WELCOME

AI in the Liberal Arts Initiative
2025 Undergraduate Conference

The AILA Conference Planning Committee welcomes you to our annual AI in the Liberal Arts Undergraduate Conference! Since AILA's founding we have been eager to create a space for undergraduates in the Five Colleges and surrounding colleges and universities to gather and discuss AI in interdisciplinary ways. We are delighted to host you today at Amherst College's Aliki Perroti & Seth Frank Lyceum as we share ideas and discussions around Artificial Intelligence.

The presentations and workshop sessions we have planned offer an interdisciplinary look at the many ways undergraduate students are experiencing and thinking about Artificial Intelligence as it becomes increasingly present in our everyday lives and work. While some of our presenters are working on the frontlines of AI in the sciences others work in more social and theoretical realms, exploring the ethical ramifications of what it means to be living and learning in this new era of AI. The breadth of work presented here today stands as a testament to the strong diverse community of learners within our Five College Community.

Thank you so much for your willingness to share your work and to connect with others interested in doing the same.

Thank you for joining us today!
-AILA Conference Committee

AMHERST COLLEGE
in the Liberal Arts

AI in the Liberal Arts Undergraduate Conference

4.19.2025

PROGRAM

9:00 a.m. CHECK IN/RECEPTION

9:30 a.m. KEYNOTE

Professor Lee Spector

10:00 a.m. PANEL 1: CULTURE & AI

- Tony Zhu (Tufts) - Can generative AI be aligned with a language's cultural memory—not just grammar?
- Jennifer Long (Wellesley College) - Fatal AI-attraction: Connecting Fictional Female AI in Video Games to Ourselves

11:00 a.m. IDEAS TO INNOVATION WORKSHOP

Ryan Ji: Capabilities and Limitations of AI Coding Assistants

11:45 a.m. PANEL 2: AI APPLICATIONS

- Dhyey Mavani (Amherst) - LogFlowAI: Transforming DevOps with Automation, Real-Time Analytics & AI
- Stephen Chen (Amherst) - Integrating Machine Learning Pipelines into Wearable Health Sensing
- Jianfei Lyu (Wellesley) - Real-time Credibility Assessment of Webpages through an LLM-powered Bot
- Reihaneh Iranmanesh (Amherst) - Shared Autonomy for Proximal Teaching

12:45 p.m. LUNCH + DAISY DEMO!

Boston Dynamics Robotics Demo - Alfeld Lab with Daisy the "Dog"

1:45 p.m. PANEL 3: AI IN CS

- Liam Davis (Amherst) - Lookahead Branching Heuristics for Efficient Neural Network Verification
- Osasikemwen Ogieva (Amherst) - High-Efficiency Online Data Generation to Improve Pretraining Scaling Laws of Deep Networks
- Soyon Choi (Amherst) - Towards a Foundational Understanding of the Effects of Neural Architecture on Adversarial Robustness

2:45 p.m. AI SAFTEY AT AMHERST WORKSHOP

Hank Hsu: Biosecurity Risks in the Age of AI
Reihaneh Iranmanesh: Overview of AI Safety and What Should We Do Now
Yilin Huang: AI and the Future of Work

Kenote Speaker

LEE SPECTOR

Professor Lee Spector is the Class of 1993 Professor of Computer Science at Amherst College. He received a B.A. in Philosophy from Oberlin College in 1984, and a Ph.D. from the Department of Computer Science at the University of Maryland in 1992. Dr. Spector teaches and conducts research in artificial intelligence, artificial life, and a variety of areas at the intersections of computer science with cognitive science, physics, evolutionary biology, and the arts. He has also directed the AI in the Liberal Arts initiative since 2020. Spector created the initiative to build an interdisciplinary community of people interested in thinking deeply about the bigger questions related to our lives and work in this new era of artificial intelligence.

About AI in the Liberal Arts (AILA): AILA aims to engage a broad, interdisciplinary community of participants in discussions and activities related to artificial intelligence, exploring and facilitating multi-way interactions between work in artificial intelligence and work across the liberal arts.

We are here to guide you toward scholarship, discussion, events, and other learning opportunities that are interdisciplinary and critically minded. We believe in the importance of asking big questions and thinking about these issues across disciplinary boundaries

Learn more about AILA:
web: https://www.liberal-arts.ai/
email: AILA@amherst.edu
@amherst_aila

AMHERST COLLEGE in the Liberal Arts

PANEL 1

AI and Culture with AILA Host: Yilin Huang

Tony Zhu (Tufts University)

Can generative AI be aligned with a language's cultural memory—not just grammar?

