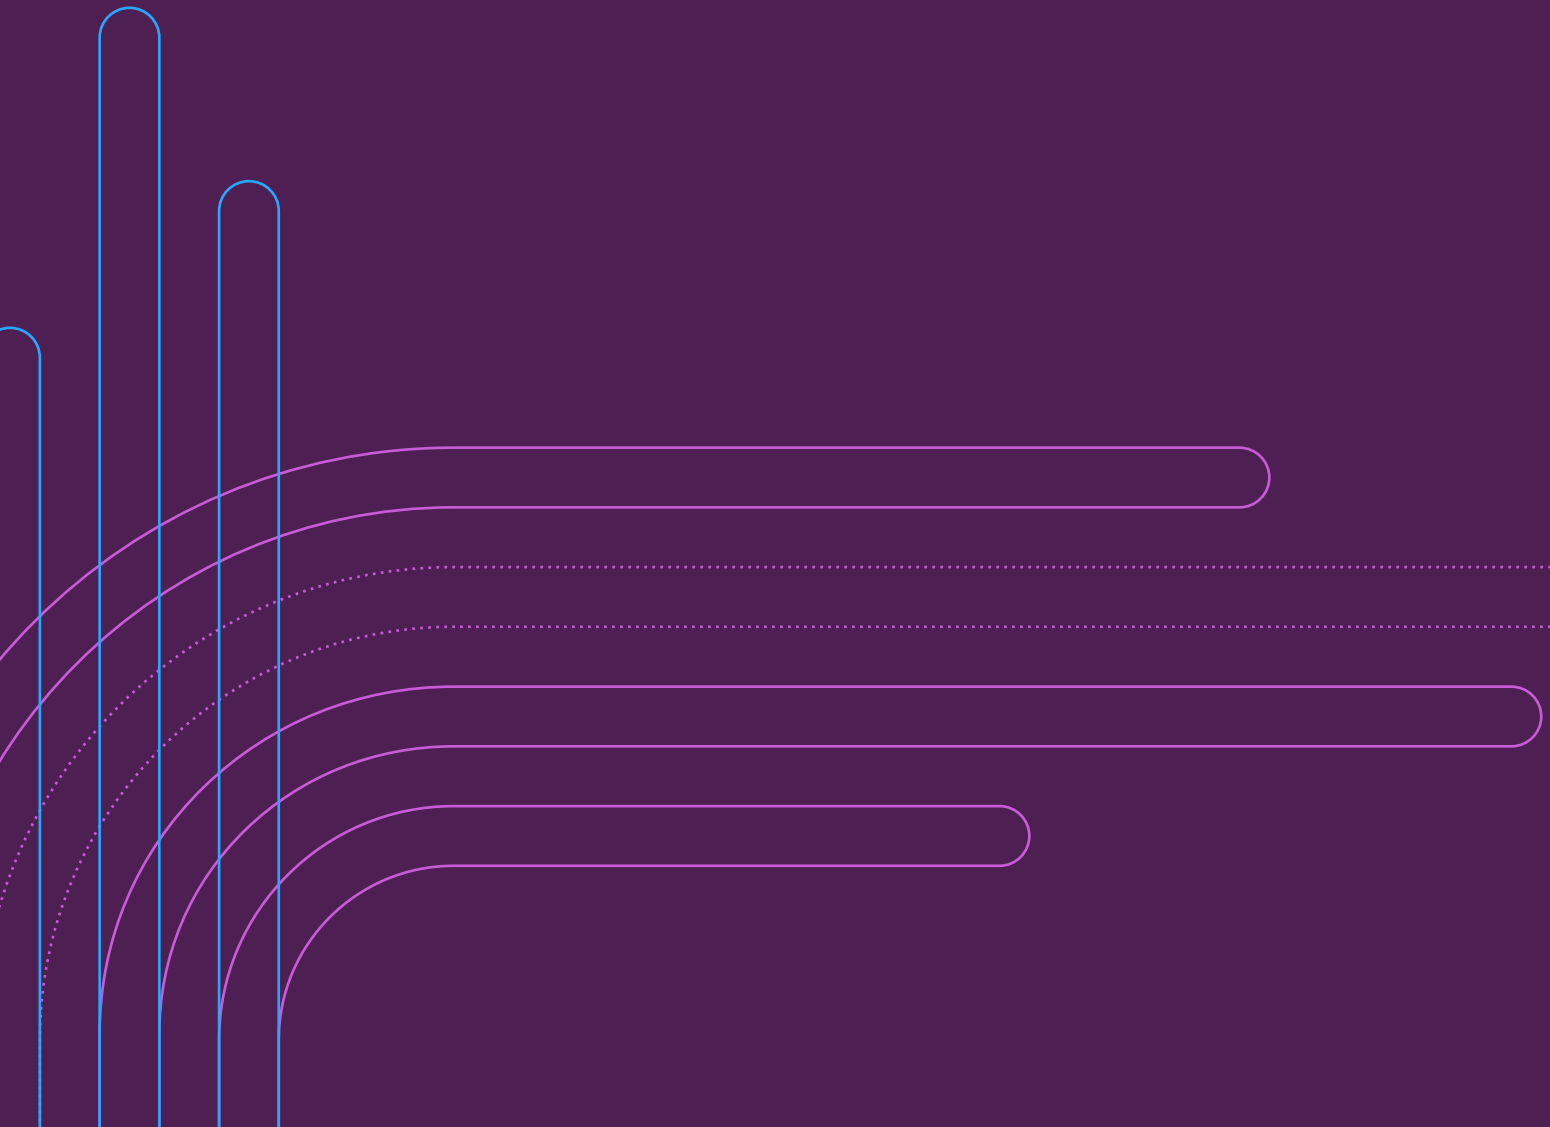
Stanford
University

Annual Report 2023

Stanford Institute for Human-Centered Artificial Intelligence

Letter from the HAI Leadership Team

In a year of unprecedented AI innovation and growth, Stanford HAI was well positioned to shape the increasingly urgent conversation about the need for a human-centered approach to AI. While the world pondered the implications of generative models like Stable Diffusion and ChatGPT, we drew on many years of prior thinking and interdisciplinary scholarship to further establish the institute as a leading voice in guiding the development of fair and ethical AI across society.

Contents

During the Academic Year 2022-23, HAI took great strides forward in our primary focus areas of Research & Education and Policy & Society to guide and build the future of AI as it applies to medicine, government, education, and many other aspects of daily life. Our network of Stanford-funded faculty continued to grow, with more than $40 million distributed to date through our grant programs. Research & Education highlights from the past year include:

— A new holistic framework for evaluating large language models across a wide range of criteria, developed by HAI's Center for Research on Foundation Models

— The latest edition of the AI Index Report, which garnered widespread media coverage and demonstrated significant growth in global AI legislation (from 25 countries in 2022 to 127 countries in 2023)

— Support for various new Student Affinity Groups and Graduate Fellows to help develop young AI scholars

— Involvement in a range of theme-specific research communities across Stanford, including in the areas of safety, healthcare, and education, to promote the importance of a human-centered AI approach

— Through our cloud credit program, scholars launched DSPy, an open-source project that enables researchers to build powerful systems with small, inexpensive language models (LMs).

In the realms of Policy & Society, we experienced a surge in demand among global policymakers for our educational programs and briefings on critical AI topics. Our breakthrough accomplishments include:

— The success of our leadership efforts with top computer science universities and a bipartisan, bicameral group of lawmakers to bring the National AI Research Resource (NAIRR) to fruition, as well as the pivotal role we played in advancing the CREATE AI Act

— Fei-Fei and Rob Reich holding a meeting with President Biden in June 2023 to encourage him to adopt a moonshot mentality when thinking about AI

— The ability to help shape EU perspectives on the EU AI Act through our technical research on foundation models

— The launch of the Tech Ethics & Policy summer fellowship program, with notable placements in the White House, Senate Committee on Homeland Security and Governmental Affairs, American Enterprise Institute, USAID, and Brookings Institution

We're excited to announce several new collaborations on campus and beyond. For example, HAI has partnered with Stanford Medicine on an initiative called Responsible AI for Safe and Equitable Health (RAISE-Health), which seeks to address critical ethical and safety issues in AI for health and medicine. We're also working with the Cyber Policy Center (CPC) on convenings and programs to address AI in the Global South.

HAI's relationship with industry partners is evolving, too. From finance and retail to food and beverage, the institute has begun to shape how industries think about AI and to play a critical role in ensuring that industry implements AI in a human-centered, responsible way. In the last academic year, we welcomed five new corporate members into our affiliate program: McKinsey, SCBx, PepsiCo, LVMH, and SAP.

Within the institute, we've expanded the HAI team to include three new faculty Associate Directors: Julian Nyarko from the Stanford Law School brings experience applying machine learning to explore contract design under U.S. and international law; Amy Zegart joins us from the Hoover Institution, where she directs the Robert and Marion Oster National Security Affairs Fellows program; and Ge Wang, Professor of Music, specializes in the art of design and computer

music. HAI is committed to growing our intellectual capital by appointing additional Senior Fellows.

As we look ahead to our fifth anniversary and the unrelenting pace of AI development, we remain true to our mission of advancing AI research, education, policy and practice to improve the human condition. We appreciate the ongoing support and encouragement of the extended HAI community, including our faculty, affiliates, students, donors, advisory council members, and above all, our staff. We realize that we could not do this important work without all of you.

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

Fei-Fei Li
Sequoia Professor, Computer Science Department
Denning Co-Director, Stanford HAI

James Landay
Anand Rajaraman and Venky Harinarayan Professor, School of Engineering, Professor of Computer Science
Denning Co-Director and Faculty Director of Research, Stanford HAI

Russell Wald
Deputy Director, Stanford HAI

HAI Research Focus

All of HAI's research initiatives are tied to at least one of three imperatives that were identified at the institute's founding: understanding the human and societal impact of AI, augmenting human capabilities, and developing AI technologies inspired by human intelligence.
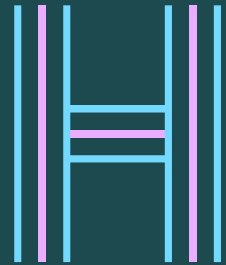
Human Impact

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

Research efforts that augment human capabilities have the potential to educate, train, and support individuals in ways that lead to progress in health care, education, sustainability, automation, and other domains.

Intelligence

Intelligence initiatives aim to develop AI that understands human language, emotions, inventions, behaviors, and interactions at multiple scales.

Spotlights

Hoffman-Yee Symposium

The Hoffman-Yee Grant Symposium highlighted the work of six research teams that received funding in 2022 to solve some of the most challenging problems in AI. Each team presented its results to date and competed for additional funding of up to $2 million over the next two years. Grant recipients from 2021 also presented their work to date and plans for the future.

RAISE-Health: Responsible AI for Safe and Equitable Health

In response to rapid advances in AI for health and medicine, HAI collaborated with Stanford Medicine to launch a new initiative that seeks to address critical ethical and safety issues. Co-led by Stanford School of Medicine Dean Lloyd Minor and Fei-Fei Li, RAISE-Health has articulated three goals:

— Enhance clinical care outcomes through responsible integration of AI

— Accelerate research to solve the biggest challenges in health and medicine

— Educate patients, care providers, and researchers to navigate AI advances

HELM: Holistic Evaluation of Language Models

In an effort to improve transparency across the commercial AI landscape, HAI worked with the Center for Research on Foundation Models (CRFM) to introduce a holistic framework for evaluating large language models. This new framework covers a broad range of use cases and evaluates applications across a comprehensive set of relevant metrics such as fairness, efficiency, robustness, and toxicity.

HAI Conferences

Each year HAI hosts an annual conference on a relevant or emerging AI topic, in order to catalyze new interdisciplinary research collaborations, discussions, and conversations. By convening leading experts from computer science, engineering, ethics, law, and the social sciences, the institute takes a holistic look at how AI is transforming our world. These conferences allow researchers who may not normally interact to exchange ideas, identify shared interests, and form new partnerships.

In November 2022, "AI in the Loop: Humans in Charge" challenged participants to rethink the phrase "human-in-the-loop" and consider a future where humans remain at the center of all AI technologies. HAI's 2023 Spring Symposium, titled "Creativity in the Age of AI" and supported by the Doris Duke Foundation, focused on the impact of these technologies on the creative field.

