



Stanford University
Human-Centered
Artificial Intelligence

2025

ANNUAL REPORT

LETTER FROM LEADERSHIP

The 2024-25 academic year was a period of significant growth and impact for Stanford HAI. Our work to advance human-centered AI reached new audiences and informed critical decisions around the world, strengthening our role as a lighthouse at a time when clarity about the technology has never been more important.

Government leaders, industry executives, educators, and civil society organizations have come to HAI for guidance — not just technical insight, but what AI means for their communities and institutions. Our faculty and fellows shared evidence-based perspectives with leaders at venues ranging from the United Nations Security Council and the World Economic Forum to state capitols and classrooms, helping them cut through the hype and place society’s well-being at the forefront of their choices.

Universities remain one of the few places where AI can be studied and built in ways that are open, accountable, and guided by long-term public needs. At HAI, we pursue this mission through interdisciplinary research, education, policy engagement, and partnerships that center human outcomes and ensure that AI serves society as a whole, not just those who develop, deploy, and sell it. Highlights from the past year include:

- **Advancing interdisciplinary research and thought leadership:** We built and studied powerful AI systems in the open; created public resources such as the AI Index, shared evaluation platforms, and transparency benchmarks; and contributed to the consequential California Report on Frontier AI Policy.
- **Extending HAI’s global reach and impact:** We developed international research collaborations; delivered executive education programs in Abu Dhabi and Singapore; and engaged with world leaders and heads of state from Davos to the UN Security Council and the French AI Action Summit.
- **Deepening HAI’s educational impact from K–12 to lifelong learning:** We created AI literacy resources for K–12 educators; expanded learning opportunities for Stanford students; and scaled human-centered AI education for global business and public-sector leaders through a robust executive education and policymaker training portfolio.
- **Integrating the Stanford Artificial Intelligence Lab (SAIL) into HAI:** We strengthened Stanford’s historic leadership in AI through this important partnership, opening new pathways for collaboration in research and education across the university.
- **Added valuable voices to our core faculty:** In 2025 we recruited Dr. Yejin Choi to HAI. Yejin’s pioneering work in language models, pluralistic values, and common-sense reasoning deepens our faculty bench and amplifies our global voice on the future of foundation models.
- **Expanding industry partnerships:** We funded more than 50 new research projects through our industry affiliate program and launched executive education offerings that brought human-centered AI training to leaders around the world.

In the year ahead, we will undertake the most significant structural change since HAI was founded by joining forces with Stanford Data Science as a single institute. Uniting human-centered AI and data-driven discovery will bring together the people who develop new algorithms, design AI systems and applications, build their infrastructure, study their impacts, and shape their governance. This integration will sharpen Stanford’s ability to help society make informed choices about the future of machine intelligence with greater clarity, evidence, and public purpose.

We are deeply grateful for the support of our extended HAI community — including our faculty, affiliates, students, donors, advisory council members and, above all, the staff. This work is only possible because of you.

Fei-Fei Li,
Denning Co-Director,
Stanford HAI

James Landay,
Denning Co-Director,
Stanford HAI

John Etchemendy,
Denning Co-Director,
Stanford HAI

Russell Wald,
Executive Director,
Stanford HAI

CONTENTS

- 00. Letter from the Leadership Team 2
- 01. Research Focus 3
- 02. Research & Education 11
- 03. Policy & Society 15
- 04. HAI Research Centers, Labs, and Collaborators 19
- 05. Industry Programs 25
- 06. People 27
- 07. HAI in the News 31
- 08. Financial Overview 33

RESEARCH FOCUS

OUR VISION FOR THE FUTURE IS LED BY THE COMMITMENT TO PROMOTE HUMAN-CENTERED USES OF AI, AND TO ENSURE THAT HUMANITY BENEFITS FROM THE TECHNOLOGY AND THAT THE BENEFITS ARE BROADLY SHARED.

In support of these goals, our research falls into three key focus areas:

HUMAN IMPACT



To develop equitable and trustworthy technology, we must understand how AI interacts with humans as well as with vital social structures and institutions.

AUGMENT HUMAN CAPABILITIES

HAI seeks to develop new human-centered design methods and tools so that AI agents and applications are designed and created with the ability to communicate with, collaborate with, and augment people more effectively, and to make their work better and more enjoyable. These breakthroughs will allow great progress in healthcare, education, sustainability, automation, and countless other domains.



INTELLIGENCE



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

HAI Co-Director James Landay presents HAI's research focus to Stanford faculty

📷: Jeanina Matias



YEAR AT A GLANCE

During the 2024-2025 Academic Year, HAI...

Distributed more than

\$12.8m

in research funding to faculty in all of Stanford's seven schools.

Designed and operationalized more than

65 new research projects benefiting over **40** Stanford faculty.

Developed in-person and virtual executive education programs for more than

3k participants and hosted over **75** industry visits on campus.

Generated new heights of engagement from the 2025 AI Index Report through **467,851** website views and **114,497** report downloads by September 2025, a

33.87% 

increase from the previous year's report viewership.

Educated more than **6k** U.S. federal, state, and local policymakers through our partnership with the U.S. General Services Administration, and another **1,200** global policymakers through a partnership with Apolitical.

Published **21** impactful policy research outputs, including nine policy briefs, two issue briefs, two white papers, and four responses to federal requests for comment. These publications involved collaborations with more than **70** faculty members and scholars across Stanford.

Secured

\$22m

for research projects and industry collaborations through our industry programs.

Engaged with **21** California executive and legislative offices,

31 U.S. federal agencies and congressional offices, and

65 foreign governments and international organizations.

Additionally...

HAI's Center for Research on Foundation Models (CRFM) evaluated **48** visual language models on **35** datasets, **17** audio language models on **20** datasets, and **78** large language models on **5** safety datasets.

HAI's Digital Economy Lab (DEL) launched the Digitalist Papers Volume I featuring a dozen original essays, and hosted more than **300** attendees at the launch event.

DEL shared preliminary results of a survey of more than **10,000** U.S. citizens as part of a research project to understand which goods and innovations contribute the most to economic growth and well-being.

Conference participant interacting with a robotic device created in the Stanford Robotics Center
 📷 : Christine Baker



SPOTLIGHT

ROBOTICS IN A HUMAN-CENTERED WORLD: INNOVATIONS AND IMPLICATIONS

As the field of robotics continues to evolve at an unprecedented pace, HAI's Spring Symposium explored the latest advances, challenges, and societal implications of robotic technology at the intersection of artificial intelligence. The conference fostered a deeper understanding of how researchers, engineers, social scientists, and policymakers can work together to navigate the evolving landscape of robotics. HAI closely collaborated with the Stanford Robotics Center to offer interactive robotic demonstrations.

“
 Unlike purely digital AI, robots interact physically with people, making user experience, safety, and ethical considerations impossible to ignore. I encourage [everyone] to not just think about what AI-powered robots can do, but how they should be designed to enhance human potential, align with our values, and distribute their benefits to society.”

—James Landay; Denning Co-Director, Stanford HAI | Anand Rajaraman and Venky Harinarayan Professor of Computer Science, Stanford University



HAI ON THE GLOBAL STAGE

Over the past year, HAI took significant steps to establish the institute internationally through our policy, education, and industry research and collaborations.

 In timeline order: FEB 2025: Paris AI Action Summit | MAR 2025: Jeanina Matias | JUN 2025: The European Union | JUL 2025: Accenture Singapore | SEP 2025: UN Photo/Eric Kanalstein

DEC 2024

HAI Co-Director Fei-Fei Li gave a [briefing to the United Nations Security Council](#).

JAN 2025

HAI Co-Director James Landay, HAI Senior Fellows Erik Brynjolfsson and Yejin Choi, HAI Center Fellow Alex “Sandy” Pentland, and HAI faculty Andrew Ng joined world leaders at the [World Economic Forum](#) in Davos.

MAR 2025



Fei-Fei Li, James Landay, and Russell Wald met with Lawrence Wong, Singapore’s prime minister, for a valuable conversation on advancing human-centered AI from innovation to governance and global collaboration.

MAY 2025

HAI’s Industry Programs team held the first Accenture Executive Education program in Abu Dhabi with James Landay and HAI faculty Diyi Yang and Michael Bernstein.

MAY 2025

HAI faculty Yejin Choi, Jiajun Wu, and Sanmi Koyejo spoke at the Asia TechX summit in Singapore, and HAI co-hosted a workshop on AI cultural values with AI Singapore.

JUN 2025



HAI Senior Fellow Susan Athey chaired a discussion with European Commission Executive Vice President Henna Virkkunen to outline [Europe’s AI Future](#).

JUL 2025

HAI’s Industry Programs team held the second Accenture Executive Education program in Singapore with James Landay, Jiajun Wu, Diyi Yang, and Russell Wald.

APR 2025

James Landay keynoted the FIND Financial Insight Summit conference in South Korea, then met with industry and academic leaders from Seoul National University Hospital, 12Labs, HD Hyundai, Samsung Research, and Hanwha Life to share his vision of human-centered AI.

SEP 2025

Yejin Choi [testified](#) to the United Nations Security Council on maintaining international peace and security, calling on the global scientific and policy communities to rethink how AI is built to ensure it benefits all of humanity.



FEB 2025

Fei-Fei Li delivered the keynote address at the [AI Action Summit](#) in Paris, while HAI Senior Fellows Erik Brynjolfsson and Rob Reich, HAI Executive Director Russell Wald, and HAI Policy Fellow Riana Pfefferkorn joined other summit week events and discussions.



GLOBAL PROJECT HIGHLIGHTS

LAUNCH OF THE GLOBAL AI VIBRANCY TOOL

One of the most comprehensive indices of AI vibrancy globally, this interactive tool offers valuable insights for understanding and fostering AI development. By providing a transparent assessment of national progress in AI across 36 countries, it serves the diverse needs of policymakers, industry leaders, researchers, and the general public.

36

countries assessed in the launch of the Global AI Vibrancy Tool

POLICY EDUCATION

HAI educated 1,234 policymakers from 80 different countries via free online training in partnership with Apolitical. The course, "AI Fundamentals for Public Servants: Opportunities, Risks and Strategies," covers topics ranging from risks and regulations to ethical AI and the future of the technology landscape.

1,234

policymakers educated from

80

different countries

AI GOVERNANCE

HAI has expanded our role as a global thought leader and convener for discussions about AI governance, spotlighting key issues of low-resource language development and AI talent flows. By partnering with international institutions and hosting policymakers from the United Nations, Europe, Asia, and Latin America, HAI is advancing cross-cultural approaches and positioning Stanford as a hub for global

dialogue. The policy white paper, [*Mind the \(Language\) Gap: Mapping the Challenges of LLM Development in Low-Resource Language Contexts*](#), highlights challenges, trade-offs, and strategies to increase investment; prioritize cross-disciplinary, community-driven development; and ensure fair data ownership.

Cover Illustration for *Mind the (Language) Gap: Mapping the Challenges of LLM Development in Low-Resource Language Contexts*

Credit: Michi Turner



2 RESEARCH & EDUCATION

HAI HAS CONTINUED TO MAKE SIGNIFICANT STRIDES IN OUR FOCUS AREAS OF RESEARCH AND EDUCATION BY FOSTERING A CULTURE OF INTERDISCIPLINARY COLLABORATION AND ENABLING THE TOP MINDS IN AI TO STUDY, DEVELOP, AND PROMOTE HUMAN-CENTERED AI THAT IS DESIGNED TO AUGMENT HUMAN CAPABILITIES.

EDUCATION

HAI provided educational programs to empower emerging leaders and decision-makers with cutting-edge research, tools, and resources to help them address the societal and ethical implications of AI. Programs included executive education for professionals, educational programs for policymakers and civil society leaders, AI literacy and career-readiness programs for K-12 educators and students, and undergraduate programs for Stanford students. HAI Education Fellow Peter Norvig co-led [CS139: Human-Centered AI](#), which explores AI's ethical issues, how people understand AI systems, and how to apply those insights to designing AI that is trustworthy, fair, and beneficial. The course is open to undergraduate and graduate students.

GRADUATE AND POSTDOCTORAL FELLOWS

HAI hosted a cohort of 29 scholars representing a diverse range of research, covering topics including education data science, digital health innovation, AI safety and ethical development, energy science, psychology geophysics, AI literacy, and precision medicine.

STUDENT AFFINITY GROUPS

HAI supported 12 teams of interdisciplinary students who worked on a variety of projects, such as "AI and Social Services," "Global AI Governance," and "An Exploration of AI Art." One group hosted a post-AI Action Summit fireside chat with Erik Brynjolfsson, director of the Stanford Digital Economy Lab, and Florence G'Sell, visiting professor of private law at the Stanford Cyber Policy Center.



 Christine Baker

RESEARCH PROJECT HIGHLIGHTS

Participants in discussion at the HAI Seed Grant Program kick-off workshop.

📷 : Jeanina Matias



SEED GRANT

Supported by the HAI Seed Grant Program and a partnership grant with the Stanford Accelerator for Learning, Assistant Professor of Education Data Science Dora Demszky and a team of education and computer science scholars developed an AI-powered platform to extend standard middle-school math curricula into personalized learning

experiences. [Coteach.ai](#) generates curriculum-aligned educational scaffolds like tailored warmups, additional practice questions, and worked examples. The team's platform won the 2024 Educational Tools Competition as one of 1,900 submissions and received \$400,000 in follow-on funding from the Gates Foundation.

GENERATIVE AI: TECHNOLOGY, BUSINESS, AND SOCIETY PROGRAM

Developed with Stanford Online, this program examines the technical, business, and societal impacts of AI with a people-first focus. As generative AI reshapes industries and creative processes, this course prepares business leaders to

grasp the implications of this transformation for their organizations, technologies, and society at large. This year, over 1,200 leaders enrolled in this course, rating it on average an 8 out of 10.

K12 PROGRAM

HAI's commitment to supporting AI literacy and career readiness in the field of AI is reflected through the AI4ALL program for rising 10th-grade students. Last year, the program expanded to nearly 100 students from 11 countries around the world.

HAI Co-Director James Landay welcoming AI4ALL alumni participants to a reunion event

📷 : Jeanina Matias





Credit: Kevin Liotti

The annual AI Index Report is a go-to resource for policymakers, industry leaders, and the general public. The most recent report included in-depth analyses of the evolving landscape of AI hardware, novel estimates of inference costs, and new analyses of AI publication and patenting trends. It also featured an expanded chapter on science and medicine, developed by a team from [RAISE Health](#), a collaboration between Stanford Medicine and HAI.

The 2025 AI Index achieved wide-reaching engagement. Within its first week of release, its findings were covered in 440 news stories, including top-tier outlets like [Forbes](#), [IEEE Spectrum](#), [The Information](#), [Nature](#), [Politico](#), and [Wired](#), as well as an exclusive broadcast segment with [Fox & Friends](#).

Briefing requests came in from policymakers, media professionals, and major tech companies, resulting in dozens of meetings from the AI Index Steering Committee members to discuss key insights from the Index’s findings.

Covered in more than

440 News Stories



In collaboration with Microsoft, LinkedIn, and the Council on Foreign Relations, HAI hosted an AI Index [workshop](#) focused on implications for the workforce and beyond. HAI also



(From left to right) Sebastian Elbaum, Yolanda Gil, James Landay, and Vijay V. Vaitheeswaran during a panel discussion at the Artificial Intelligence in 2025 workshop.

met with industry leaders at JPMorganChase in New York to discuss insights from the 2025 report and perspectives from industry. To educate policymakers, HAI held AI Index roll-out events in Sacramento with California policymakers and in San Francisco at the British Consulate.

Three related HAI-published articles broadened awareness of the AI Index by translating its key findings into clear, public-facing insights, including [AI Index 2025: The State of AI in 10 Charts](#).



The AI Index is not just a report, it’s a compass for policymakers, researchers, and the public to navigate the future of AI from a place of knowledge. As a trusted source, it provides the foundation for informed decisions and further exploration.

– Vanessa Parli, Managing Director of Programs and External Engagement, Stanford HAI



📷 : Christine Baker

3 POLICY & SOCIETY

HAI Co-Director Fei-Fei Li
briefing the United Nations
Security Council

📷 : UN Photo/Loey Felipe

In the 2024–25 academic year, HAI advanced human-centered AI governance from concept to practice by pairing interdisciplinary scholarship with direct engagement of policymakers and civil society leaders. From California to Washington, D.C., and on the international stage, HAI informed critical debates, trained government leaders, and convened diverse stakeholders to shape how AI is developed and governed. These efforts elevated HAI as a trusted, authoritative hub for policy insight and a global leader in ensuring that AI serves society.



HEALTHCARE AI POLICY

📷 : Tracy Navichoque

HAI's Healthcare AI Policy Steering Committee is a group of academics, physicians, lawyers, computer scientists, and ethicists who see an urgent need to review existing regulatory frameworks to ensure AI tools are safe, fair, and secure for clinical use. In June 2025, the committee held its Healthcare AI Workshop with 60 technologists, medical professionals, and policy experts to discuss integrating patient perspectives into the development of healthcare AI.

HAI also published two healthcare-related policy briefs, *The Complexities of Race Adjustment in Health Algorithms* and *Increasing Fairness in Medicare Payment Algorithms*.



SOCIAL SECTOR AI PROGRAM

Participants in discussion during a breakout session in the Social Sector AI Program.

📷 : Patrick Beaudoin



HAI hosted an inaugural cohort of 32 social sector practitioners for a two-day workshop on advancing human-centered AI in the social and education sectors.

CONGRESSIONAL BOOT CAMP ON AI



(From left-to-right) Andrew Grotto, Colin Kahl, Amy Zegart, and Christopher Manning speaking at the Congressional Boot Camp on AI during the "AI and National Security"

📷 : David Gonzales

25 congressional staffers participated in the fifth annual boot camp, along with 43 faculty and scholars. HAI provided lawmakers with in-depth, bipartisan training on emerging AI policy issues, giving participants a conceptual framework to address the current technology landscape and better anticipate the challenges of tomorrow.



The program met my expectations and went beyond, offering in-depth insights into human-centered design principles, ethical considerations, and practical AI applications. I really appreciated the interdisciplinary approach.

– Lieselot Declercq, D-teach

The speakers shared great examples that helped expand our horizons and reflect on what's possible. The program was geared to navigating a specific project or need, which made it very helpful and practical.

– Michelle Rojas-Soto, KIPP SoCal Public Schools

The workshop was insightful in showing how AI can support the public sector, proving it to be a valuable tool rather than something to fear.

– Lily Fraser, ThinkData Ed

TRACKING, ANALYZING, AND INFORMING U.S. EXECUTIVE ACTION ON AI

HAI scholars submitted [initial input](#) on drafting the new U.S. AI Action Plan and analyzed the published plan. The policy team also analyzed the prior administration's implementation of executive AI initiatives and responded to four federal agency requests for comment.

Through critical scholarship and global dialogue, we inform human-centered AI governance and engage government, industry, and civil society leaders on emerging policy issues. Our efforts also center on empowering underrepresented communities to shape the design, development, and oversight of AI technologies. In the 2024–25 academic year, HAI reinforced our role as a trusted adviser in informing California’s AI policy landscape and invested in cultivating the next generation of AI policy leaders by placing students in government and civil society organizations through the Tech Ethics & Policy Summer Fellowships.



📷 : SF Photo Agency

TECH ETHICS & POLICY (TEP) SUMMER FELLOWSHIPS

HAI cultivates the next generation of AI policy leaders by embedding students in government and civil society. TEP Fellows have advanced legislation, tested AI safety practices, and supported Congress in improving operations. The program has built a pipeline of leaders who bridge technical expertise and policy responsibility.

Last summer, HAI placed eight students in Washington, D.C. Some snapshots:

- **Alberto Tono** drafted bills for U.S. Senator Todd Young and helped provide comments on the White House AI Action Plan.
- **Evelyn Yee** led an extensive red-teaming project for Georgetown’s Center for Security and Emerging Technology.
- **Clara Yi** helped the Senate Appropriations Committee identify opportunities to use AI to streamline the appropriations process.

TEP FELLOW PLACEMENTS

- US Senator Todd Young (R-IN)
- Senate Appropriations Committee
- Senate Homeland Security Committee Minority
- American Enterprise Institute
- Brookings Institution
- Center for Democracy and Technology
- Georgetown Center for Security and Emerging Technology

INFORMING CALIFORNIA’S AI POLICY LEADERSHIP

As California positions itself at the forefront of AI innovation and governance, HAI has contributed research, training, and public dialogue that inform the state’s approach. From analyzing frontier model governance to developing tools that help reduce regulatory barriers, HAI connects scholarship to practice and fosters evidence-driven discussions that can guide policymaking in California and beyond.

The policy brief “Cleaning Up Policy Sludge: An AI Statutory Research System” explained work from HAI Senior Fellow Daniel Ho and researchers at the Regulation, Evaluation and Governance Lab (RegLab). In an exciting real-world application, Professor Ho worked with the city of San Francisco to implement the statutory research system, as highlighted in Politico.

For example, at the request of California Governor Gavin Newsom and as part of the Joint California Policy Working Group on AI Frontier Models, HAI researchers contributed to the [California Report on Frontier AI Policy](#), and they published “Advancing Science- and Evidence-Based AI Policy” in [Science](#). The report was a reference point for policy debates in Sacramento that ultimately led to the successful passage of SB-53, California’s first significant frontier AI legislation.

The California State Boot Camp on AI, held in conjunction with the Stanford Institute for Economic Policy Research and RegLab, brought 44 state executive and legislative branch staffers to campus for a daylong, immersive education event. HAI also held AI Index rollout events in Sacramento and San Francisco, and produced a policy highlights document to accompany the AI Index report.

Participants at the California State Boot Camp on AI.

📷 : Ryan Zhang



4 RESEARCH CENTERS, LABS & COLLABORATORS

HAI's research centers, labs, and collaborators are all critical in achieving our mission.

- **Center for AI Safety**
- **Center for Research on Foundation Models** (Spotlighted on page 22)
- **Golub Capital Social Impact Lab**
- **One Hundred Year Study of Artificial Intelligence**
- **Open Virtual Assistant Lab**
- **RAISE Health**
- **Regulation, Evaluation, and Governance Lab**
- **Stanford Artificial Intelligence Lab**
- **Stanford Digital Economy Lab** (Spotlighted on pages 20-21)
- **Stanford Emerging Technology Review**

STANFORD DIGITAL ECONOMY LAB

Led by Erik Brynjolfsson and Christie Ko



📷 : Christine Baker

The Stanford Digital Economy Lab (DEL) aims to build the intellectual and technical infrastructure for an inclusive digital economy—advancing research, measurement, and policy models that empower people. In the 2024-25 academic year, DEL focused on developing the Economics of Transformative AI agenda, launching public-facing scholarship, piloting AI-enabled civic technologies, and producing new measures of well-being and user value.

The team launched [The Digitalist Papers](#), volume I, to national attention, secured a National Bureau of Economic Research (NBER) special volume on the Economics of Transformative AI, and advanced the [research agenda](#) for this project with a completed white paper draft and an economist roundtable to shape next steps.

DEL also moved core platforms from concept to pilots: The lab upgraded and piloted an AI-

assisted civic deliberation platform called [Deliberation.io](#) that features preference visualization and refined voting modules. External testing began in Spring 2025 with partners in Washington, D.C., and Ghana to evaluate how LLMs can support consensus-building and diverse perspective-sharing. Meanwhile, DEL's Measuring the Digital Economy program completed a 10,000-person U.S. survey to determine which goods and innovations contribute the most to economic growth and well-being, with early results shared at a March 2025 workshop.

Across the portfolio, DEL combined rigorous measurement, platform prototyping, and policy engagement to inform better governance and broaden economic opportunity in an AI-driven economy.

LOYAL AGENTS

In early 2025, DEL collaborated with Consumer Reports to convene 15 leaders from academia, industry, and policy to launch a coalition developing legal and technical foundations for trustworthy, user-authorized AI agents. The initiative is defining standards, prototyping agent systems, and building evaluation toolkits to align agentic AI with consumer interests and platform accountability.



Tobin South, DEL Research Fellow, leads a workshop on agentic AI at the Loyal Agents convening.

ARTIFICIAL INTELLIGENCE AND THE FUTURE OF WORK

In November 2024, the National Academies released a consensus report co-chaired by DEL Director Erik Brynjolfsson and Visiting Scholar Tom Mitchell, which evaluated the impacts of AI on productivity, employment, and skills. The report concludes that AI can enhance human labor and create new forms of valuable work—but this outcome is not guaranteed. It recommends that the United States build the capacity to rapidly detect and respond to AI developments and to track workforce impacts in real time, enabling adaptive policy, education, and training. This work provides a practical framework for employers, educators, and policymakers to align innovation with broad-based opportunity.

ECONOMICS OF TRANSFORMATIVE AI



DEL launched a two-week course for 40 Ph.D students from the world's top university economics programs designed to provide a deep dive into the economics of transformative AI. The course received high marks from participants.

Participants of the ETAI workshop at the closing event
: Elsa Conde

CENTER FOR RESEARCH ON FOUNDATION MODELS

Led by Percy Liang



: Christine Baker

In the 2024-25 academic year, the Center for Research on Foundation Models (CRFM) pushed for greater impact on both the research world and in broader society:

- On the technical front, the center expanded [Holistic Evaluation of Language Models \(HELM\)](#), its foundation model evaluation project, to support multimodal models and to conduct comprehensive safety evaluations.
- On the policy front, CRFM continued its work on evidence-based AI policy by articulating a clear vision for the field that was published in [Science](#), along with direct impact on global policy in multiple jurisdictions through the [California Report on Frontier AI Policy](#), the EU AI Act [Code of Practice](#), and the [International Scientific Report on Advanced AI](#).

MARIN

CRFM established its third large ongoing initiative with this open lab for building foundation models to realize a vision of truly open-source AI. Marin is now included on the HELM leaderboard platform, as well as in the Foundation Model Transparency Index ([FMTI](#)), a composite index for developer transparency.

AI POLICY COLLABORATION

Scholars from Stanford CRFM and HAI collaborated with academic colleagues from UC Berkeley, the Carnegie Endowment for International Peace, Harvard, Princeton, and others to publish “Advancing Science- and Evidence-based AI Policy” in [Science](#). The paper presents a clear vision for the future of AI policy and governance by identifying how policymakers can actively catalyze the growth of the evidence base through implementing evidence-generating policies.

HELM SAFETY

CRFM introduced a new public leaderboard for transparent, standardized, and reproducible safety evaluation of language models, which addresses deficits in mainstream evaluations that only focus on language model capabilities.

MULTIMODAL MODELS

The CRFM evaluation team built new holistic and reproducible evaluations for multimodal models: [VHELM](#) for vision language models and [AHELM](#) for audio language models. These new evaluations help fill the gap in public leaderboards for critical modalities that otherwise lack evaluation infrastructure comparable to what has become customary for text-to-text language models.

ADDITIONAL COLLABORATORS

In addition to CRFM and DEL, HAI encompasses several other centers and labs, all of which are a part of the institute and critical partners in achieving our mission.

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.

GOLUB CAPITAL SOCIAL IMPACT LAB

The lab focuses on developing, applying, and testing cutting-edge digital tools to improve the effectiveness of the social sector and catalyze innovation directed at social challenges.

ONE HUNDRED YEAR STUDY OF ARTIFICIAL INTELLIGENCE (AI100)

The One Hundred Year Study of Artificial Intelligence (AI100) is a longitudinal study to evaluate and anticipate how the effects of artificial intelligence will ripple through every aspect of how people work, live, and play.

OPEN VIRTUAL ASSISTANT LAB

The lab builds LLM-augmented cognition systems (LLM-CSys) that combine the expressive language skills of LLMs with the interpretability and reliability of software systems, as well as connecting LLMs with databases, compilers, and executive functions to create assistants that are more trustworthy and helpful.

RAISE HEALTH

RAISE Health is a joint initiative between Stanford Medicine and HAI to guide the responsible use of AI across biomedical research, education, and patient care.

REGULATION, EVALUATION, AND GOVERNANCE LAB (REGLAB)

Stanford's RegLab builds the evidence base and technology for effective governance.

STANFORD ARTIFICIAL INTELLIGENCE LAB (SAIL)

SAIL has been a center of excellence for artificial intelligence research, teaching, theory, and practice since its founding in 1963.

STANFORD EMERGING TECHNOLOGY REVIEW

The Stanford Emerging Technology Review helps America's public and private sectors better understand transformational technologies so that the United States can seize opportunities, mitigate risks, and ensure its innovation ecosystem continues to thrive.

“ UNIVERSITIES REMAIN ONE OF THE FEW PLACES WHERE AI CAN BE STUDIED AND BUILT IN WAYS THAT ARE OPEN, ACCOUNTABLE, AND GUIDED BY LONG-TERM PUBLIC NEEDS.

— HAI LEADERSHIP: FEI-FEI LI, JAMES LANDAY, JOHN ETCHEMENDY & RUSSELL WALD

INDUSTRY PROGRAMS

HAI BUILDS MEANINGFUL COLLABORATIONS BETWEEN STANFORD RESEARCHERS AND INDUSTRY LEADERS, TRANSLATING STANFORD'S AI EXPERTISE INTO REAL-WORLD SOLUTIONS THAT ADVANCE RESPONSIBLE INNOVATION AND CREATE LASTING SOCIETAL IMPACT.

In the 2024-25 academic year, HAI focused on long-term institutional memberships while growing our executive education offerings and cultivating targeted gifts. Through these efforts, HAI has continued to prove the ability to sustain industry trust, deliver meaningful value via research collaborations, and position the institute as the premier academic partner for human-centered AI. Notable accomplishments include:

HAI INDUSTRY MEMBERS AND SUPPORT



AFFILIATE MEMBERSHIPS

HAI funded more than 50 new research projects. These projects benefitted more than 40 Stanford faculty members in many areas of applied AI research, including AI/LLM applications in HCI, UX, fintech, insurance, and marketing. At the first HAI annual affiliates meeting, members came together to learn more about our programs, exchange ideas, and hear directly from HAI scholars Michael Bernstein, Erik Brynjolfsson, James Landay, Sandy Pentland, and Melissa Valentine on exciting new research.

NEW COLLABORATIONS

HAI launched the Self-improving, Efficient and Accelerated Models and Systems ([SEAMS](#)) affiliate program, with IBM and Felicis as members, and integrated three new members through the collaboration with the Stanford AI Lab ([SAIL](#)).

EXECUTIVE EDUCATION, ON CAMPUS AND ABROAD

HAI welcomed 475 executives and leaders from over 70 companies to semi-monthly executive education programs and hosted 75 industry visits on campus.

OFF-CAMPUS ENGAGEMENT

HAI hosted human-centered research meetings at several conferences, including the International Conference on Machine Learning (ICML), the Conference on Neural Information Processing Systems (NeurIPS), and a workshop at the International Conference on Learning Representations (ICLR) 2025.

HAI's Industry Programs bridge academia and industry to advance the responsible development and adoption of AI. Through executive education, research collaborations, and strategic partnerships, HAI equips global leaders to apply human-centered principles to innovation and governance.

GLOBAL EXECUTIVE EDUCATION PROGRAMS

In the 2024-25 academic year, HAI's global executive education reach expanded, delivering two tailored international programs for HAI affiliate member Accenture's clients in Abu Dhabi and Singapore.

In May 2025, HAI delivered the UAE Foundation Model Scholar Program in Abu Dhabi. Designed for UAE executives, the program combined hands-on training with a Stanford-Accenture joint certificate and a focus on how leaders can scope, staff, and scale responsible generative-AI initiatives.

Core sessions covered highlights from the 2025 AI Index, a pragmatic introduction to foundation models and generative agents from HAI Senior Fellow Michael Bernstein, Responsible AI from Stanford researcher Anka Reuel, Human-Centered AI from HAI Co-Director James Landay, Opportunities & Risks of Generative AI and Human-AI Interaction from HAI faculty Diyi Yang, value-led AI transformation and AI adoption/case studies with industry practitioners, and a capstone panel and graduation ceremony.



TECHNICAL PAPERS

HAI's affiliate program supported HAI scholars on the following technical papers and projects:

[Attention Factors for Statistical Arbitrage](#), by Markus Pelger



MARKUS PELGER
Associate Professor of Management Science & Engineering, Stanford University; Chambers Faculty Scholar, School of Engineering | Faculty Affiliate, Stanford HAI

[An Adaptive Responsible AI Governance Framework for Decentralized Organizations](#), by Mykel Kochenderfer



MYKEL KOCHENDERFER
Associate Professor of Aeronautics and Astronautics, Stanford University | Senior Fellow, Stanford HAI | Director, Stanford Intelligent Systems Laboratory (SISL)

[Generative Agents: Interactive Simulacra of Human Behavior](#), by Michael Bernstein



MICHAEL BERNSTEIN
Associate Professor of Computer Science | Senior Fellow, HAI | STMicroelectronics Faculty Scholar, Stanford University

[Collaborative Gym: A Framework for Enabling and Evaluating Human-Agent Collaboration](#), by Diyi Yang



DIYI YANG
Assistant Professor, Computer Science Department, Stanford University | Faculty Affiliate, Stanford HAI

FACULTY SPOTLIGHT



DIYI YANG has been a key contributor to HAI's Industry Program activities over the last year. As a go-to, top-rated voice across multiple executive education programs—domestic and international—she bridges cutting-edge human-computer interaction (HCI) and natural language processing (NLP) with real-world impact. Her flagship work on AI agents and socially aware, human-centered language technologies—covering safety, fairness, and collaborative AI—has been catalyzed by research tokens from the HAI affiliate program. These collaborations have seeded new pilots, informed partner playbooks, and drawn fresh practitioners into our ecosystem.



PEOPLE

DENNING CO-DIRECTORS

The Stanford HAI leadership team helps the institute realize the vision of artificial intelligence that serves the collective needs of humanity.



JOHN ETCEMENDY

Denning Co-Director, Stanford HAI | Stanford Provost Emeritus | Patrick Suppes Family Professor in the School of Humanities and Sciences, Stanford University



JAMES LANDAY

Denning Co-Director, Stanford HAI | Anand Rajaraman and Venky Harinarayan Professor of Computer Science, Stanford University



FEI-FEI LI

Denning Co-Director, Stanford HAI | Sequoia Professor of Computer Science, Stanford University

ASSOCIATE DIRECTORS & SENIOR FELLOWS



RUSS ALTMAN

Kenneth Fong Professor and Professor of Bioengineering, of Genetics, of Medicine, of Biomedical Data Science, and Professor, by courtesy, of Computer Science, Stanford University



SURYA GANGULI

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



CHRISTOPHER MANNING

Thomas M. Siebel Professor of Machine Learning, Professor of Linguistics and of Computer Science, Stanford University



JULIAN NYARKO

Professor, Stanford Law; Center Fellow, Stanford Institute for Economic Policy Research (SIEPR), Stanford University



GE WANG

Associate Professor of Music and Associate Professor, by courtesy, of Computer Science, Stanford University



AMY ZEGART

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

SENIOR FELLOWS



SUSAN ATHEY

The Economics of Technology Professor, Stanford Graduate School of Business | Senior Fellow, Stanford HAI



MICHAEL S. BERNSTEIN

Associate Professor of Computer Science | Senior Fellow, Stanford HAI | STMicroelectronics Faculty Scholar, Stanford University



ERIK BRYNJOLFSSON

Jerry Yang and Akiko Yamazaki Professor | Senior Fellow, Stanford HAI | Senior Fellow, SIEPR | Professor, by courtesy, of Economics; of Operations, Information & Technology; and of Economics, Stanford Graduate School of Business



YEJIN CHOI

Dieter Schwarz Foundation HAI Professor | Professor of Computer Science | Senior Fellow, Stanford HAI



ANGÈLE CHRISTIN

Associate Professor of Communication and, by courtesy, of Sociology, Stanford University | Senior Fellow, Stanford HAI



MICHELE ELAM

William Robertson Coe Professor in the Humanities | Senior Fellow, Stanford HAI | Bass University Fellow in Undergraduate Education, Stanford University



CARLOS ERNESTO GUESTRIN

Professor of Computer Science, Stanford University | Senior Fellow, Stanford HAI



DANIEL E. HO

William Benjamin Scott and Luna M. Scott Professor of Law | Professor of Political Science and, by courtesy, of Computer Science | Senior Fellow, Stanford HAI | Senior Fellow, Stanford Institute for Economic and Policy Research | Director of the Regulation, Evaluation, and Governance Lab (RegLab)



MYKEL KOCHENDERFER

Associate Professor of Aeronautics and Astronautics, Stanford University | Senior Fellow, Stanford HAI | Director, Stanford Intelligent Systems Laboratory (SISL)



CURTIS LANGLOTZ

Senior Associate Vice Provost for Research | Professor of Radiology (Integrative Biomedical Imaging Informatics), of Medicine (Biomedical Informatics Research), of Biomedical Data Science | Senior Fellow, Stanford HAI



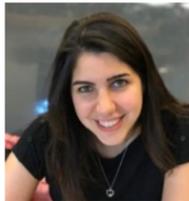
PERCY LIANG

Associate Professor of Computer Science, Stanford University | Director, Stanford Center for Research on Foundation Models | Senior Fellow, Stanford HAI



ROB REICH

McGregor-Girard Professor of Social Ethics of Science and Technology | Senior Fellow, Stanford HAI | Professor, by courtesy, of Education and of Philosophy | Senior Fellow, by courtesy, Freeman Spogli Institute for International Studies



DORSA SADIGH

Associate Professor of Computer Science and of Electrical Engineering, Stanford University | Senior Fellow, Stanford HAI



MELISSA VALENTINE

Associate Professor of Management Science and Engineering, Stanford University | Senior Fellow, Stanford HAI | Co-Director of Science of Work, Technology, and Organization (WTO), Stanford University

POLICY FELLOWS



RISHI BOMMASANI

Senior Research Scholar, Stanford HAI



JASON ELLIOTT

California Policy Fellow, Stanford HAI



JENNIFER KING

Privacy and Data Policy Fellow, Stanford HAI



ROHINI KOSOGLU

Policy Fellow, Stanford HAI



RIANA PFEFFERKORN

Policy Fellow, Stanford HAI



MARIETJE SCHAAAKE

International Policy Fellow, Stanford HAI

POSTDOCTORAL FELLOWS

STAFF LEADERSHIP

Ruyu Chen

Jennifer Chien

Duncan Eddy

Ziv Epstein

Basil Halperin

Andy Haupt

Wanheng Hu

Christina Langer

Connor Lawless

Jiaxin Pei

Justin Shin

Philip Trammell

Gabriel Unger

Luca Vendraminelli

Dian Wang

Lio Wong



RUSSELL WALD

Executive Director

Elena Cryst: Director of Policy and Society

Christie Ko: Executive Director, Stanford Digital Economy Lab

Anh Le: Director of Finance and Administration

Carolyn Lehman: Director of Media and Communications

Vanessa Parli: Managing Director of Programs and External Engagement

Ahmad Rushdi: Director of Industry Programs

Daniel Zhang: Chief of Staff

HAI IN THE NEWS

ASSOCIATED PRESS

[US Ahead in AI Innovation, Easily Surpassing China in Stanford's New Ranking](#)

(11/21/24 | Global AI Vibrancy Tool)

BLOOMBERG

[AI Seeks Out Racist Language in Property Deeds for Termination](#)

(10/17/25 | Daniel Ho)

FORBES

[Why Everyone — From Technologists to Creatives — Needs to Guide AI Design](#)

(1/11/25 | James Landay)

FOX & FRIENDS

(4/7/25 | Russell Wald)

IEEE SPECTRUM

[AI Godmother Fei-Fei Li Has a Vision for Computer Vision](#)

(12/12/24 | Fei-Fei Li)

KQED

[AI Could Soon Shop for You. Can We Trust It With Our Credit Cards?](#)

(5/2/25 | Jennifer King)

NATURE

[AI Race in 2025 Is Tighter Than Ever Before](#)

(4/7/25 | 2025 AI Index)

NEWSWEEK

[Entry-Level Jobs for Gen Z Are Disappearing: Experts](#)

(6/19/25 | Melissa Valentine)



NEW YORK TIMES

[AI Might Take Your Job. Here Are 22 New Ones It Could Give You.](#)

(6/17/25 | Erik Brynjolfsson)

NO PRIORS

[Ep. 117 | With Co-Director of Stanford's HAI & Founder of World Labs Dr. Fei-Fei Li](#)

(6/5/25 | Fei-Fei Li)

PBS FIRING LINE WITH MARGARET HOOVER

[Fei-Fei Li | Full Episode 5.23.25](#)

(5/23/25 | Fei-Fei Li)

PBS

[The Future of AI](#)

(1/2/25 | James Landay and Erik Brynjolfsson)

POLITICO

[San Francisco Wants to Use AI to Save Itself from Bureaucracy](#)

(6/5/25 | Daniel Ho)

TECH BREW

[Are Open-Source AI Models Worth the Risk?](#)

(10/31/24 | Rishi Bommasani)



TIME

[AI Will Devastate the Future of Work. But Only If We Let It](#)

(6/4/25 | Erik Brynjolfsson)



UPSTARTS

[What AI Startup Founders Thought About Stanford's 'Wake-Up Call' Report](#)

(4/10/25 | 2025 AI Index, Russell Wald, Vanessa Parli)

THE VERGE

[Making Human Music in an AI World](#)

(11/17/24 | Ge Wang)

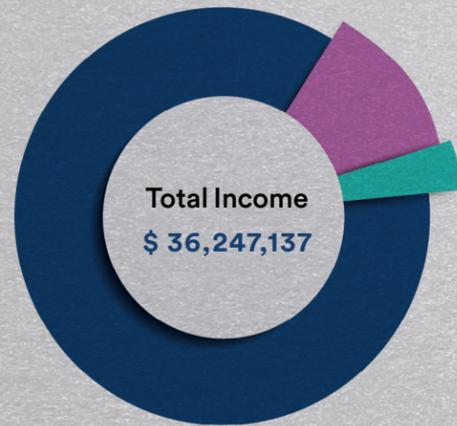
WASHINGTON POST

[23andMe's DNA Database Is Up for Sale. Who Might Want It, and What For?](#)

(3/25/25 | Jennifer King)

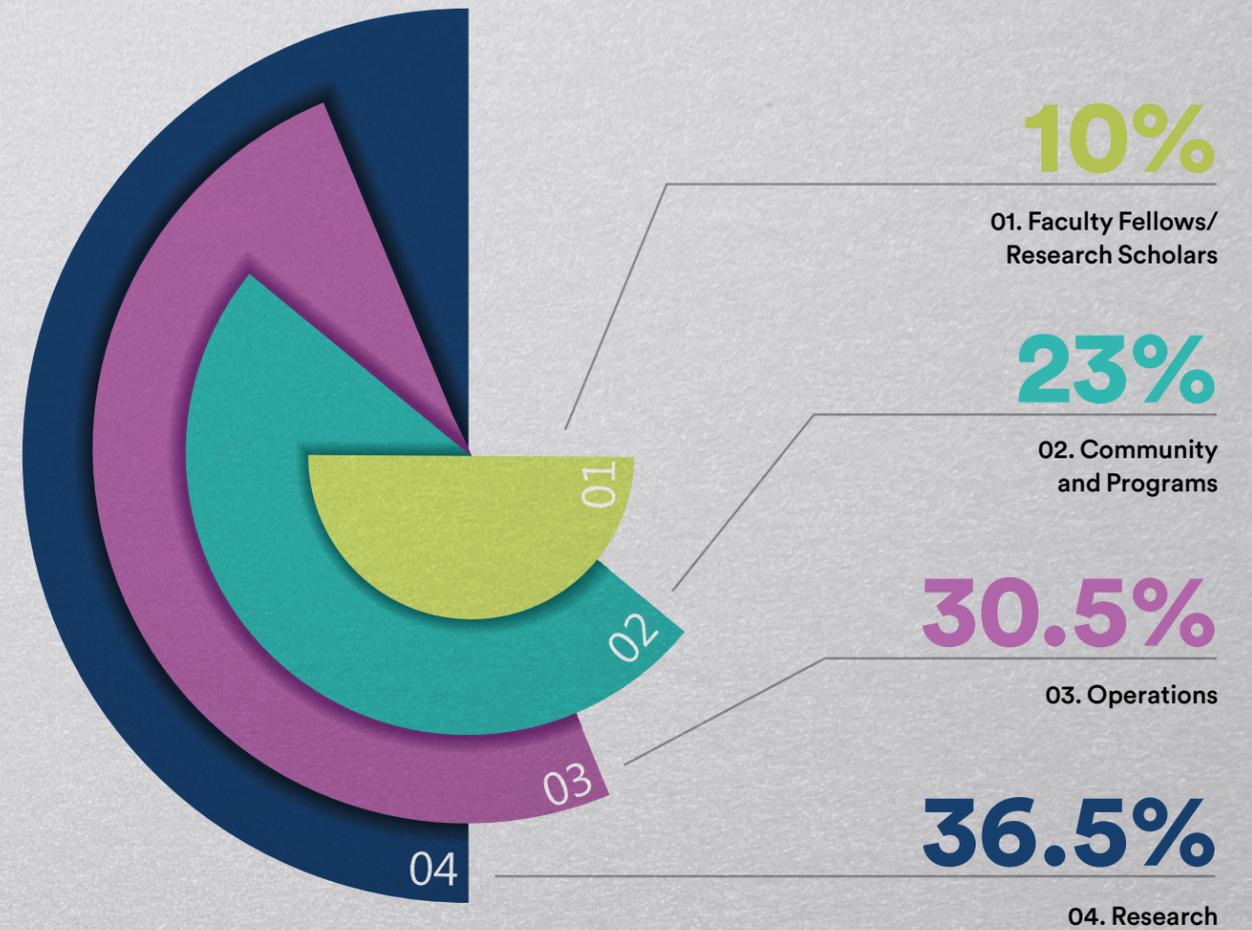
FINANCIAL OVERVIEW

SOURCE OF INCOME FY25



● Gifts and other income \$ 30,265,699	83.5%
● Endowment payouts \$ 4,623,323	12.5%
● Sponsored research \$ 1,358,114	4%

● Hoffman-Yee Grants \$ 3,000,000	53%
● Seed Grants \$ 2,285,649	40.5%
● Cloud Credit Grants \$ 373,050	6.5%



HAI FINANCIALS FY25

● Faculty Fellows/Research Scholars	\$ 3,517,466
● Community and Programs	\$ 7,902,950
● Operations	\$ 10,537,567
● Research	\$ 12,559,104

TOTAL **\$ 34,517,087**

\$12.8m

in funds given to
Stanford researchers

THANK YOU

Stanford HAI's mission is to advance AI research, education, policy, and practice to improve the human condition. We believe AI should be guided by its human impact, inspired by human intelligence, and designed to augment, not replace, people. Our interdisciplinary faculty conducts research focused on guiding the development of AI technologies intended to enhance human capabilities while ensuring its ethical, fair, and transparent use.

Contact us: hai-institute@stanford.edu

Learn more: hai.stanford.edu