This project introduces an undergraduate-developed NLP toolkit for Classical Chinese, a culturally rich and under-resourced language often overlooked by mainstream AI models. By building a cleaned, annotated corpus and training a Word2Vec model to capture nuanced semantic relationships, the project addresses challenges posed by elliptical syntax and symbolic references. Its centerpiece is a fine-tuned GPT model designed to preserve thematic and philosophical coherence in generative tasks. As a working prototype, this toolkit highlights the potential of culturally sensitive NLP approaches and contributes to the development of more inclusive, context-aware AI systems.

Jennifer Long (Wellesley College)

Fatal AI-ttraction: Connecting Fictional Female AI in Video Games to Ourselves

How do we approach women who speak, think, and behave as human, but are explicitly labeled the opposite? Be it robots, cyborgs, or chimeras, popular media has long been populated with representations of non-human, human-made females. My research investigates the relationship between cultural imagination of yet-to-be technology and their real-life counterpart – more specifically, how portrayal of and treatment towards women in video games (as both robotic and human) may explain our actual perspective and usage of artificial intelligence. My work is divided into two parts: 1) a dating simulation made with RPGMaker MZ (within a larger installation space defined by graphic adverts), which narratively and formally subverts its genre expectations, and 2) an analysis of how cultural perspectives engender sexualized fictional female robots, gratuitously impractical within their narrative context.

PANEL 2

AI Applications with AILA Host: Jiaqi Huang

Dhyey Mavani (Amherst College)

LogFlowAI: Transforming DevOps with Automation, Real-Time Analytics & AI

LogFlowAI is an AI-powered DevOps tool that uses predictive analytics and real-time monitoring to proactively manage system performance. Our research lies at the intersection of agentic-AI, DevOps, and machine learning, specifically exploring how neural networks and statistical time series models can forecast system failures from log data. We've employed streaming data processing via Kafka, Zookeeper, and GPU-accelerated computations, and the project is ongoing, currently focused on extending our methods and refining the predictive accuracy of our models under guidance of Professor Riondato as well. This project won HackMIT 2023 (Y Combinator Challenge - 3rd Place), and was recognized by judges from NVIDIA, Meta, OpenAI, Walmart, Liquid AI, Tesla, and more. I also received a special Prize from Tesla: 48-hour demo drive for groundbreaking AI DevOps innovation.

Stephen Chen (Amherst College)

Integrating Machine Learning Pipelines into Wearable Health Sensing

This project explores the use of digital biomarkers from wearable sensors to classify handedness and quantify upper-limb movement. Using inertial measurement units worn on each wrist, we collected data from healthy subjects performing daily motor tasks, then extracted movement features grounded in clinical research. Our machine learning pipeline achieved strong performance in distinguishing dominant from nondominant hands (AUROC = 0.76, F1 = 0.75), with velocity- and speed-related features emerging as key predictors. This work contributes to the development of simple, interpretable biomarkers for understanding complex human motion in real-world settings.

PANEL 2 CONT.

AI Applications with AILA Host: Jiaqi Huang

Jianfei Lyu (Wellesley College)

Real-time Credibility Assessment of Webpages through an LLM-powered Bot (*Zoom)

Our research is in the field of credibility assessment and web literacy. In this research, we explore how LLM can be incorporated into a web interface as an educational tool for users to assess unknown websites' credibility and learn how to better detect misinformation. We employ experimental design by developing a prototype for Credbot, an AI-powered Chrome extension that automatically evaluates a website's credibility based on established criteria. Credbot is then evaluated on accuracy and latency across a set of websites with known credibility labels from Media Bias Fact Checker.

Reihaneh Iranmanesh (Amherst College)

Shared Autonomy for Proximal Teaching

My research is in the field of HRI (human-robot interaction).
My main question was: How can we decompose a complex motor control task (e.g., high-performance racing) into interpretable sub-skills? Our research asks whether weakly supervised unsupervised techniques (modified CompILE with language annotations) can yield human-interpretable skill segments that align with intuitive aspects of driving (such as braking, steering, and throttle control).

AMHERST COLLEGE in the Liberal Arts

PANEL 3

AI in Computer Science with AILA Host: Bilal Tariq

Liam Davis (Amherst College)

Lookahead Branching Heuristics for Efficient Neural Network Verification

This project focuses on enhancing the efficiency of neural network verification, a crucial step in ensuring the safety of AI systems, by optimizing the branching decisions within verification algorithms. We've developed a lookahead branching heuristic that simulates potential search paths to identify the most promising constraints, and initially implemented it in the Marabou verification framework, yielding significant performance improvements. We are currently extending this implementation to the Alpha Beta Crown verifier to gather more comprehensive results and are in the process of preparing the full paper.