The 2022 Fall Conference, organized by James Landay and Russ Altman on the topic of "AI in the Loop, Humans in Charge." Photos: Christine Baker

Research & Education

Led by Vanessa Parli, Director of Research, HAI continued to seek out and fund projects with the potential to develop AI that is collaborative, augmentative, and enhancing to human productivity and quality of life.

Education

HAI is committed to fostering a diverse range of voices, which includes an intentional investment in the next generation of AI thinkers. HAI is developing the following programs to strengthen the young scholar AI community.

Student Affinity Groups

In January 2023, HAI launched a pilot student affinity group program that invites students from any of Stanford's seven schools to identify a topic of interest to pursue and recruit their own small group of interdisciplinary teammates. Projects run for the academic year and are eligible for funding to cover basic expenses. Results from the inaugural six teams included a comedian-moderated AI panel, an interview series profiling individuals working at the intersection of AI and creative expression, and a position paper on intelligent writing agents, which was accepted into In2Writing, a workshop on intelligent writing assistants.

Graduate Fellows

HAI offers a three-quarter fellowship program for Stanford graduate students to encourage interdisciplinary research conversations, facilitate new collaborations, and grow the HAI community of graduate scholars who are working in the area of AI. The program fosters collaboration among engineers, social scientists, humanists, and others researching the future of purposeful, intentional, and human-centered AI.

AI Literacy in K-12 Schools

HAI's K-12 programs focus on enhancing AI literacy for educators and students from diverse backgrounds. For example, the AI4ALL summer camp immerses underrepresented students in hands-on AI research and mentoring at Stanford to increase diversity in the field's future leaders. Together, these K-12 initiatives equip teachers and youth with the knowledge, skills, and access needed to shape AI innovations in ways that reflect community needs.

Isabelle Levent and Lila Shroff, participants in this year's HAI Student Affinity Group pilot project. Photo: Jeanina Casusi

Human-Centered AI Course by Peter Norvig

In its second year, CS139: Human-Centered AI emphasized the intersection of AI development and human impact, delving into how advanced tools must be designed with consideration for their effects on users and communities. Open to graduate and undergraduate students, the curriculum aims to foster a thorough understanding of the importance of safety and positive outcomes in the field of AI. The course was recorded and is intended to become a signature HAI online course for professionals.

Executive Education

Offered in-person and virtually, HAI's executive education programs are designed to cultivate synergy and strengthen connections between academic endeavors and industry in alignment with HAI's mission. These programs provide a platform for faculty, postdoctoral researchers, graduate students, and other participants to engage with real-world use cases. The aim is to not only test but also develop human-centered AI applications, fostering a collaborative environment that aligns academic expertise with practical, industry-driven challenges. In 2022-23, HAI delivered executive education programs to a diverse set of organizations, including Barclays London, Accenture, Wells Fargo, and SCBx.

Course Highlight

HAI's inaugural on-campus executive education program, "Executive Upskilling for HR," was conducted for Barclays London over three consecutive days. Attendees included Barclays Executive Committee in HR, totaling 15 executives. Topics ranged from human-centered AI, AI for financial services, data privacy, and security to structuring AI Initiatives.

Seminar Series

Each academic year, HAI organizes research seminars to bring the AI community together around common interests. In AY 2022-23, HAI hosted 15 seminars featuring scholars from Stanford and beyond across a wide range of topics, from collaborative robots to ethical AI and health care AI applications to embodied intelligence. This year's seminars included computer scientist and entrepreneur Yoav Shoham's discussion around "Understanding, Understanding," and HAI Visiting Artist Lauren Lee McCarthy's innovative work titled "I Heard Talking Was Dangerous." On average, about 300 attendees joined for each seminar (virtually and in-person) representing over 100 countries throughout the year.

Theme-specific Research Communities

HAI supported the following activities to grow theme-specific interdisciplinary research communities across campus and beyond:

AI+Safety

To foster human-centeredness and values of safety, trust, and responsibility when developing AI methods and solutions, HAI continues to partner with labs and centers across campus. This work includes the 2023 Stanford Data Science Conference, for which HAI judged and sponsored an award to Xiao Ge for her work on How Emotion Shapes Creative Collaborative Work; Stanford's Trust and Safety Research Conference 2023, where HAI organized and moderated a panel on Mental Health and AI; and Stanford's Center for AI Safety, for which HAI supported many collaborative activities including the center's Annual Meeting.

AI+Health Care

HAI co-organized the second AI+Health Conference, in collaboration with the Center for Artificial Intelligence in Medicine and Imaging (AIMI), and the Center for Continuing Medical Education (CME). The December 2022 event brought together leading experts from academia, industry, government, and clinical practice to explore critical and emerging issues related to AI's impact across the spectrum of health care.

AI+Education

HAI partnered with the Stanford Accelerator for Learning to organize the AI+Education Summit: AI in the Service of Teaching and Learning in February 2023. The summit showcased the latest research in AI and education from Stanford faculty and researchers, offering networking opportunities with industry partners, VCs, and government and civic society organizations. Given ChatGPT was just released, the summit asked how AI can transform teaching and learning in an ethical, equitable, and safe manner. The conference offered a platform to inspire structured roundtable discussions and team formation toward HAI/GSE seed funding, where HAI contributed $225K to fund 10 faculty-lead projects. Other conference outcomes included a Seminar on Generative AI and Education, a student-led paper reading group, and student affinity groups.

State of AI Reports

AI100

The One Hundred Year Study of Artificial Intelligence (AI100) is a 100-year effort to study and anticipate how AI will affect every part of human life. The study is administered by HAI and managed by a standing committee of AI leaders from institutions around the world, chaired by Vincent Conitzer at Carnegie Mellon University. AI100 releases reports every five years with the goal of documenting AI developments and their ethical implications for society, as well as providing a through line that marks the trajectory of the field.

In 2023, AI100 launched an essay competition to hear directly from the next generation of AI scholars as a way of laying the groundwork for the next report scheduled to be released in 2026. After reviewing 54 submissions, the AI100 Standing Committee selected University of Texas at Austin Assistant Professor Samantha Shorey's essay as the winner. Shorey's essay advocates for greater attention to the labor of AI integration performed by essential workers. Shorey will serve on the next study panel tasked with writing the 2026 AI100 Report. In addition to the winner, AI100 released a collection of five essays that thoughtfully consider AI at the intersection of morality, regulation, love, labor, and religion.
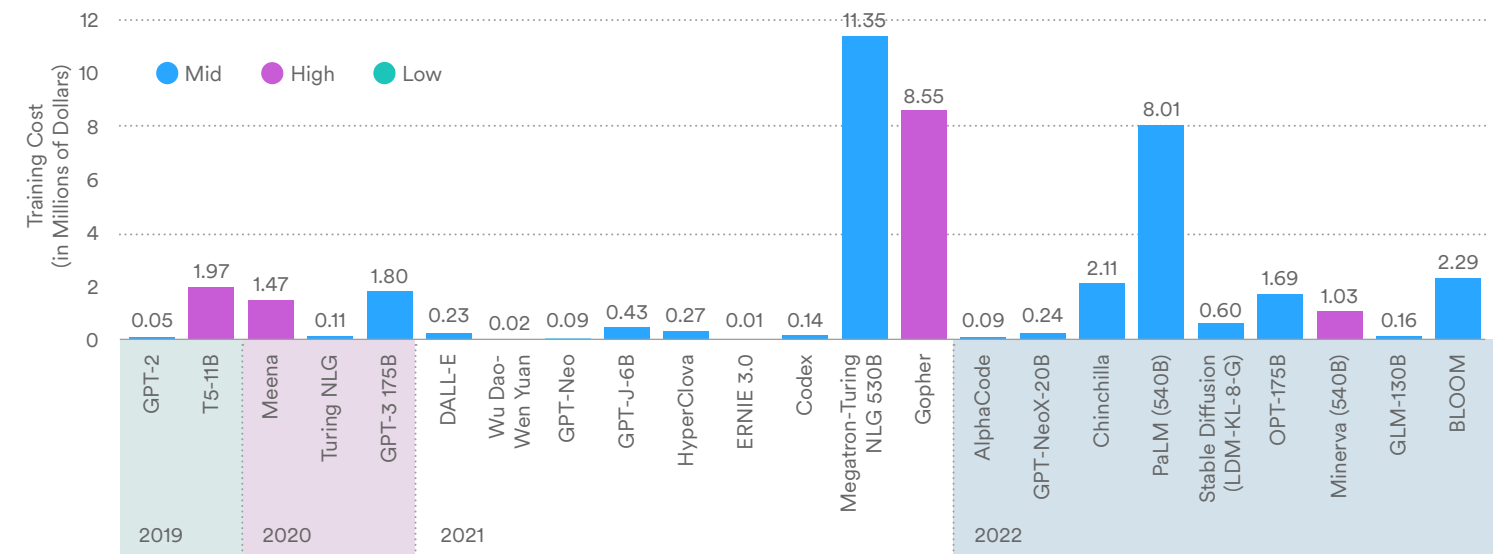
AI Index

HAI and the AI Index Steering Committee published the 2023 AI Index report in April 2023. 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 latest report included new analysis on foundation models, including their geopolitics and training costs, the environmental impact of AI systems, K-12 AI education, and public opinion trends in AI. The AI Index also broadened its tracking of global AI legislation from 25 countries in 2022 to 127 in 2023.

The AI Index received positive press coverage around the time of the launch and throughout the rest of the academic year. Highlights included pieces from *The Washington Post*, Bloomberg, IEEE Spectrum, and the *San Francisco Chronicle*. In 2023, website visits and report downloads also tripled.

Estimated Training Cost of Select Large Language and Multimodal Models The AI Index estimates validate popular claims that large language and multimodal models are increasingly costing millions of dollars to train. For example, Chinchilla, a large language model launched by DeepMind in May 2022, is estimated to have cost $2.1 million, while BLOOM's training is thought to have cost $2.3 million.

Estimated training cost of select large language & multimodal models

Legend: ● Mid ● High ● Low

Model	Year	Cost
GPT-2	2019	0.05
T5-11B	2019	1.97
Meena	2020	1.47
Turing NLG	2020	0.11
GPT-3 175B	2020	1.80
DALL-E	2021	0.23
Wu Dao-Wen Yuan	2021	0.02
GPT-Neo	2021	0.09
GPT-J-6B	2021	0.43
HyperClova	2021	0.27
ERNIE 3.0	2021	0.01
Codex	2021	0.14
Megatron-Turing NLG 530B	2021	11.35
Gopher	2021	8.55
AlphaCode	2022	0.09
GPT-NeoX-20B	2022	0.24
Chinchilla	2022	2.11
PaLM (540B)	2022	8.01
Stable Diffusion (LDM-KL-8-G)	2022	0.60
OPT-175B	2022	1.69
Minerva (540B)	2022	1.03
GLM-130B	2022	0.16
BLOOM	2022	2.29

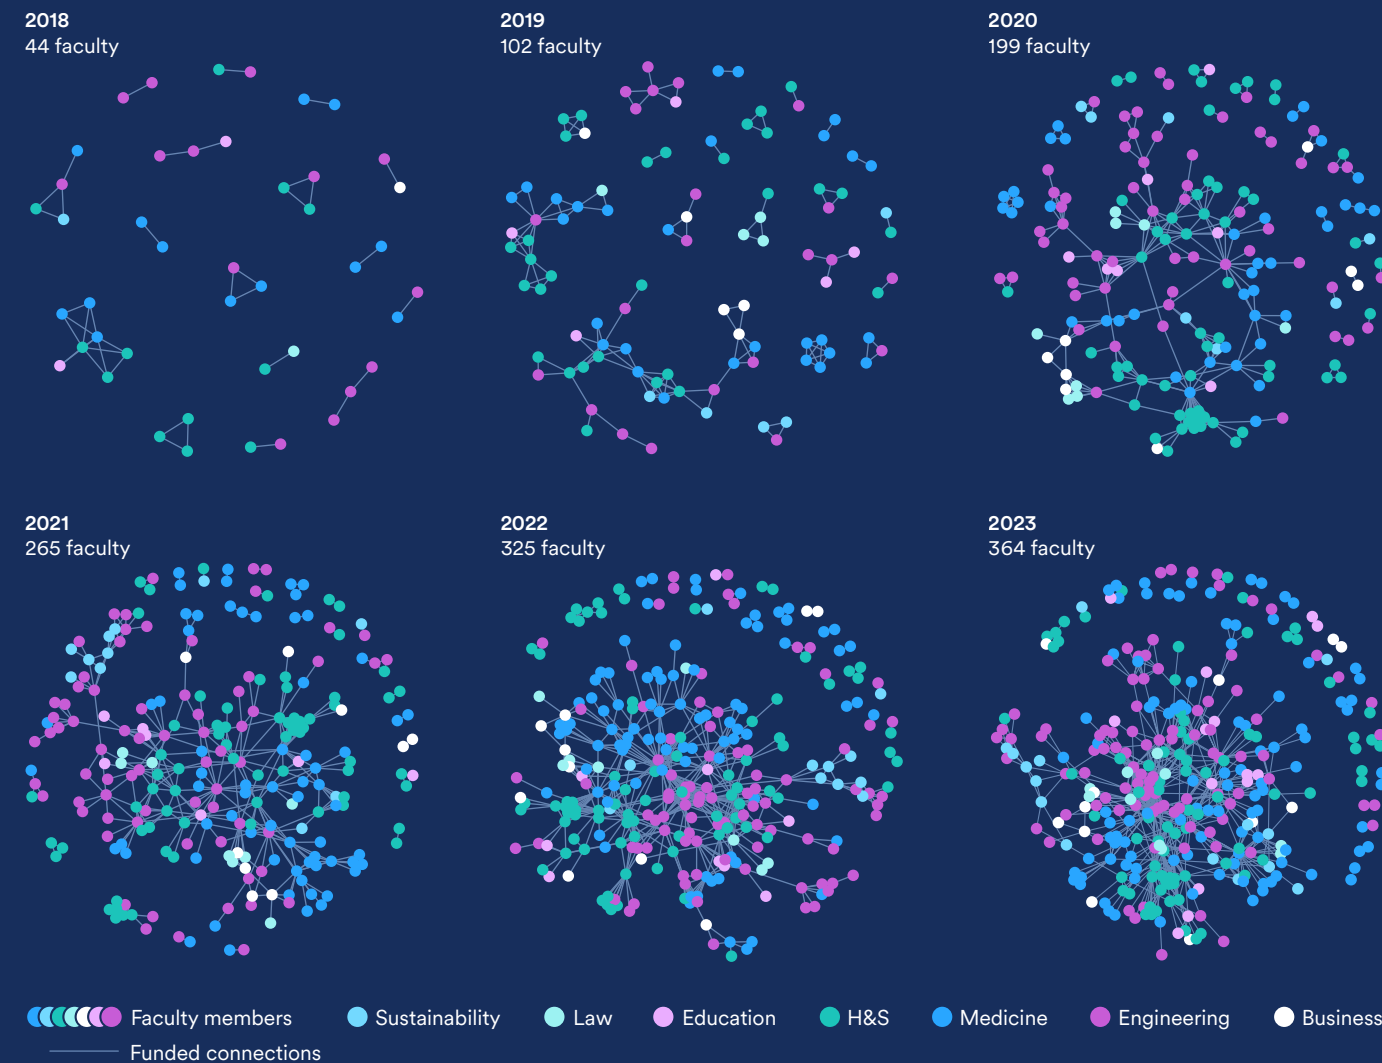Y-axis: Training Cost (in Millions of Dollars)

Source: AI Index, 2022 I Chart: 2023 AI Index Report

Grant Programs for AI Research

Since its founding, HAI's grant programs have supported over 300 faculty members from all seven Stanford schools focusing on projects aligned with three research focus areas: human impact, augment human capabilities, and intelligence. Many of these projects span multiple departments, in keeping with the commitment to support interdisciplinary AI research.
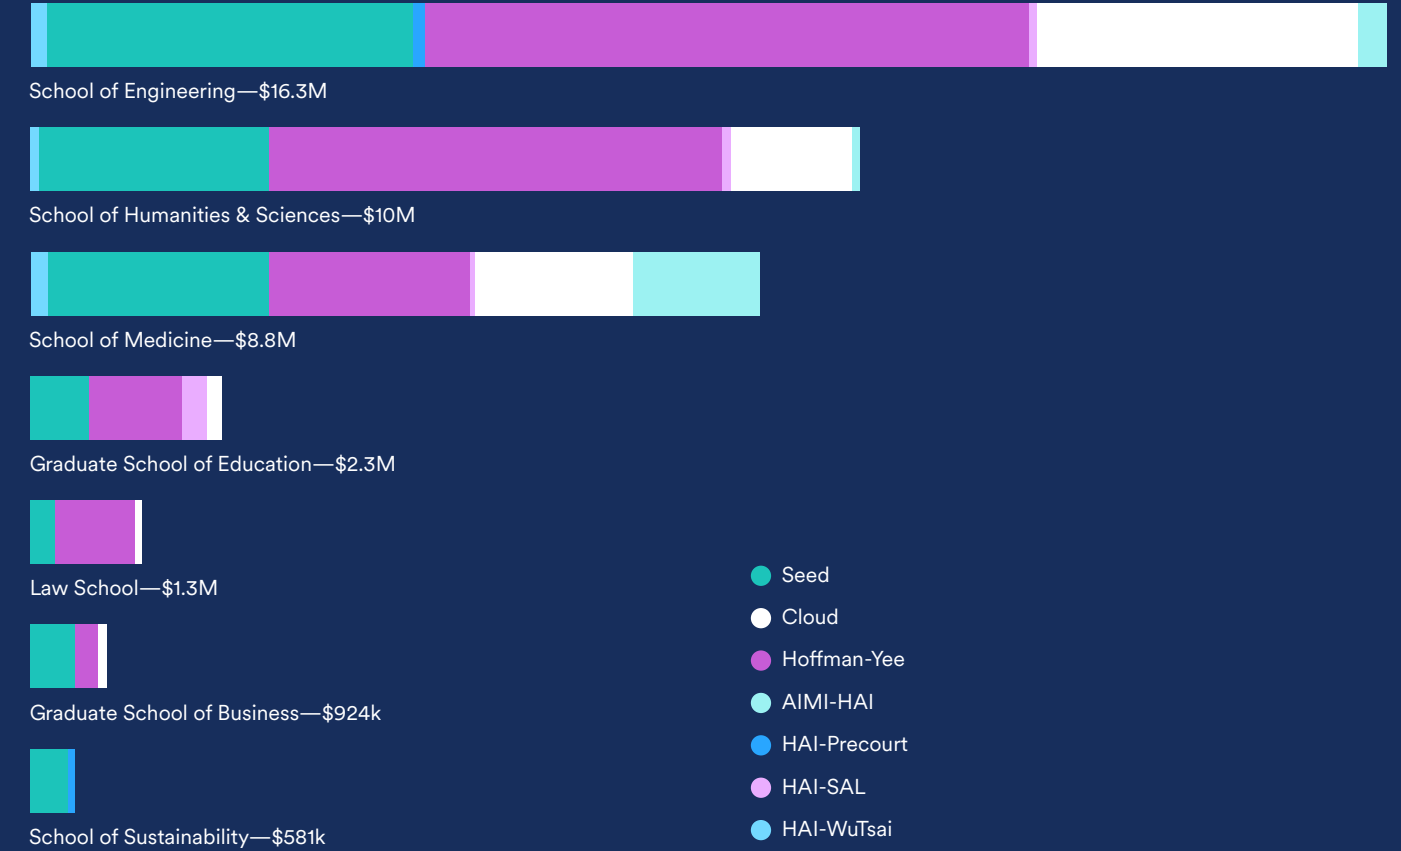
The HAI network of funded faculty continues to grow

The chart below illustrates the interdisciplinary growth of the HAI research network. Over the years, more dots (faculty) appear and they cluster closer together demonstrating that they are more often working together. Importantly, the clusters are more than one color meaning that the connections are increasingly interdisciplinary. Through the grant programs, HAI has distributed over $40 million in funding across Stanford.

2018
44 faculty

2019
102 faculty

2020
199 faculty

2021
265 faculty

2022
325 faculty

2023
364 faculty

● Faculty members ● Sustainability ● Law ● Education ● H&S ● Medicine ● Engineering ● Business
— Funded connections

2018-2023, >$40M to faculty in all seven schools

Funding by School

School of Engineering—$16.3M

School of Humanities & Sciences—$10M

School of Medicine—$8.8M

Graduate School of Education—$2.3M

Law School—$1.3M

Graduate School of Business—$924k

School of Sustainability—$581k

● Seed
○ Cloud
● Hoffman-Yee
● AIMI-HAI
● HAI-Precourt
● HAI-SAL
● HAI-WuTsai

Many faculty who have been involved in the grant programs express how impactful this funding has been in incentivizing the interdisciplinary work that's so essential to human-centered AI:

> An injection of funds into the research ecosystem, especially for interdisciplinary projects, is much needed and appreciated.
>
> **Mike Frank**, Psychology

> The best part was being able to think big and abstract and have a supportive group helping to make a solution a reality.
>
> **Christian Rose**, Emergency Medicine

> The best part of the program is that it provides funding for new and riskier projects.
>
> **Jackelyn Hwang**, Sociology

HAI runs four types of grant programs with specific and disparate objectives:

Seed Research Grants

The Seed Research Grants, supported early on by Steve and Roberta Denning and for the past four years by Dalio Philanthropies, are dedicated to supporting new, ambitious, and speculative ideas with the objective of getting initial results. Since the program's start in 2018, approximately $12 million in awards has attracted over $17 million in external funding creating a multiplier effect. This outcome demonstrates the program's promise in identifying high-reward, speculative, and early-stage research. Additionally, due to the breadth of proposals the institute receives from across campus, the seed grants provide a bellwether for what is to come in the field of AI.

Project Highlight

The Mineral-X initiative was developed based on relationships that emerged from an HAI seed grant project. The project uses Partially Observable Markov Decision Processes to build intelligent agents for two use cases: (1) safely and efficiently storing CO_2 and (2) safely and efficiently discovering battery metals required for technologies such as electric vehicles. This initiative is working to develop an Intelligent Agent by formulating sequential decision problems in a real-world setting that can address the speed and urgency of climate change.

Partnership Grants

In 2023, HAI partnered Artificial Intelligence in Medicine and Imaging (AIMI) and the Stanford Accelerator for Learning on grant programs. These partnership grants enable the institute to quickly initiate new research in key areas as AI advancements continue to accelerate. For example, in February 2023, after the launch of ChatGPT, in collaboration with the Stanford Accelerator for Learning, HAI quickly launched a grant program focused on exploring how generative AI can be applied in novel ways to support learning or investigate critical issues in learning contexts.

Project Highlight

In July 2023, a team of researchers from HAI and AIMI published their work on "The shaky foundations of large language models and foundation models for electronic health records." The project examined 84 foundation models trained on non-imaging EMR data and created a taxonomy delineating their architectures, training data, and potential use cases. Researchers found that most models are trained on small, narrowly scoped clinical datasets or broad, public biomedical corpora and are evaluated on tasks that do not provide meaningful insights on their usefulness to health systems. Considering these findings, the authors proposed an improved evaluation framework for measuring the benefits of clinical foundation models that is more closely grounded to metrics that matter in health care.

Hoffman-Yee Grants

The Hoffman-Yee Grants, generously supported by Reid Hoffman and Michelle Yee, fund "asteroid shot" ideas to address significant scientific, technical, or societal challenges. HAI believes the results of these projects will play a significant role in defining future work in AI from academia to industry, government, and civil society. In FY 2022-23, 12 funded projects produced over 35 publications in top AI journals and conferences, including NeurIPS, IEEE, ICLR, and ICML.

Project Highlight

AddBiomechanics is part of a project titled, "Intelligent Wearable Robotic Devices for Augmenting Human Locomotion." This team aims to build an AI-enabled human exoskeleton. In order to do this, a model first needs to be created that fully understands how humans move. AddBiomechanics.org is an effort to construct such a dataset, which includes physical details such as force magnitudes and body properties. AddBiomechanics provides an automated platform that helps practitioners process motion capture data; in exchange, they agree to publish the anonymized processed motion results under a Creative Commons license. Within the first few weeks of the project, over 100 universities across a dozen countries had uploaded 4,000 motion files totaling 13Gb of data.

Photos: Christine Baker

Cloud Credit Grants

The Cloud Credit Grants support emerging research that requires advanced computational resources provided by the commercial cloud. In FY 2022-23, HAI distributed $1.2 million in cloud credits with generous gifts from Microsoft Azure and Google Cloud. While the institute is only able to support a small portion of the requests it receives, the program continues to make an outsized impact across the social sector, government, and industry, resulting in widely adopted models (BioMedLM) and algorithms (FlashAttention), leveraging social media to measure well-being, and detecting neighborhood gentrification at a large scale.

Project Highlight

DSPy is an open-source project that enables researchers to build powerful systems with small, inexpensive language models (LMs). It is a framework for algorithmically optimizing LM prompts and weights, especially when LMs are used one or more times within a pipeline. DSPy is rapidly gaining popularity with a large and growing community of contributors.

Ethics and Society Review

As a HAI grant program standard, all research proposals undergo multiple rounds of rigorous faculty review as well as an Ethics and Society Review to assess for any potential negative impact on society before being green-lighted for funding. Due to its demonstrated success, the Ethics and Society Review is now scaling to other research domains at Stanford, beyond HAI and AI.

Research Centers

Collaboration with labs, centers, and institutes across and beyond the Stanford campus continues to be essential to the realization of HAI's mission.

Photo: Christine Baker

Regulation, Evaluation, and Governance Lab (RegLab)

Daniel E. Ho, Faculty Director | Christine Tsang, Executive Director

Stanford RegLab builds high-impact partnerships for data science and AI in the public sector. RegLab envisions a future where government can leverage state-of-the-art advances to address society's most urgent challenges.

Over the past year, RegLab made significant progress across five areas: Environmental Governance, Tax Administration, Health Equity, Mass Adjudication, and AI Policy.

Research Highlights

Modernizing Tax Administration

Developed an active learning system to reduce the annual $500B tax gap and uncovered the first direct evidence of differences in audit rates by race, sparking an IRS "overhaul" of auditing around the Earned Income Tax Credit

Improving Worker Protections

Designed and evaluated an NLP assistance tool to improve the adjudication of workers' compensation claims in a randomized trial with U.S. Department of Labor, affecting a $3-4B claims system

Detecting Pollution

Piloted a near real-time system for identifying illegal dumping under the Clean Water Act using computer vision and aerial imagery to modernize environmental compliance

Benchmarking Law

Developed a test suite showing that LLMs generate "legal hallucinations" 60-90% of the time and a benchmark for legal reasoning (LegalBench) in LLMs

AI Policy Highlights

— Testimony, U.S. Senate Committee on Homeland Security and Governmental Affairs

— Testimony, U.S. House Subcommittee on Cybersecurity, Information Technology, and Government Innovation

— Service on National AI Advisory Committee and the Administrative Conference of the United States

— Letters to OMB on risk management and open innovation

— The AI regulatory alignment problem

— The Bureaucratic Challenges to AI Governance

Summer Institute

The RegLab Summer Institute is an experiential learning fellowship designed to train the next generation of leaders in public sector AI. We hosted 11 PhD, masters, and undergraduate students from 8 universities, with research areas ranging from LLMs to computer vision and satellite imagery to responsible AI.

Research Leadership

Kit Rodolfa
Research Director

Derek Ouyang
Research Manager

RegLab has a core research staff of over twenty individuals and has worked with hundreds of law, graduate, and undergraduate students across all of Stanford's seven schools.

Daniel E. Ho, Director of RegLab

Center for Research on Foundation Models (CRFM)

Percy Liang, Director

The Center for Research on Foundation Models (CRFM) was founded in August 2021, bringing together more than 400 students, postdocs, engineers, and faculty across Stanford to make fundamental advances in the study of foundation models. In AY 2022-23, the center made progress across three areas: technical research, applied research, and societal research.

CRFM scholars celebrating the one-year anniversary of its founding. Photos: Jeanina Casusi

Technical Research

CRFM makes core technical contributions toward building, using, and understanding foundation models:

— Holistic Evaluation of Language Models

— Alpaca: A Strong, Replicable Instruction-Following Model

— Generative Agents: Interactive Simulacra of Human Behavior

— Hungry Hungry Hippos: Toward Language Modeling with State Space Models

— Sophia: A Scalable Stochastic Second-order Optimizer for Language Model Pre-training

Applied Research

CRFM builds foundation models and other artifacts that provide value in specific application domains such as medicine and law:

— BioMedLM

— Language-Driven Representation Learning for Robotics

— Anticipatory Music Transformer: A Controllable Infilling Model for Music

— LegalBench: A Collaboratively Built Benchmark for Measuring Legal Reasoning in Large Language Models

— Levanter—Legible, Scalable, Reproducible Foundation Models with JAX

Societal Research

CRFM shapes policy and clarifies the societal impact of foundation models to facilitate the responsible governance of this new technology:

— Ecosystem Graphs: The Social Footprint of Foundation Models

— Stanford-Princeton Response to the US National Telecommunications and Information Administration (NTIA) on AI accountability policy, drafted in partnership with Princeton University's Center for Information Technology Policy (CITP)

— Do Foundation Model Providers Comply with the Draft EU AI Act?

— Foundation Models and Fair Use

— Whose Opinions Do Language Models Reflect?

Percy Liang, director of CRFM. Photo: Jeanina Casusi

Percy Liang
Associate Professor of Computer Science and, by courtesy, of Statistics; Senior Fellow, Stanford HAI; Director, CRFM, Stanford University

Rishi Bommasani
Society Lead, CRFM; PhD Candidate of Computer Science, Stanford University

David Hall
Research Engineering Lead, CRFM, Stanford University

Yifan Mai
Research Engineer, CRFM, Stanford University

Ivan Zhou
Masters Student in Computer Science, Stanford University

Julian Quevedo
Student, Computer Science and Mathematics, Stanford University

Josselin Somerville
Masters Student in Computer Science, Stanford University

Tony Lee
Research Engineer, CRFM, Stanford University

Stanford Digital Economy Lab (DEL)

Erik Brynjolfsson, Director | Christie Ko, Executive Director

The Stanford Digital Economy Lab (DEL) is an interdisciplinary research group studying how digital technologies are transforming work, organizations, and the economy—toward a world of shared prosperity. DEL leverages the expertise and interests of Stanford faculty, researchers, and students to initiate and support a wide range of projects on the most pressing topics related to the digital economy. As an integral part of Stanford HAI, DEL pursues a multidisciplinary approach that encourages dialogue and collaboration across the domains of economics, business, technology, policy, and beyond.

Erik Brynjolfsson
Director

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

DEL has made significant strides over the past year. It welcomed a new lead research scientist to further the pursuit of providing better measures of the digital economy and launched the Pol.is project, aimed at enhancing transparency and consensus building in democratic processes.

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

Digital Platforms and Society

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

Research Highlights

ADP National Employment Report

This ongoing collaboration with the ADP Research Institute measures the monthly changes in U.S. employment and pay based on payroll data from ADP. It also features the Pay Insights report, which provides a view of the wages and salaries of nearly 10 million employees over a 12-month period.

Generative AI at Work

This working paper explores the productivity impacts of large language models on knowledge work, such as customer support, and found measurable gains in the percent of issues resolved per hour.

Beyond the Turing Test—New Benchmarks for AI

This continuing research effort builds on Erik Brynjolfsson's exploration of augmentation versus automation in creating AI benchmarks to guide technology development toward more equitable economic outcomes. The work included a roundtable of technologists and economists, and the Lab continues to shape a strategy for the development of new AI benchmarks.

NSF Pilot on National Network for Critical Technology Assessment

DEL joined this pilot, aimed at advancing national strategy and policy surrounding AI investments and planning.

Student Engagement

DEL launched a new Stanford course, "The AI Awakening: Implications for the Economy and Society," which hosted visionary researchers and industry leaders in technology, economics, and business on topics including foundation models, vision and robotics, work and employment, bias and explainability, AI and geopolitics, and a world without work. The course culminated in a session where student teams pitched their policy or research proposals and business plans in front of an audience with a "shark tank"-like panel of guest judges.

Cross-Campus Collaboration

Stanford University became a founding partner of the Project Liberty's Institute, a consortium of experts in law, policy, social sciences, economics, ethics, and technology working together to shape a new internet designed for the common good. Stanford's activities related to the institute are enabled by a generous gift from Frank McCourt, Jr., and anchored in three university hubs: the Stanford Digital Economy Lab, the McCoy Family Center for Ethics and Society, and the Freeman Spogli Institute's Cyber Policy Center.

Support from Project Liberty will allow DEL to expand its research into how digital technologies, particularly AI, affect society and the economy—and shape the global conversation about transforming social media and other institutions for the betterment of society.

Advisory Group

The Stanford Digital Economy Lab's Advisory Group convened twice during the academic year. In November 2022, the group discussed economic measurement and decision making in the digital age and AI and the future of work. In April 2023, members explored digital platforms, society, and progress as well as opportunities and risks of generative AI. In AY 2022-23, Steve Denning became chairman emeritus, and Sarah Friar (CEO, Nextdoor) took the post as co-chair with James Manyika. The group also welcomed four new members: Dario Amodei (Anthropic), Jon Levin (Stanford GSB), Mira Murati (OpenAI), and Nela Richardson (ADP).

Visiting Faculty

Alex "Sandy" Pentland
MIT Media Lab

Research Team

Wang Jin Research Scientist	**Jae Joon Lee** Postdoctoral Researcher
David Nguyen Research Scientist	**Gabriel Unger** Postdoctoral Researcher
Andrew Wang Research Scientist	**Wajeeha Ahmad** PhD Candidate
Max Feng Data Engineer	**Rebecca Janßen** Visiting PhD Student, University of Zurich
Ruyu Chen Postdoctoral Researcher	

Policy & Society

The policy and society initiative at HAI continued to guide the responsible development and governance of AI worldwide. The team engaged with global policymakers on cutting-edge, multidisciplinary AI research, accompanied by custom educational training, multi-stakeholder convening, and briefings on key AI policy issues. The team also supported the needs of leaders in mission-driven civil society organizations to leverage AI better for social impact.

Policy Engagement

In response to an unprecedented surge in inbound requests from policymakers, HAI further strengthened its engagement with policymakers across California, the United States, the European Union, and beyond with the goal of enriching the AI policy discussion with evidence-based research from Stanford.

National AI Research Resource (NAIRR)

In late 2019, Fei-Fei Li and John Etchemendy were the first to call for the U.S. government to create a close partnership between academia, government, industry, and civil society to provide researchers equitable access to high-end computational resources, large-scale government datasets. Since then, HAI has led efforts with 22 top computer science universities and a bipartisan, bicameral group of lawmakers to bring the NAIRR to fruition. This includes the National AI Research Resource Task Force Act of 2021, which established a federal task force in which Fei-Fei Li participates. In addition, she was a member of a policy practicum with Stanford Law School that led to a white paper with an exhaustive study on precisely how the United States can create, implement, and maintain an NAIRR.

As of 2023, one of the most significant U.S. policy advances in AI is the introduction of the legislation to authorize the creation of NAIRR. The U.S. Congress has introduced the bipartisan and bicameral "Creating Resource for Every American to Experiment with Artificial Intelligence Act of 2023 (CREATE AI Act)," an important step toward establishing the NAIRR to provide federally funded compute resources and access to government datasets for academic research. Senator Martin Heinrich, the co-chair of the Senate AI Caucus, recognized the important role HAI's study played in advancing the concept.

Rep. Anna Eshoo (CA-16), Fei-Fei Li, Russ Altman, and Jennifer King speaking at an HAI Seminar, Unlocking Public Sector AI Innovation: Next Steps for the National AI Research Resource. Photo: Don Feria

Meeting with President Biden

Fei-Fei Li and Rob Reich met with President Biden at the consequential meeting on AI in June 2023. Fei-Fei Li set the tone of the conversation by encouraging the President to adopt a moon shot mentality when thinking about AI.

MEETING ON ARTIFICIAL INTELLIGENCE

U.S. Engagement

HAI Policy in Washington, D.C.: The HAI policy and society team led Stanford faculty members and researchers to hold briefings and congressional testimonies for policymakers, including the following:

U.S. Federal Agencies

Central Intelligence Agency

Department of Commerce (National Institutes of Standards and Technology, National Telecommunications and Information Association)

Department of Defense

Department of Labor

Department of State

Federal Trade Commission

National Economic Council

National Geospatial-Intelligence Agency

National Institutes of Health

National Science Foundation

National Security Council

Office of the U.S. Trade Representative

U.S. Agency for International Development

White House Domestic Policy Council

White House Office of Management and Budget

White House Office of Science and Technology Policy

Congressional Offices and Committees

Congressional and Senate AI Caucuses

Congressional Black Caucus

House Committee on Energy and Commerce

House Committee on Oversight and Government Reform

House Committee on Science, Space, and Technology

House Committee on the Judiciary

Office of Representative Jake Auchincloss

Office of Representative Jamaal Bowman

Office of Representative Yvette Clarke

Office of Representative Anna Eshoo

Office of Representative Derek Kilmer

Office of Representative Frank Lucas

Office of Representative Jay Obernolte

Office of Senate Majority Leader Chuck Schumer

Office of Senator Cory Booker

Office of Senator Chris Murphy

Office of Senator Alex Padilla

Office of Senator Gary Peters

Office of Senator Mike Rounds

Senate Committee on Appropriations

Senate Committee on Commerce, Science, and Transportation

Senate Committee on Health, Education, Labor and Pensions

Senate Committee on Homeland Security and Governmental Affairs

Senate Committee on the Judiciary

Congressional Testimony

Daniel E. Ho testified before the Senate Committee on Homeland Security and Governmental Affairs in May 2023 on Artificial Intelligence in Government.

Fei-Fei Li testified before the Senate Committee on Homeland Security and Governmental Affairs in September 2023 on Governing AI through Acquisition and Procurement.

DANIEL HO

California Governor's Office

California Governor Gavin Newsom signed an executive order on AI, which officially named HAI as a partner of the State of California on AI and announced a co-hosted summit to take place in 2024.

European Union (EU) Engagement

The HAI policy and society team worked with the institute's Center for Research on Foundation Models (CRFM) and Marietje Schaake, HAI's international policy fellow, to guide the world's first comprehensive AI regulation, the EU AI Act, as it enters the final negotiations stage. This effort included a series of meetings in Brussels, during which Marietje Schaake, Rishi Bommasani, and Daniel Zhang presented their recent research on the EU AI Act, with key players leading negotiations on the act in the European Parliament, the European Commission, and the European Council.

HAI also hosted meetings at Stanford with delegations from the European Commission Directorate-General for Communications Networks, Content, and Technology; members of the European Parliament; the Presidency of the European Commission; and Spanish Secretary of State for Digitalization and Artificial Intelligence.

International Governments and Organizations

In August 2023, Fei-Fei Li was appointed to the U.N. Secretary General's Scientific Advisory Board. Also during the academic year, Marietje Schaake was appointed to the U.N. High-Level Advisory Body on Artificial Intelligence.

HAI also hosted a number of visits from high-level representatives to engage and inform government officials on AI policy, including Singapore's deputy prime minister and the Ministry for Communication and Information; U.K.'s Department for Science, Innovation and Technology; Brazil's Ministry of the Economy; and other tech representatives from governments of Japan, South Korea, Estonia, and Norway.

Education and Training for Policymakers

HAI's policymaker education program has made significant strides in bridging the gap between the rapid advancements in AI and a comprehensive understanding of the technology at the policymaking level.

Congressional Boot Camp on AI

In August 2023, HAI hosted the 2023 Congressional Boot Camp on AI with a diverse cohort of 28 senior, bipartisan, and bicameral staffers at Stanford and participation from more than 20 Stanford faculty across disciplines, as well as Silicon Valley leaders and pioneers from civil society organizations. HAI designed the boot camp to equip congressional staff, who play a key role in shaping and developing AI policy, with the comprehensive knowledge needed to think critically about regulating and governing this emerging technology. The event attracted more than 140 applications, a 40 percent jump from the previous year.

Online Training Series for Federal Employees with the White House

In October 2022, HAI held a six-session Online Training Series for Federal Employees with the White House, in partnership with the White House Office of Management and Budget (OMB) and General Services Administration (GSA). This was a unique, high-impact opportunity for HAI to help fulfill an important congressional mandate of the AI Training Act of 2022. HAI was the sole academic institute to help teach more than 8,000 registered federal employees.

Emma Brunskill presenting at the 2023 Congressional Boot Camp. Photo: Christine Baker.

Policy Research

HAI leverages the university's strength across disciplines—including computer science, business, economics, education, law, literature, medicine, neuroscience, philosophy, sustainability, and more—to create innovative policy research on AI. We translate cutting-edge, multidisciplinary AI research for the policy audience and produce original AI-related policy research to equip policymakers with tools to understand and govern the technology. Such initiatives facilitate the communication of research findings to policymakers, ensuring that policies are enriched with cutting-edge research insights and enabling them to be more effectively aligned with real-world needs and challenges.

Policy Briefs

Promoting Algorithmic Fairness in Clinical Risk Prediction

Led by Nigam Shah, this brief discusses how fairness techniques could impact clinical predictive models and suggests the best way forward in developing fairer ML tools for a clinical setting.

Using Algorithm Audits To Understand AI

Led by Jeffrey Hancock, this brief reviews the history of algorithm auditing, describes its current state, and offers best practices for conducting algorithm audits today.

Balancing Fairness and Efficiency in Health Plan Payments

Led by Sherri Rose, this brief discusses how to realign the health care market's incentives in favor of patients, recommending interventions that shape companies' incentives around the pricing models they deploy.

Algorithms and the Perceived Legitimacy of Content Moderation

Led by Michael Bernstein, this brief dives into this problem by surveying people's views of Facebook's content moderation processes, providing a pathway for better online speech platforms and improving content moderation processes.

Using AI To Understand Residential Solar Power

Led by Ram Rajagopal, this brief showcases how to use AI to analyze variations in solar panel adoption in U.S. households and help policymakers tailor incentive mechanisms.

Whose Opinions Do Language Models Reflect?

Led by Tatsu Hashimoto, this brief introduces a quantitative framework that allows policymakers to evaluate the behavior of language models to assess what kinds of opinions they reflect.

Responses to Federal Requests for Information (RFIs)

Response to FTC's Advanced Notice of Proposed Rulemaking on Commercial Surveillance and Data Security

Led by Jennifer King, the research and filing of a response to the FTC on its ANPR is a first step toward the agency potentially adopting specific rules to advance consumer privacy and protect the public from commercial surveillance practices. Co-authors include Daniel Zhang, James Zou, and Stanford graduate students Caroline Meinhardt, Abel Ribbink, and Pete Warden.

Responses to NTIA's Request for Comment on AI Accountability Policy

Led by researchers from CRFM and the Princeton University's Center for Information Technology Policy (CITP), this response argues that pervasive opacity compromises accountability for foundation models and recommends the federal government invest in digital supply chain monitoring for foundation models, invest in public evaluations of foundation models, and incentivize research on guardrails for open-source models.

Foundation Model Policy Research

In collaboration with the CRFM, HAI presented a series of issue briefs on key policy issues related to foundation models. Drawing on the latest research and expert insights from the field, this series aims to provide policymakers and the public with a clear and nuanced understanding of these complex technologies and to help them make informed decisions about how best to regulate and govern their development and use.

The first two briefs in this series are Improving Transparency in AI Language Models: A Holistic Evaluation—led by Percy Liang, this brief introduces Holistic Evaluation of Language Models (HELM) as a framework to evaluate commercial application of AI use cases—and Do Foundation Model Providers Comply with the Draft EU AI Act?—also led by Percy Liang, this research identified early adoption of the most recent version of the EU AI Act and how current foundation model providers would not be compliant under the proposed regulations.

Faculty-led Original Research

Implementation Challenges to Three Pillars of America's AI Strategy

In collaboration with the RegLab and led by Daniel E. Ho, this white paper systematically examined the implementation status of two AI-related executive orders and the AI in Government Act. The authors performed a comprehensive search across more than 200 federal agencies to assess implementation of key requirements to identify regulatory authorities pertaining to AI and to enumerate AI use cases.

ChatGPT and Physicians' Malpractice Risk

Led by Michelle Mello, this research analyzes how physicians should evaluate applied AI applications, such as large language models, for clinical use, and how clinicians will need to navigate uncharted liability risks.

Open Source Analysis of Newly Available Thermal Satellite Imagery

Led by Amy Zegart, this research compiled the first-ever training dataset to detect nuclear proliferation activities based on commercially available thermal satellite imagery.

The Privacy-Bias Trade-off: Data Minimization and Racial Disparity Assessments in the U.S. Government

Led by Daniel E. Ho and Jennifer King, who co-authored this award-winning paper at the 2023 FAccT conference, this paper examined the U.S. government's recent efforts to introduce government-wide equity assessments of federal programs. It further considered a range of policy solutions to address the tension between data minimization principles and racial disparity assessments.

The Stanford Tech, Ethics, and Policy (TEP) Fellowship

Co-hosted with the McCoy Family Center for Ethics and Society, this fellowship aims to build a talent pipeline between tech students at Stanford and Washington, D.C.

The fellowship trains 10 Stanford graduate students studying a variety of STEM subjects—from AI to computer science to bioengineering—on public policy via a spring course and places them in congressional offices, government agencies, and think tanks in Washington, D.C., for a summer fellowship to apply their technical knowledge to real-world policy challenges and applications. This is a unique opportunity for students to gain firsthand experience at the intersection of technology, ethics, and policy; contribute to the development of policies and regulations that will shape the future of technology; and develop their own professional networks and gain valuable skills in communication, collaboration, and problem-solving.

For our inaugural 2023 cohort, placements include the White House Office of Management and Budget, Senate Committee on Homeland Security and Governmental Affairs, Senate Appropriations Committee, Georgetown University's Center for Security and Emerging Technology, Brookings Institution, American Enterprise Institute, USAID, and more.

"
I always knew I wanted to work in tech policy in D.C., but I didn't know I'd be part of this program with other fellows and a network of connections, building a foundation for the future. I think there's an appetite for this on campus, and it's a model that means students end up working in really exciting places.

Liana Keesing, who worked at the Senate Committee on Homeland Security and Governmental Affairs

"
Coming to D.C. imbued me with a sense of possibility that there's an intersection to be found here and there are collaborations to be made. So, did it make me want to work in and with government public policy and adjacent areas for the public and civic interest? Definitely yes.

Regina Ta, who worked at the Brookings Institution

Policy Convenings

HAI hosts a diverse array of policy events and convenings, which serve as pivotal platforms for thought leadership in the realm of AI and its impact on society. These gatherings bring together a unique blend of experts from academia, industry, government, and civil society to engage in deep, insightful discussions to foster an environment of interdisciplinary collaboration and open dialogue.

In 2023, we hosted the following events:

National AI Advisory Committee Field Hearing
HAI hosted NAIAC's second hearing and first field hearing with Stanford Law School.

Advancing AI Audits: Project Showcase and Lessons Learned
HAI hosted an event, led by Marietje Schakke, that concluded the Stanford AI Audit Challenge with presentations from winners of the challenge.

AI Index Reception
HAI collaborated with the British Consul General's residence to officially roll out the 2023 AI Index report. The event was attended by 70 foreign government officials, including consuls general and tech ambassadors.

Decoding the EU Artificial Intelligence Act
HAI hosted a panel joined by Marietje Schaake, Rishi Bommasani, Irene Solaiman of Hugging Face, and Dragos Tudorache of the EU for a virtual program to explain the latest development in the EU AI Act. There were 600+ attendees for the program.

The National AI Advisory Committee Field Hearing. Photo: Christine Baker

Congressional staff members engaged in a session during the 2023 Congressional Boot Camp. Photo: Christine Baker

Society

In early 2023, HAI's policy team expanded its engagement to civil society organizations, both at home in the U.S. and across the world in global south countries. Building on the success of the policy portfolio, HAI's emerging society activities are driven by the objective of ensuring that diverse communities that are otherwise underrepresented in AI feel empowered to influence the design, development, and governance of this technology. To do this effectively, the sector needs to learn more about cutting-edge AI issues and inform the conversation from the lens of social impact.

Through closed-door multi-stakeholder convenings, course development, educational webinars, and custom training programs, the HAI society portfolio aims to meet the needs of community-serving organizations, such as nonprofits and philanthropic organizations. This work also is expanding HAI's footprint and impact for policy and society stakeholders in global south countries that are home to some of the largest democracies in the world, with a majority youth population that is actively experimenting with AI.

Global South Convenings and Workshops

Program Mapping

In partnership with the Freeman Spogli Institute's Cyber Policy Center (CPC), HAI was awarded an in-kind donation as part of The Rockefeller Foundation's competitive Bellagio Convening Center Program to host a closed-door global convening in Italy in July 2023. The event brought together 22 senior scholars and civil society leaders from countries in Asia, Africa, and Latin America to discuss the role in the global south. The goal of the convening was to catalyze new programming at HAI and CPC and further solidify our relationships with current and prospective funders. Insights from this effort led to the development of the program team's plan for a Global Technology Regulation initiative at Stanford and ongoing fundraising efforts.

The Asia Foundation Workshops

HAI launched the first-of-its-kind strategic partnership with the largest international development organization in Asia to expand the global footprint beyond Europe and North America. The goal of this partnership with The Asia Foundation is to co-design and implement custom educational training for policymakers and civil society organizations in Asia and the Pacific. Additional activities will include joint publications and events on AI issues of relevance to the policy and society audiences in Asia and of strategic importance for the United States. HAI's policy and society initiative launched this partnership on October 16, with AI Primer Workshop for The Asia Foundation's senior leadership team across country offices in the Asia Pacific region, including Nepal, Vietnam, Cambodia, Pakistan, Thailand, Fiji, Bangladesh, and the Philippines.

Special thanks to the Patrick J. McGovern Foundation for its generous gift that helps fund most of the policy and society initiatives at HAI.

Right: Nobel Peace Prize Winner and journalist Maria Ressa shares remarks on the impact of AI on the information ecosystem at a civil society roundtable hosted by HAI at Stanford. Photo: Jeanina Casusi

About HAI

HAI Leadership Team

John Etchemendy

Denning Co-Director

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

John is a professor of philosophy and symbolic systems, teaching the Philosophy of Artificial Intelligence to undergraduate students. Along with Russ Altman, he co-chaired the search for HAI Junior Fellows, which resulted in the hiring of Shriti Raj, HAI's newest Junior Faculty Fellow, in July 2023. He also serves on the AI Index Steering Committee. In his 35 years as a Stanford faculty member, John has served as the director of the Center for the Study of Language and Information, chair of the Philosophy Department, senior associate dean for the School of Humanities and Sciences, and was Stanford University's twelfth provost. He is also the author or co-author of seven books and numerous articles in logic, and has been co-editor of the Journal of Symbolic Logic and on the editorial board of several other journals.

Fei-Fei Li

Denning Co-Director

Sequoia Professor, Computer Science Department

During the AY 2022-23 Fei-Fei held public conversations with many leaders in AI, including Demis Hassabis, Eric Horvitz, and Geoffrey Hinton. She also met and advised President Biden in June 2023 during a convening on AI, as well as testified to the U.S. Senate Committee on the subject of Governing AI Through Acquisition and Procurement. Fei-Fei also won the Intel Innovation Lifetime Achievement Award, which recognized her for her significant contributions and advancements in the field of AI. Notably, Fei-Fei published her memoir, *The Worlds I See: Curiosity, Exploration, and Discovery at the Dawn of AI*, which details her upbringing as an immigrant in America, where her childhood love of science evolved into her adult calling to help advance the AI movement into what it is today. She also co-hosted the "AI+Education Summit: AI in the Service of Teaching and Learning," which showcased the latest research in AI and education from Stanford faculty and researchers and examined how AI can be used to advance human learning through new pedagogies, new modalities of assessment, new foundation models, and more.

James Landay

Denning Co-Director and Faculty Director of Research

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

James co-hosted the 2022 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 also co-hosted the 2023 HAI Spring Symposium with HAI Senior Fellow Michele Elam on "Creativity in the Age of AI: AI Impacting Arts, Arts Impacting AI," a conversation about creativity in the age of AI from aesthetic, technical, social, ethical and legal perspectives. James was also a panelist at a *TIME100* talk where he emphasized the important role that humans play in AI safety.

Russell Wald

Deputy Director

Russell oversees HAI's research, education, communications, administrative activities, industry programs, and policy and society hub. He collaborates with HAI's co-directors and faculty leaders to shape the strategic vision and human-centered mission of the organization. From 2020 to 2022, he served as HAI's first Director of Policy and later as Managing Director for Policy and Society. During this time, he established HAI's policy and society initiative, engaged with policymakers at the state, national, and international levels, created the impactful congressional boot camp, and co-authored several AI policy publications, including *Enhancing International Cooperation in AI Research: The Case for a Multilateral AI Research Institute* and *The Centrality of Data and Compute for AI Innovation: A Blueprint for a National Research Cloud*.

Associate Directors & Senior Fellows

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, "AI in the Loop: Humans in Charge," 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, in July 2023. During the AY 2022-23, Russ organized the Senior Fellow by-laws and process documents, working with leadership. In May 2023, he participated in the 2023 HAI Symposium on "Creativity in the Age of AI: AI Impacting Art, Art Impacting AI." He also was appointed to the faculty leadership group of RAISE-Health, a joint effort of HAI and the School of Medicine focusing on responsible use of AI in medicine.

Surya Ganguli

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

Surya was awarded and named to 2023 Schmidt Science Polymaths Program cohort, where he will pursue an interdisciplinary research program to understand how nonlinear elements within neural networks, both biological and artificial, interact to yield remarkable emergent computational capabilities. "Beyond neural scaling laws: beating power law scaling via data pruning," a paper that Surya co-authored, won the Neural Information Processing Systems Foundation's Outstanding Paper Award in November 2022. He also published "Catalyzing next-generation Artificial Intelligence through NeuroAI" in Nature Communications.

Curtis P. Langlotz

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

As the director of the Center for Artificial Intelligence in Medical Imaging (AIMI Center) and associate director of Stanford HAI, Curtwas instrumental in launching RAISE Health, a joint AI initiative of HAI and the School of Medicine. Curt also helped organize the 2022 Stanford AI+Health conference, which convened experts and leaders from academia, industry, government, and clinical practice to explore critical and emerging issues related to AI's impact across the spectrum of health care. Curt serves as a reviewer for HAI-AIMI grant programs and the HAI Startup series. Curt was elected as the President of the Radiology Society of North America in November 2023.

Christopher Manning

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

In AY 2022-23, Chris was awarded an honorary doctorate from the University of Amsterdam for his pioneering research contributions in the field of Natural Language Processing. His work on DetectGPT, at ICML 2023, has been covered in *Science* and *The Financial Times*. "Backpack Language Models," and "Recursive Deep Models for Semantic Compositionality over a Sentiment Treebank," both co-authored by Chris, won the Outstanding Paper Award and Test-of-Time Paper Award at ACL 2023, respectively. Chris also appeared in and helped organize the short video "AI at Stanford: 1962-2022," which received a Northern California Emmy Award. Additionally, the Stanford Chirpy Cardinal team, advised by Chris, won first place for scientific innovation in the Alexa Prize Socialbot Grand Challenge 5.

Julian Nyarko

Associate Professor of Law and Center Fellow at the Stanford Institute for Economic Policy Research Center Fellow

Julian Nyarko is dedicated to exploring the intersection of computational methods and legal studies. He brings his expertise in quantitative methods, natural language processing, and machine learning to the leadership team. His research delves into the intricacies of contract design under U.S. and international law, encompassing commercial and consumer contracts, smart contracts, and treaties. Julian's commitment to interdisciplinary collaboration is further evidenced in his work with fellow researchers across a wide array of additional topics, including causal inference, algorithmic fairness, and criminal justice.

Ge Wang

Associate Professor, Center for Computer Research in Music and Acoustics (CCRMA); Department of Music and, by courtesy, of Computer Science, Stanford University

Ge Wang has made significant contributions to the intersection of computer science and music, with a focus on the creation of computer-mediated tools, instruments, games, and toys that enable new forms of creative expression. Ge is the author of the ChucK music programming language and the founder of the Stanford Laptop Orchestra (SLOrk). Ge is also the author of *Artful Design: Technology in Search of the Sublime*, a photo comic book about how we shape technology—and how technology shapes us. His creative endeavors have reached millions of users, and he has been recognized as a 2016 Guggenheim Fellow. Among other Stanford courses, Ge teaches Music 356/CS 470: "Music and AI (a critical making course)". Ge participated in the "HAI Embedded Ethics Conference," which was co-sponsored by HAI and McCoy Family Center for Ethics in Society, as well as the "AI+Education Summit: AI in the Service of Teaching and Learning," which was co-sponsored by HAI and the Stanford Accelerator for Learning.

Amy Zegart

Senior Fellow at the Freeman Spogli Institute for International Studies; Morris Arnold and Nona Jean Cox Senior Fellow at the Hoover Institution; Professor, by courtesy, of Political Science

Amy Zegart brings a wealth of experience and insights in the fields of emerging technologies, U.S. national security, and global political risk management. Her distinguished academic career at Stanford and UCLA includes several award winning books about U.S. intelligence agencies and the recent bestseller, *Spies, Lies, and Algorithms: The History and Future of American Intelligence*, which was nominated by Princeton University press for a Pulitzer Prize. She has advised senior US foreign policy officials for more than two decades, served on the National Security Council staff, and writes frequently in Foreign Affairs, the Atlantic, and other media outlets.

Senior Fellows

Michele Elam

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

After cycling off her term as Associate Director, Michele continues to serve HAI as a Senior Fellow. In her role at HAI, Michele served as the faculty lead of the HAI Visiting Artist program and for the Arts+Tech Salon series. Michele co-led the HAI 2023 Spring Symposium, "Creativity in the Age of AI: AI Impacting Arts, Arts Impacting AI," which featured conversation among technologists, scholars, and creatives from both commercial and non-commercial sections. This conference focused on creativity in the age of AI from aesthetic, technical, social, ethical and legal perspectives, including two HAI Visiting Artists. In AY 2022-23, she taught Black Mirror: Representations of Race & Gender in AI, Mixed-Race Politics and Culture, James Baldwin & Twentieth Century Literature, and A.I.-Activism-Art.

Daniel E. Ho

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

Daniel serves as the director of the Stanford RegLab, which partners with government agencies to design and evaluate programs, policies, and technologies that modernize government. In 2022, the lab collaborated with the U.S. Department of the Treasury to uncover the first direct evidence of differences in audit rates by race, sparking the IRS to overhaul tax administration around the Earned Income Tax Credit. He spoke at the Stanford HAI Congressional Boot Camp on Artificial Intelligence and testified in the U.S. Senate Committee on Homeland Security and Governmental Affairs on AI policy. He additionally serves on the National AI Advisory Committee to advise the White House on AI policy.

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 helped organize the inaugural "Embedded Ethics Conference," which was co-sponsored by HAI and the Stanford McCoy Family Center for Ethics in Society. This conference offered broad discussions applicable to creating effective ethics education programs, covering the full life cycle of program creation, development, and implementation. He also met with President Biden during an AI convening in July 2023. Additionally, he spoke at the "AI+Education Summit: AI in the Service of Teaching and Learning," which showcased the latest research in AI and education from Stanford faculty and researchers and examined how AI can be used to advance human learning through new pedagogies, new modalities of assessment, new foundation models, and more. This summit was co-sponsored by HAI and the Stanford Accelerator for Learning.

Faculty & Fellows

HAI currently has more than 200 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 the 2022-23 academic year, HAI welcomed 16 graduate fellows from an array of disciplines, facilitating collaboration between engineers, social scientists, humanists, and others researching the future of human-centered AI. HAI also partnered with the McCoy Family Center for Ethics in Society to host four Embedded Ethics postdocs, a program that embeds the teaching of ethics directly into the core undergraduate courses of Stanford's computer science curriculum.

HAI Faculty Fellows

Peter Norvig

Distinguished Education Fellow

Peter is the HAI Distinguished Education Fellow, where he helps to 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.

Hariharan Subramonyam

Ram and Vijay Shriram HAI Faculty Fellow

Hari 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 HAI. Hari's research resides at the intersection of Human-Computer Interaction (HCI) and the Learning Sciences. He studies methods to enhance human learning through AI by engaging in cognitively informed design practices, co-designing with learners and educators, and developing transformative AI-enabled learning experiences. Additionally, through this research, Hari contributes tools and methodologies that emphasize ethical considerations, responsible design practices, and human values in the creation of AI experiences.

Johannes Eichstaedt

Ram and Vijay Shriram HAI Faculty Fellow

Johannes is a computational social scientist jointly appointed as a Ram and Vijay Shriram HAI Faculty Fellow and Assistant Professor of Psychology. His lab in Computational Psychology & Well-Being is currently focused on using LLMs to develop interventions for well-being and mental health to augment therapy and support therapists. Johannes also analyzes big social text data (Twitter, Facebook, Reddit) with NLP and AI to understand and measure psychological processes for individuals and populations—for example, to build early warning systems for the opioid epidemic and mental health. Across these lines of work, Johannes' overall focus is on using AI and NLP to improve the well-being and health of the population.

Shriti Raj

HAI Faculty Fellow

Shriti is a human-centered computing researcher. Her work focuses on augmenting health-related data reasoning and decision-making. She is jointly appointed as a HAI Faculty Fellow and an Assistant Professor in the Department of Medicine Center for Biomedical Informatics Research. Shriti received her PhD in Information from the University of Michigan, where she was advised by Dr. Mark Newman. Shriti's research develops and evaluates human-centered methods and tools to help people make health data and algorithms actionable.

HAI Policy Fellows

Jennifer King
Privacy and Data Policy Fellow, Stanford HAI

Rohini Kosoglu
Policy Fellow, Stanford HAI

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

2022–23 HAI Fellows

Ruyu Chen
Stanford Digital Economy Lab Postdoctoral Fellow

Nicholas A. Coles
CSLI Fellow and Director, Psychological Science Accelerator

Yuchen Cui
HAI Postdoctoral Fellow

Dapeng Feng
HAI Postdoctoral Fellow

Caroline Glidden
HAI Postdoctoral Fellow

Alex Goldmark
HAI Journalist-in-Residence

Farnaz Jahanbakhsh
HAI Postdoctoral Fellow

Christopher Kelly
HAI Postdoctoral Fellow

JJ Lee
Stanford Digital Economy Lab Postdoctoral Fellow

Juan Pava
Tech, Ethics & Policy Rising Scholar

Gabriel Unger
Stanford Digital Economy Lab Postdoctoral Fellow

Luca Vendraminelli
HAI Postdoctoral Fellow

Kaveh Waddell
HAI-JSK Fellow

Luwen Wan
HAI Postdoctoral Fellow

Zhecheng Wang
HAI Postdoctoral Fellow

Benji Xie
HAI-EIS Embedded EthiCS Fellow

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.

Research Centers and Partners

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 (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.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.

Responsible AI for Safe and Equitable Health (RAISE Health)

Artificial intelligence stands to revolutionize health and medicine—from advancing drug discovery to personalizing patient care. As AI innovation keeps accelerating, we must act now to ensure it advances with intention and purpose.

A joint initiative between **Stanford School of Medicine** and the Stanford Institute for **Human-Centered Artificial Intelligence** guides the responsible use of AI across biomedical research, education, and patient care.

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

Black in AI Innovation and Research (BlackAIR) Center for AI in Medicine & Imaging

Center for Comparative Studies in Race and Ethnicity Center for Ethics in Society

Center for International Security and Cooperation Cyber Policy Center

Digital Civil Society Lab

Hoover Institution

Institute for Diversity in the Arts

John S. Knight Journalism Fellowships

Office of the Vice President for the Arts

Stanford Aging and Ethnogeriatrics (SAGE) Research Center

Stanford Institute for Economic Policy Research (SIEPR)

External Partners

AI4All

HAI Events

Each year, HAI convenes a wide range of participants from academia, government, and industry through its two signature conferences and numerous workshops and seminars. During the 2022-23, a total of 119 events reached 27,655 attendees, including both in-person venues and virtual platforms.

Creativity in the Age of AI

Several HAI events on the theme of Creativity & AI, featuring HAI Visiting Artist Lauren Lee McCarthy:

2023 Spring Symposium

Artist showcase

Arts & Tech Salon

HAI Seminar with Lauren Lee McCarthy

Performance by Enongo Lumumba-Kasongo, Professor of Music at Brown University, at the 2023 Spring Symposium. Photo: Christine Baker

Photo Credit: Christine Baker

October 22

Inaugural HAI Orientation

The inaugural Stanford HAI Orientation gave attendees the opportunity to learn more about HAI's work, meet faculty and peers, and find out more about the ways in which to get involved.

Photo Credit: Christine Baker

November 22

HAI Fall Conference

HAI signature event

During this conference, the participants were challenged to rethink the phrase "Human-in-the-loop" and consider a future where humans remain at the center of all AI technologies.

Photo Credit: Christine Baker

January 23

Tanner Lecture: AI and Human Values with Seth Lazar

The 2023 Tanner Lecture was given by Seth Lazar, Professor of Philosophy at the Australian National University, an Australian Research Council (ARC) Future Fellow, and a Distinguished Research Fellow of the University of Oxford Institute for Ethics in AI.

April 23

Fei-Fei Li in Conversation with Demis Hassabis: Using AI to Accelerate Scientific Discovery

In this discussion, Demis Hassabis, co-founder and CEO of DeepMind Technologies, shared his company's efforts on the topic of how AI is rapidly changing the speed and breadth of scientific discovery. This was followed by a conversation with Fei-Fei Li on the future of AI.

May 23

AI and Human Values: A Conversation with Fei-Fei Li and Eric Horvitz

In this conversation, Tanner Lecturers Fei-Fei Li and Eric Horvitz discussed the topics of AI and Human Values. This talk featured topics covered during both Fei-Fei and Eric's Tanner Lectures.

July 23

Decoding the EU AI Act

This webinar brings together a distinguished panel of leading AI policy researchers and policymakers from the United States and the European Union who shed light on the key provisions in the draft EU AI Act, implications for the tech industry and transatlantic relationship, and what's next as EU enters trilogue negotiations among the Parliament, Commission, and Council to finalize the EU AI Act.

October · November · December · January · February · March · April · May · June · July · August

---

December 22

AI+Health Conference

At this wide-ranging online event, HAI, Artificial Intelligence in Medicine and Imaging (AIMI), and Center for Continuing Medical Education (CME) convened experts and leaders from academia, industry, government, and clinical practice to explore critical and emerging issues related to AI's impact across the spectrum of healthcare.

February 23

AI+Education Summit

Photo Credit: Trever Tachis

This summit showcased the latest research in AI and education from Stanford faculty and researchers, examined how AI can be used to advance human learning through new pedagogies, new modalities of assessment, new foundation models, and more. In doing so, the summit asked the question of how AI can transform teaching and learning in an ethical, equitable, and safe manner.

March 23

Embedded Ethics Conference

Photo Credit: Trever Tachis

The conference offered broad discussions applicable to creating effective ethics education programs, covering the full life cycle of program creation, development, and implementation. Panelists from different universities addressed how to get institutional buy-in to create a new ethics initiative and how to structure a successful program, while instructors presented live teaching demonstrations and ethics curricula.

June 23

Advancing AI Audit: Project Showcase and Lessons Learned

This event spotlighted the four award winning-entries of the AI Audit Challenge and featured insightful discussions from participants and AI experts on effective practices and valuable lessons learned from the Challenge. The Challenge, launched by HAI and the Cyber Policy Center, is an initiative that invited teams from around the world to submit their models, solutions, and tools with the goal of improving our ability to evaluate AI systems.
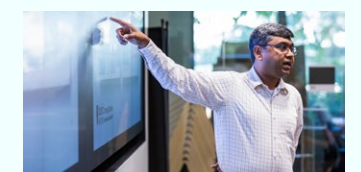
Aug 23

Congressional Boot Camp on AI

Photo Credit: Trever Tachis

During the bicameral, bipartisan Boot Camp, participants engage in sessions on mitigating AI risk, understanding and regulated foundation models, and impacts of AI on healthcare, education, climate, and online harm. The goal is to leave all participants with the conceptual framework needed to address the emerging technology landscape today and better anticipate the challenges of tomorrow.

Industry Programs

Led by Panos Madamopoulos-Moraris, Managing Director for Industry Programs and Partnerships, HAI continued to help advance collaborations with industry and helped to ensure that AI is deployed in human-centered ways. In AY23, HAI's community of industry funders included Accenture, Google, IBM, McKinsey, PepsiCo, SCBx, and Wells Fargo. Google and Microsoft provided cloud credits for the HAI cloud grant program.

Industry visits, outreach to senior decision-makers, and engagements with executives took place multiple times each week, engaging over 500 C-suite executives and industry leaders. Cross-cutting programming in AY 2022-2023 comprised six Executive Breakfasts featuring faculty, including Nigam Shah, Steve Blank, Rishee Jain, Gordon Wetzstein, Riitta Katila, and Jiajun Wu. The HAI Startup Series kicked off, showcasing 14 startups over 5 sessions and laying the foundation for an AI-focused entrepreneurial community at Stanford.

HAI launched the "Foundation Model Scholar Program," which brought faculty, postdocs, and graduate students together to demystify foundation models and their applications during an intensive three-day, on-campus program. Focusing on generative AI, its main objective was to elucidate the technical aspects of foundation models and delve into their far-reaching impact on AI development across various research domains and applications including health care and the digital economy. With participation from diverse global locations, including Mumbai, NYC, SF, Dublin, and Sydney, the program has been acclaimed as an "Oscar Winning moment" by Accenture's senior leadership. This collaboration highlighted the importance of a strategic partnership between academia and industry in advancing accessible and impactful learning experiences, and the program will be recurring and welcomes Accenture's senior leadership and ecosystem partners.

Research Collaborations and Faculty Engagement

HAI operationalized research collaborations with industry affiliates, fellow-mentor-advisor relations, and visiting scholars. The areas of collaboration spanned AI and society/economy, generative AI, responsible/trusted AI, foundation models, embodied AI, health care, fintech, and sustainability.

Funds allocated to research collaborations almost tripled year-over-year, reaching $4.5M and covering more than 35 research collaborations with over 250 beneficiaries at the schools of Engineering, Medicine, Humanities & Sciences, and Sustainability, including faculty, postdocs, and graduate students.

Member Testimonials

"Generative AI offers a whole new world of opportunities to simplify and improve our lives. As the AI landscape continues to evolve at a rapid scale, SAP focuses on building AI into core business processes to drive productivity and growth for companies worldwide. Our close collaboration with the Stanford Institute for Human-Centered Artificial Intelligence reinforces our strong commitment to building and using relevant, reliable, and responsible AI. It also provides us with access to the latest research and allows us to engage with talents and experts in the field as we create Business AI for our customers.

Juergen Mueller
Chief Technology Officer and Executive Board Member, SAP

"Collaborating with the team at HAI and their renowned list of talented academics will allow us to gain insights that will help shape the future of artificial intelligence. This is a unique relationship that will help us grasp the implications and significance of this technology and its impact in key industry areas including supply chains, organizational design, omnichannel consumer experience, and sustainability. As we look to the future, our goal is to drive innovation with AI at the forefront and guide industry leadership on how AI can benefit our society.

Athina Kanioura
Chief Strategy and Transformation Officer, PepsiCo

"Many of our clients have a deep appetite for learning and foresight into how to best cope with the AI disruption in their organizations. Joining Stanford HAI's corporate affiliate program gives us an opportunity to open up Silicon Valley with its leading-edge thinking, concentration of talent, and ecosystem of partners to all of our clients.

Ben Ellencweig
Senior Partner, McKinsey

"We chose the Stanford HAI corporate affiliate program as our flagship university collaboration, and the opportunity to work with world-class researchers and faculty members, as well as the dedicated staff from HAI, has been truly rewarding so far. We cannot be prouder of the immense impact of the Accenture Stanford Foundation Model Scholar Program that has become instrumental to achieve Accenture's AI ambition.

Lan Guan
Chief AI Officer, Accenture

"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
Chief Scientist, Google DeepMind and Google Research

"The HAI Corporate Program continues to be a great enabler for us to engage with the broader ecosystem across research and industry adjacencies. The insights and value we derive from it have helped us tune our new to market offerings and reach a better product market fit.

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

"As the HAI Corporate Program has matured over the last couple of years, we find increased opportunities for our researchers to engage with faculty and students to advance areas critical to the future of AI. In addition to the benefits of engaging across a broad community on technology, ethics, and policy, we are also beginning to see specific advances in AI techniques and approaches that are of mutual interest.

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

Content & Communications

Led by Carolyn Lehman, Director of Media and Communications, HAI's content team published 118 blog posts, 26 announcement stories, and 73 videos, which generated 1.1M page views and 362k YouTube views. The team also produced a series of law, policy, and AI updates; a series featuring new faculty leadership; and a foundation models video explainer. As a result of these efforts and others, HAI's social media following grew 84 percent in the academic year.

Media Relations

HAI continued to make a significant impact in the media by sharing individual and group insights regarding human-centered AI. HAI's executive team, assistant directors, fellows, and lab executive leaders are sought out regularly for analysis and comment on everything AI, from policy-related implications to the impact of research on medicine, education, and more.

From September 2022 to August 2023, HAI secured more than 856 unique stories in local, national, and global media outlets, including print and broadcast news, involving HAI faculty members, staff, and fellows.

The Economist

"How Foundation Models Are Turbo-Charging AI Progress"

In a cover story, the Economist published HAI's first major announcement of the Center for Research on Foundation Models—putting the term "foundation model," coined by CRFM, on the map.

Washington Post

"Stanford University Just Schooled Congress on AI"

This is one of about a dozen stories about HAI's AI Policy Boot Camp, which brings congressional aides to campus.

New York Times

"In Reversal Because of A.I., Office Jobs Are Now More at Risk"

This is one of 16 stories highlighting Erik Brynjolfsson's latest productivity and AI research.

Notably, the federal government also has been turning to HAI for its broad AI expertise. In June, President Joe Biden convened a group of technology leaders to debate what he called the "risks and enormous promises" of artificial intelligence. Fei-Fei Li and Rob Reich were two of only 10 invited guests to brief the president.

The 2023 AI Index succeeded in building engagement with media and influencers. Media coverage increased 53 percent over the 2022 AI Index coverage. HAI had 130 original pieces of coverage and 179 syndications from outlets including Bloomberg, Fortune, Fox News, The Guardian, and Insider. This provided a total reach of 723 million—based on the publications' unique visitors per month.

Content and Social Media

The HAI Blog continues to be a significant engagement driver. Some of the most viewed blogs include the following:

— AI Will Transform Teaching and Learning. Let's Get it Right

— 2023 State of AI in 14 Charts

— Scholars Build AI Detection Tool

— Fei-Fei and Demis Hassabis video

— What Is a Foundation Model? An Explainer for Non-Experts

HAI has continued to build a loyal online community—boosting visibility and impact among AI influencers, interested members of the public sector, education, research, and medicine.

Support for HAI

Stanford HAI is grateful for the generosity of those who provide support to make its work possible. The impact of this philanthropic support extends beyond individual projects; it ripples wider as our researchers seek further discoveries, share their findings, and make strides toward new therapeutic strategies for the enhancement of artificial intelligence and its role in shaping a better future for humanity.

The institute welcomes gifts made online. Please contact Stacey Sickels Boyce, Senior Associate Director of Development, with questions.
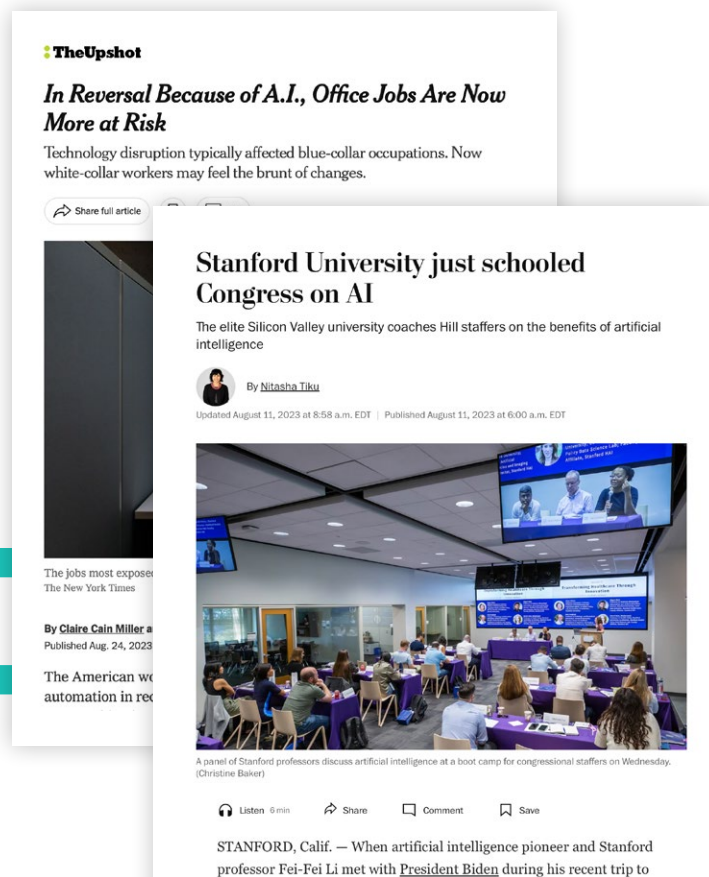
✖ @StanfordHAI
80,519
followers,
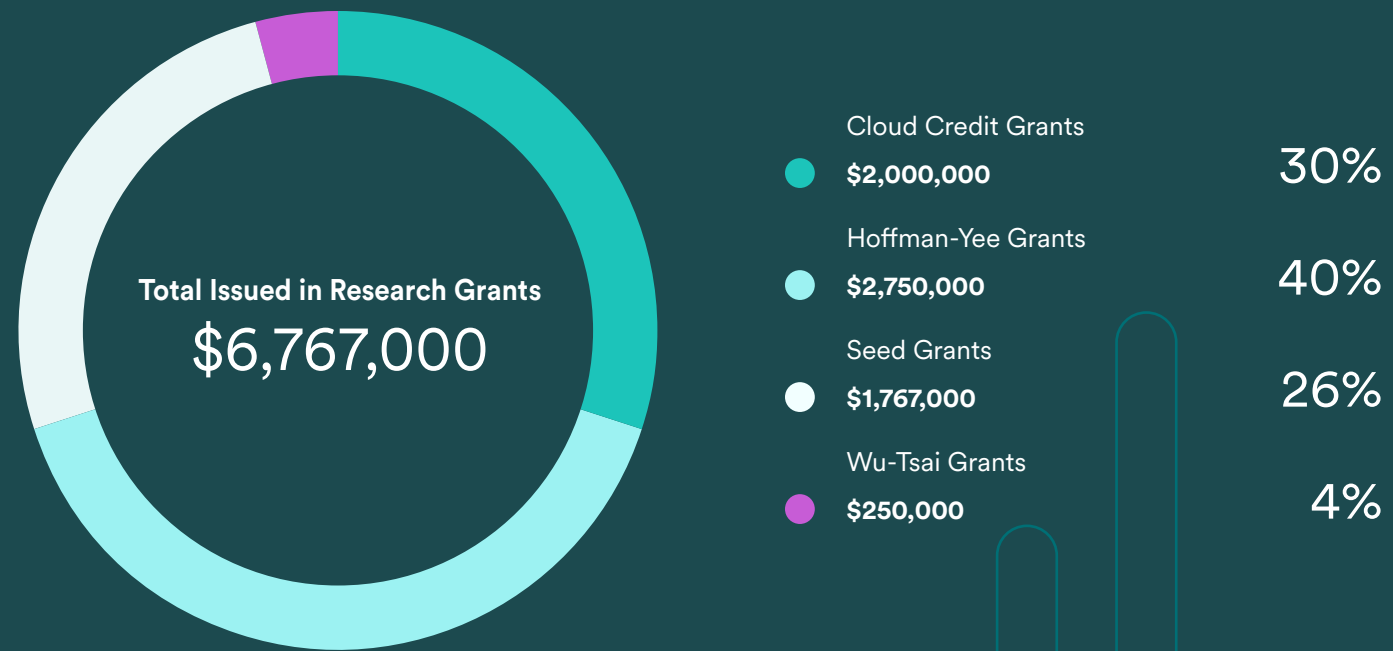and is growing by at least
25%
year over year since 2019

HAI's key audience on Twitter/X is academics, AI enthusiasts, and journalists.

in /stanfordhai
75,540
followers,
and is growing rapidly by
107%
year over year since 2019

Our engagement rates are higher than industry standards.

▶ @stanfordhai7626
300,000
views during the academic year,
and trackle
42,000
hours of watch time

Research Grant Support in 2023

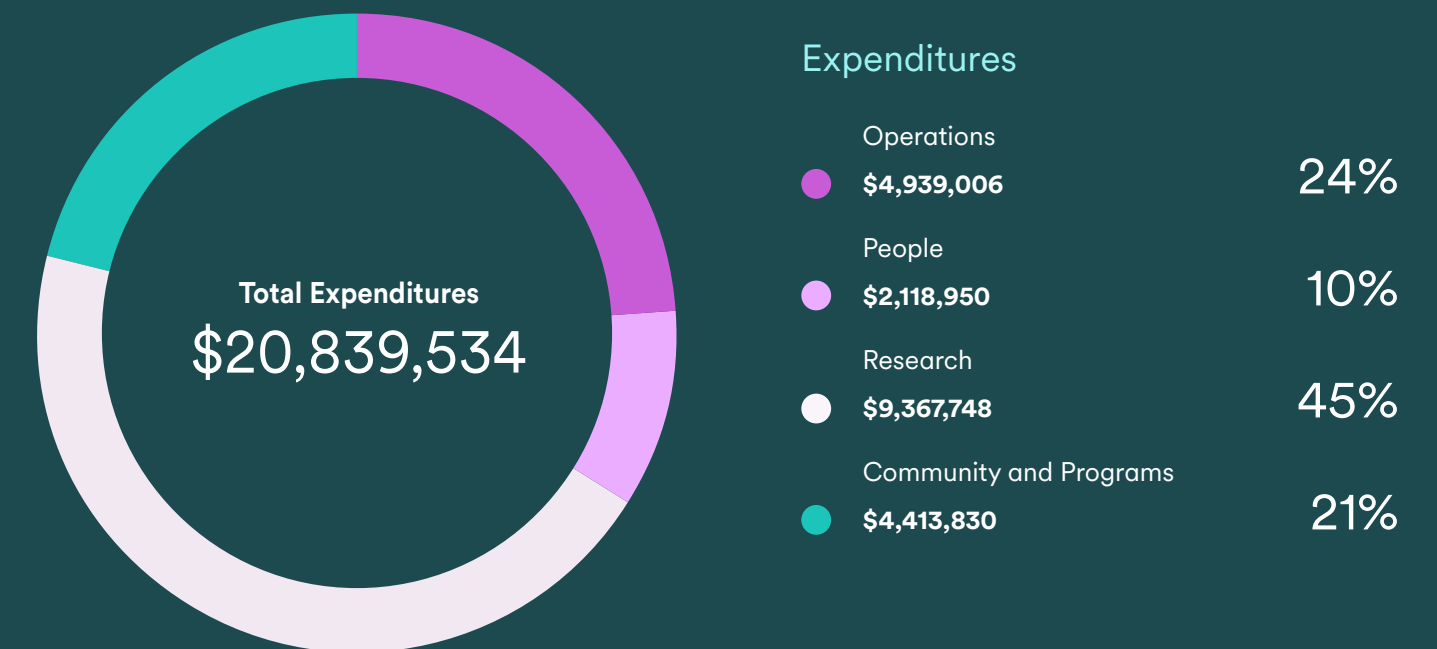**Total Issued in Research Grants**
$6,767,000

	Grant	Amount	Percent
●	Cloud Credit Grants	$2,000,000	30%
●	Hoffman-Yee Grants	$2,750,000	40%
●	Seed Grants	$1,767,000	26%
●	Wu-Tsai Grants	$250,000	4%

Financials in 2023

Sources of Income[1]

	Source	Amount	Percent
●	Gifts and Other Income	$64,884,932	84%
●	Endowment Payouts	$9,038,018	12%
●	Sponsored Research	$3,047,434	4%

1 Income does not include increases in endowment principal

Total Income
$76,970,385

Total Expenditures
$20,839,534

Expenditures

	Category	Amount	Percent
●	Operations	$4,939,006	24%
●	People	$2,118,950	10%
●	Research	$9,367,748	45%
●	Community and Programs	$4,413,830	21%

Thank You

Thank you for following Stanford HAI's efforts to realize our vision of human-centered AI. We appreciate the support of our extended community 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

Follow us on social media
@StanfordHAI

HAI Stanford University
Human-Centered
Artificial Intelligence

Stanford
University