Osasikemwen Ogieva (Amherst College)

High-Efficiency Online Data Generation to Improve Pretraining Scaling Laws of Deep Networks

Training deep neural networks at scale is increasingly constrained by fundamental problems of data scarcity. In this work, we propose a fast and cheap pipeline for online data generation for supervised learning models. Our approach introduces a warm-up phase where data is procedurally generated and transformed on-the-fly, eliminating disk-based data loading overhead. This synthetic data is adapted to model weaknesses through an adversarial framework. We explore the use of reinforcement learning to guide the sampler's policy, enabling dynamic adaptation of the data generation process to maximize training efficiency.

PANEL 3 CONT.

AI Applications with AILA Host: Bilal Tariq

Soyon Choi (Amherst College)

Towards a Foundational Understanding of the Effects of Neural Architecture on Adversarial Robustness

This project investigates the link between neural network architecture and adversarial robustness. We begin by formally analyzing linear multi-layer perceptrons (MLPs), deriving bounds on robustness as a function of model structure. These insights are extended to non-linear MLPs through controlled empirical studies, including the introduction of a tunable non-linearity parameter. Building on these results, we propose a novel defense method against deployment-time attacks that approximates attacker behavior using a linear proxy. Our approach improves both robustness and accuracy on synthetic benchmarks, even in the absence of an attacker.

AMHERST COLLEGE in the Liberal Arts

WORKSHOPS

IDEAS TO INNOVATION

The Capabilities and Limitations of AI Coding Assistants
Hosted by: Ryan Ji (Amherst College)

An interactive workshop exploring both the capabilities and the many limitations of AI coding assistants. What these tools are great for, where they fall short? Ryan will give us a brief overview of these tools followed by a live demo where he will take a program idea from the audience and attempt to build it on the spot using an AI assistant, all under a time constraint!

AI SAFETY AT AMHERST

Topics in AI Safety and Society
Hosted by: Hank Hsu, Reihaneh Iranmanesh, & Yilin Huang

- **Biosecurity Risks in the Age of AI**
- **Overview of AI Safety and What Should We Do Now**
- **AI and the Future of Work**

DEMO

MEET DAISY

A Boston Dynamics Robot Dog

EXPLORE AI AT AMHERST

The Amherst College AI Team:

Academic Technology Services (Project Manager)
AI in the Liberal Arts
Center for Teaching and Learning
Community Standards
Frost Library
IT Support Services
Provost's Office
Writing Center

The AI Team is a collaborative group of departments from across Amherst College offering multiple avenues to engage with Generative AI through education, exploration, and consultation.

We invite students, faculty, and staff to explore the options listed below.
Our matrix team will continue to develop these services as we navigate the landscape of Gen AI together with the Amherst community.

Questions about the AI Team?
Please contact Amherst IT at askIT@amherst.edu

AI at Amherst - Learning Lab

A series of talks and interactive sessions designed to help faculty and staff develop an understanding of AI and its implications for higher education, by providing a forum to connect with experts and fellow educators on a number of topics around the use of AI.

Explore AI at Amherst

Experiment with ChatGPT4 at various locations around campus, including the ChatGPT4 station in the Seeley Mudd IT lab.

Keep an eye out for **AmherstAI Café** events around campus!

AI in the Liberal Arts (AILA)

A broad, interdisciplinary community of participants in discussions and activities related to artificial intelligence, exploring and facilitating multi-way interactions between work in artificial intelligence and work across the liberal arts.

EXPLORE AI AT AMHERST CONT.

GenAI Task Force — Fall 2023 Work

Learn about the members of the College's Task Force and read the report.

AI at Amherst — Emerging Examples From Amherst College

Explore examples of Generative AI use at Amherst College

AI at Amherst — Curated Pedagogy Resources

Pedagogy strategies for a GenAI context

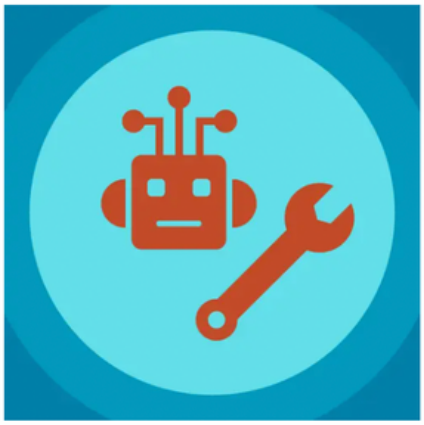

AI Tools

A resource for learning more about the digital accessibility, privacy, and security of some popular generative AI tools.

AI at Amherst - Consultations

Invite us to your department! Request consultations or demos with AILA, ATS, CTL, Frost Library, IT, or the Writing Center.

AILA UNDERGRADUATE CONFERENCE

Amherst College
April 19th, 2025

Notes:

AILA UNDERGRADUATE CONFERENCE

Amherst College
April 19th, 2025

Notes: