CRA Summit Report on AI Undergraduate and Graduate Education and Research

November 2025

CRA
Computing Research Association

Authors

Sri Yash Tadimalla, Computing Research Association

Stephanie T. Jones, Computing Research Association

Noah Cowit, Computing Research Association

Mary Lou Maher, Computing Research Association

Jeff Forbes, Computing Research Association

Tracy Camp, Computing Research Association

Acknowledgments

This material is based upon work supported by the U.S. National Science Foundation under Award No. 2533723. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

CRA gratefully acknowledges the support of the the generous sponsors of the CRA Summit on AI Graduate Education and Research: Google, AccessComputing, Center for Inclusive Computing, NCWIT, Sandia National Laboratories, Parkview Health, IEEE-CS, ACM, Dolby Labs, and Microsoft.

Suggested Citation

Tadimalla, S.Y., Jones, S. T., Cowit, N., Maher, M., Forbes, J., and Camp, T. (2025). CRA Summit Report on AI Undergraduate and Graduate Education and Research. Washington, D.C. Computing Research Association (CRA).
https://cra.org/wp-content/uploads/2025/11/CRA-Summit-Report-on-AI-Undergraduate-and-Graduate-Education-and-Research.pdf

Foreword

The 2025 CRA Summit was more than a series of co-located events — it was a moment for our community to come together, reflect, and reimagine the future of computing research and education in the age of artificial intelligence (AI). Across four days in San Francisco, participants from academia, industry, and government brought thoughtful perspectives, deep expertise, and a shared commitment to shaping AI's impact responsibly.

This report offers a synthesis of the key themes and insights that emerged from the Summit's two central convenings: *CRA Summit on AI Undergraduate Education* and *CRA Summit on AI Graduate Education and Research*. Together, they reveal both the opportunities and challenges facing our community as we work to educate the next generation, strengthen our research enterprise, and ensure computing continues to serve society's needs.

We extend our heartfelt thanks to all who made the CRA Summit possible — our speakers, moderators, and participants — whose engagement and candor made the discussions both rich and productive. We are especially grateful to the U.S. National Science Foundation (NSF) for its support of the CRA Summit on AI Undergraduate Education, and to our generous sponsors for supporting the CRA Summit on AI Graduate Education and Research: Google, AccessComputing, the Center for Inclusive Computing, NCWIT, Sandia National Laboratories, Parkview Health, IEEE Computer Society, ACM, Dolby Laboratories, and Microsoft.

The Computing Research Association (CRA) is deeply appreciative of the collaborative spirit that defined this event. The conversations that began at the CRA Summit continue to guide our collective efforts — to expand access to AI education, uphold research integrity, and build stronger connections across sectors. We look forward to carrying this momentum forward into the 2026 CRA Summit and beyond.

— Tracy Camp, PhD
Executive Director and CEO
Computing Research Association (CRA)

About the Computing Research Association (CRA)

The Computing Research Association (CRA) represents nearly 300 North American academic units, laboratories, centers, and companies engaged in computing research. Since its founding in 1972, CRA has brought together academia, industry, and government to strengthen the computing research community and its contributions to society.

CRA's mission is to catalyze computing research by leading the community, informing policymakers and the public, and promoting the development of an innovative and responsible computing research workforce. Through its programs and initiatives, CRA supports researchers across career stages and helps shape the future directions of the field.

Table of Contents

Executive Summary

The Computing Research Association (CRA) hosted the **2025 CRA Summit** on July 28-31, 2025. Over the course of the week, attendees from academia, industry, and government explored varied perspectives on how artificial intelligence (AI) is reshaping computing research, undergraduate and graduate education, and workforce development. This report summarizes the key discussions and takeaways from the two core academic convenings of the 2025 CRA Summit: **the 2025 CRA Summit on AI Undergraduate Education and CRA Summit on AI Graduate Education and Research**.

CRA Summit on AI Undergraduate Education

The CRA Summit on AI Undergraduate Education was focused on the strategies for incorporating AI into the undergraduate curriculum. Leaders debated whether to integrate AI concepts into existing core courses (like algorithms and systems), create new mandatory AI fundamentals courses for all CS majors, or develop more specialized AI-centric tracks and degrees. Participants shared models from their institutions, weighing the pros and cons of each approach based on faculty resources, student preparedness, and program accreditation. A significant portion of the Undergraduate AI education day was dedicated to the practical impact of AI tools on teaching and learning. There was a focus on the dual nature of tools like Microsoft Copilot and OpenAI ChatGPT, which are powerful productivity assistants but also present challenges to academic integrity and the development of core programming skills. Opportunities were discussed for designing assignments that encourage the ethical use of these tools for brainstorming and debugging, rather than for bypassing the learning process. A recurring tension was the need to balance teaching foundational computer science principles with the practical skills required to use and apply generative AI tools. Leaders debated how to ensure students graduate with a deep conceptual understanding that will outlast any specific technology, while also being equipped with the hands-on experience that employers demand. There was agreement on the critical importance of embedding AI ethics into the curriculum across academic and industry panelists. The conversation moved beyond advocating for a single, isolated ethics course to a model where ethics, fairness, accountability, and transparency are also woven into technical courses across the entire curriculum. This approach ensures students consider the societal implications of their work at every stage of development.

Key Takeaways

1. Educators and researchers need to **develop and disseminate best practices, pedagogical strategies, and ethical guidelines** for integrating generative AI into computing curricula to enhance learning while upholding academic integrity.

2. The community should leverage the forthcoming report on "Tracking the State of AI Undergraduate Education" to identify gaps, inconsistencies, and opportunities for standardization where appropriate.

3. Academic institutions may reflect with employers on **the distinct value propositions** of both specialized AI degrees and traditional CS degrees with AI concentrations, providing clear pathways for careers for students.

4. Universities and industry partners should **forge new, more agile models of collaboration** to ensure curricula remain relevant and that graduates possess the skills needed in a rapidly evolving, AI-driven industry.

5. The academic community is invited to **pursue NSF CISE funding opportunities** to innovate and expand AI education initiatives, build infrastructure, and support faculty development.

6. The community needs to **design and implement new methods for assessing student work** that account for the use of AI tools and **develop sustainable instructional staffing models** to support the growing demand for AI education.

7. A critical task is to **move beyond singular, standalone ethics courses to weave principles of responsible and ethical AI** throughout the AI undergraduate curriculum.

8. Institutions should **invest in building educational capacity** to make foundational AI knowledge accessible across all programs, ensuring a broad base of AI literacy among graduates.

CRA Summit on AI Graduate Education and Research

The CRA Summit on AI Graduate Education and Research addressed the dual role of AI in research: AI as a research field and AI as a tool for scientific discovery. A significant point of concern was the massive cost and scale of computational infrastructure required for state-of-the-art AI research. There was significant discussion around how research at academic institutions can continue to participate in high-quality AI research despite computational cost challenges, and the need to explore models for shared national research infrastructure, cloud computing partnerships, and consortia. The conversation on curricula focused on the need for advanced, post-foundational courses. Topics included the design of courses on large-scale model training, scalable AI systems, interpretable machine learning, and responsible AI deployment. An identified goal was preparing MS and PhD students not just for technical roles, but for leadership positions where they can help define the future of AI development and policy. A central topic was the evolving nature of the computer science master's degree in the age of AI. Leaders questioned what core competencies a new MS graduate in AI must possess

compared to traditional CS students. Discussions covered the need for deep specialization versus a broader, systems-level understanding and the importance of interdisciplinary collaboration as AI becomes integral to fields from biology to the humanities. The intense competition with industry for top-tier talent was also a critical point of discussion. Leaders shared strategies for recruiting and retaining faculty in high-demand AI fields, including competitive compensation and protected research time. They also addressed the "brain drain" of PhD students to industry by exploring the idea of more flexible faculty roles and stronger industry-academia partnerships, such as co-advised theses and collaborative, sabbatical-oriented, research experiences.

Key Takeaways

1. The research community should continue to **address key challenges in AI research**, such as improving reproducibility, ensuring rigorous evaluation of models, and developing more resource-efficient methods.
2. Researchers were challenged to **look beyond the current paradigm of Large Language Models (LLMs)** and explore novel architectures and approaches to drive the next wave of AI innovation.
3. Academic leaders should consider **rethinking computing graduate programs** to prepare students for a future where AI is a core component of research and industry, not just a specialization.
4. The computing community has a responsibility to **proactively analyze the impact of AI on the technology workforce** and adapt graduate education to equip students with the durable, high-level skills needed to thrive in their careers.
5. Universities should **strategically develop and expand professional master's programs** in AI to meet the industry demand for specialized, highly-skilled practitioners.
6. The community must **foster an environment that supports diverse research strategies**, encouraging both ambitious, high-risk "moonshots" and large-scale, collaborative projects to ensure a healthy and innovative research ecosystem.

Suggested citation:

Tadimalla, S.Y., Jones, S. T., Cowit, N., Maher, M., Forbes, J., and Camp, T. (2025). CRA Summit Report on AI Undergraduate and Graduate Education and Research. Washington, D.C. Computing Research Association (CRA).
https://cra.org/wp-content/uploads/2025/11/CRA-Summit-Report-on-AI-Undergraduate-and-Graduate-Education-and-Research.pdf

Introduction

The Computing Research Association (CRA) hosted the **2025 CRA Summit** on July 28-31, 2025. Over the course of the week, attendees from academia, industry, and government explored varied perspectives on how artificial intelligence (AI) is reshaping computing research, undergraduate and graduate education, and workforce development. This report summarizes the key discussions and takeaways from the two core academic convenings of the 2025 CRA Summit: the **CRA Summit on AI Undergraduate Education** and the **CRA Summit on AI Graduate Education and Research**. The CRA AI Education Fellows captured the insights from the Summit to provide key takeaways and summaries for the computing research community. You can view the full event agenda here, The slides presented in each session can be found here.

CRA Summit on AI Undergraduate Education

The speakers and panelists at the AI Undergraduate Education Summit examined the evolving AI curriculum, and embedding responsible AI practices and ethics across the student experience. Participants addressed challenges in integrating and scaffolding generative AI in teaching, developing accessible instructional models, and navigating the shifting relationship between academia and industry to prepare AI-ready graduates. The topics included:

- The current landscape of U.S. undergraduate education in technical AI.
- Discussions on CS vs. AI degrees.
- Integrating Generative AI into teaching computing.
- The evolving industry-university relationship in AI integration within CS undergraduate education.
- NSF CISE opportunities for AI education.
- Challenges in student evaluation and staffing to prepare AI-enabled students for professional success.
- Embedding responsible AI across the undergraduate experience.
- Building AI educational capacity across all computing undergraduate programs.

CRA Summit on AI Graduate Education and Research

Voices from academia, industry, and government came together at the AI Graduate Education and Research Summit to explore how the research ecosystem is adapting to rapid advances in AI. Sessions examined emerging trends and challenges in AI research, the shifting priorities of federal funders, and the need for reproducibility, rigorous evaluation, and clear communication of research impact to policymakers and the public. Speakers reflected on the global landscape of AI R&D, the pressures on graduate education to prepare students for a workforce where the

jobs of the next decade are not yet known, and strategies for aligning AI innovation with societal needs and ethical imperatives. The topics included:

- Reflections on NSF and the future of U.S. research.
- Trends and challenges in AI research.
- The impact of LLMs.
- Rethinking computing graduate education in light of the "AI Mandate."
- The impact of AI on the workforce.
- Strategic directions for academic leaders in building professional Master's programs in computing and AI.

The remainder of this report provides detailed summaries of each session held at the 2025 CRA Summit. For every session, we outline the central themes, key takeaways, and core insights shared by speakers and participants. All session leads and presenters were given the opportunity to review and refine their sections to ensure accuracy and clarity. We extend our appreciation to all contributors for their time, expertise, and support of both the 2025 CRA Summit and this report.

CRA Summit on AI Undergraduate Education - Plenary Sessions

The Landscape of the U.S. Undergraduate Education in Technical AI

Introduction by: **James Allan**, Chair, CRA Board of Directors; Associate Dean of Research and Engagement, Director of the Center for Intelligent Information Retrieval, Professor, University of Massachusetts Amherst

Speaker: **Carla Brodley**, Professor, Khoury College of Computer Sciences, Northeastern University, Founding Executive Director, Center for Inclusive Computing (CIC)

Key Takeaways

- Within the traditional computer science sequence, it can be difficult for students to take AI courses before the 4th semester in undergraduate programs due to challenging pre- and co-requisites in math.
- Across the nation, the number of higher education AI learning opportunities is increasing, including majors, concentrations/specializations, minors, and electives. These programs have large amounts of variation with regard to technical and applied AI curriculum.
- In-CS-department math curricula like “Math for AI” is one way that departments are beginning to address the layered math prerequisites required for AI courses.

Summary:

The CIC project “Tracking the State of AI Undergraduate Education” aims to create a system-wide understanding of how U.S. universities are preparing undergraduates in computing for AI. Questions the study seeks to answer include: ‘How well are universities responding to this need? Where (and what) are the innovations? What are the barriers for universities and innovations? In order to answer these questions, the CIC is creating a database that contains all computing bachelor’s programs in North America with regards to AI majors, minors, concentrations, required courses, and/or electives. This database will help in making sense of prerequisite courses for AI (Math and CS) and how many terms are required before a student can take an AI course, and in addition give a system-wide view on the state of U.S. AI education at the undergraduate level.

Currently, the study comprises 207 unique schools each of which graduates 150 or more CS major students every year (no certificates). Of these schools, 45 had an AI major, 38 had an AI concentration/specialization, 27 had an AI minor, 126 had AI as an elective, and 6 had no explicit AI offerings. Individual schools could have more than one of these types of programs.

To create this landscape, search queries were made with OpenAI and Perplexity using a structured prompt in order to find the URLs to curriculum pages from search results. Results were validated with a set of 100 universities’ programs (collected manually). In any case where the scraper found no AI course offered, the institution was manually checked. Exclusion cases for programs/schools were: colleges not in Arts & Sciences, Engineering, Science, or Computing, schools with more than three terms per year, and/or the curriculum page was broken. Current accuracy was 100 percent on finding majors, concentrations, and minors, whereas finding electives was around 75 percent accurate due to the different website and catalog structures that are less likely to be indexed.

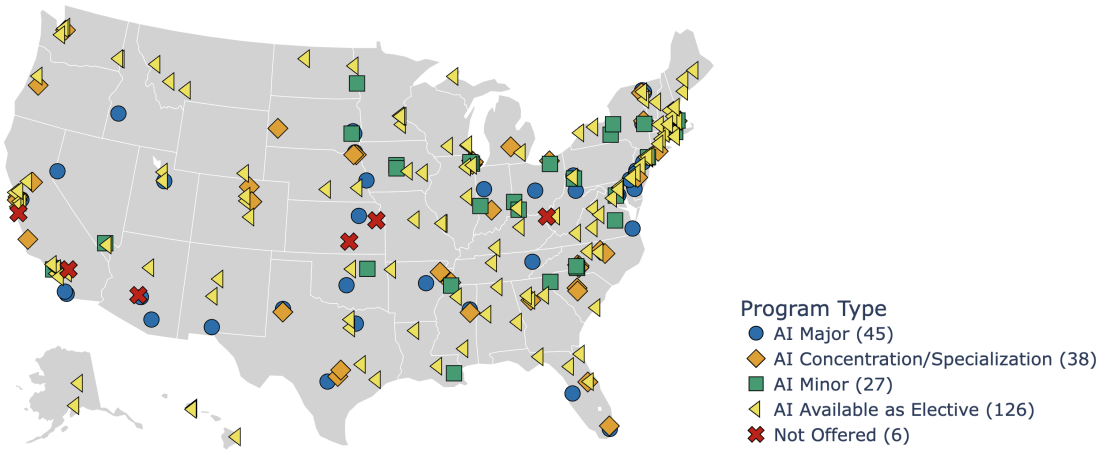

Figure 1: Landscape of U.S. Undergraduate Technical AI Programs - July 2025

In addition to understanding what programs were offered, the research group aimed to understand what courses were offered, what the prerequisites were, and the number of semesters before a student could take their first AI course in the best and worst cases. Of 236 programs in the 207 universities studies, only 113 were analyzed thus far; some websites were poorly formatted and thus the CIC made recommendations for formatting the department's information in order to be a part of this data set. Using an internally built tool, the CIC evaluated the prerequisites for AI classes. They found that there was variation in explicitly or implicitly requiring CS1, CS2, or Object-Oriented Design, Data Structures and Algorithms, or other advanced courses as prerequisites. Similarly, for math, there was variation in implicit or explicit prerequisites of discrete, calculus 1-3, linear algebra, probability, and statistics. Across the 113 schools, the mean amount of time to be able to take an AI class was 4.4 semesters, and to take an ML class was 4.22 semesters. The range spanned 3 to 7 semesters, which means that at some schools, students are not able to take their first AI course until the end of their senior year. Because of the delay in when students can take AI classes, many faculty try not to make assumptions about what AI knowledge students may bring, and so basic concepts like logistic regression are taught in multiple AI courses.

In order to continue building out this data set, the CIC team highlighted factors that made curriculum pages more accessible to both humans and machines. These included: Drop-down menus that share course information (including a description of what's covered), Math and Computer Science prerequisites / co-requisites explicitly listed, and hyperlinks to the pre-/or co-requisites and/or pop-ups with this information.

This report is a part of the ongoing work of the Center for Inclusive Computing (CIC, CIC.northeastern.edu) to understand trends in higher education regarding Artificial Intelligence (AI) education. The initiative to study changes in AI education builds on over 5 years of work aimed at understanding and creating pathways for systemic and sustainable changes. Their research suggests 1) changes to Intro Computer Science (CS) sequences curriculum and assessment approaches, 2) changes to student support approaches, and 3) changes to the CS major. The community-based efforts of the CIC are part of an effort to create resources so that universities can build consensus in approaches to changing computing education practices. The final report on the US Higher Education AI Landscape will be hosted on the CIC website (https://cic.northeastern.edu/).

CS Degrees vs. AI Degrees: How Should We Prepare Future Generations?

Moderator:

- **Eric Eaton,** Research Associate Professor in Computer and Information Science, primary member of the General Robotics, Automation, Sensing and Perception (GRASP) Lab, University of Pennsylvania

Panelists:

- **Antonio Delgado,** Vice President of Innovation and Tech Partnerships, Miami Dade College
- **Vasant Honavar,** Professor of Informatics and Intelligent Systems and Dorothy Foehr Huck and J. Lloyd Huck Chair in Biomedical Data Sciences and Artificial Intelligence, Pennsylvania State University
- **Lisa Meeden,** Professor, Computer Science Department, Swarthmore College
- **Peter Stone,** Truchard Foundation Chair in Computer Science, University Distinguished Teaching Professor, Department of Computer Science, The University of Texas at Austin

Key Takeaways:

- **The Core Conflict:** There are diverse views on the question of whether AI is a subfield built upon foundational CS principles or a distinct field that requires an independent or separate interdisciplinary degree.
- **Curriculum Design:** Team CS advocated for a stable CS core with flexible AI electives to withstand tech "hype cycles." Team AI called for a new paradigm—an adaptable AI degree that integrates human-centered considerations and ethics from the ground up.
- **Jobs and Accessibility:** The push for a new AI degree is driven by a changing job market and declining CS enrollment. Proponents argue that a dedicated degree creates an "AI-ready" workforce and can broaden participation, citing one program that attracted a broad computing population including 40% women.
- **Points of Agreement:** *AI literacy* for all undergraduate students. *Early exposure* to AI concepts for majors and non-majors alike. *Integrated ethics* woven throughout the curriculum, not siloed in a single course.

The 2025 Curriculum Debate: CS vs. AI Degrees

The discussion was framed not as a dry academic panel but as a "no holds barred academic cage match," complete with wrestling personas and a lively, confrontational tone designed to highlight the passionate divisions on a critical topic in higher education. Moderated by Eric Eaton of the University of Pennsylvania, the debate pitted "Team CS" against "Team AI." Team CS, the "reigning champion," featured Peter Stone from UT Austin and Lisa Meeden from Swarthmore College, arguing for the primacy of the traditional Computer Science degree. The challengers, "Team AI," comprised Vasant Honavar from Penn State University and Antonio Delgado from Miami-Dade College, who championed the creation of new, dedicated AI degrees. An initial poll of the audience revealed a divided room, leaning towards integrating AI

into already existing Computer Science degrees, setting the stage for a compelling exploration of the future of technology education with expert panelists from a variety of educational institutions.

Team CS: The Primacy of the Foundational Degree

The core argument from Team CS was built on the principle that a deep and rigorous grounding in the fundamentals of Computer Science is an indispensable prerequisite for any meaningful work in Artificial Intelligence. Peter Stone opened by asserting, "As you study computing, you need to understand computing. You need to understand the theory of what is computable and how computers are built and the architecture." From this perspective, creating an AI degree without this core suggests that foundational elements of computing are somehow non-essential to AI, a premise they firmly rejected.

Lisa Meeden expanded on this, arguing that the AI revolution itself is a direct result of core CS concepts. "We could not have gotten LLMs, with their billions of parameters, to run efficiently without systems researchers creating GPUs and figuring out how to effectively cache huge amounts of data." she stated, also crediting the invention of the transformer to CS principles. Initially framing AI as the "stepchild of CS," the team later nuanced this position. Peter Stone agreed with Vasant Honavar's point that AI could be seen as the "mother" of CS, as many foundational computing concepts originated from early AI research. However, he turned this point into an argument for his own side, suggesting that "computer science is what AI has left behind." This "residue" of solved problems—compilers, operating systems, networks—has become the stable, essential foundation upon which new AI innovations must be built.

From a curricular standpoint, Team CS advocated for a model that preserves a strong CS core while offering flexibility through electives and concentrations. Peter Stone described the UT Austin curriculum, which has six required foundational CS courses, leaving ample room for students to specialize in AI. They argued that this structure is more resilient and adaptable. Lisa Meeden expressed concern about the current AI "summer," noting that the field has gone through intense hype cycles before. "I think we might be reaching a plateau," she cautioned, "To jump from CS and jump on the AI bandwagon seems a little too quick for me." By keeping AI as a series of electives within a CS major, the curriculum can remain nimble, adapting to new AI trends as they emerge and fade without having to fundamentally restructure a degree program built on potentially transient topics. They argued that core CS requirements like programming, data structures, algorithms, and systems have stood the test of time, providing a durable education, while the specific "hot" topics in AI change rapidly.

Team AI: The Case for a New, Interdisciplinary Paradigm

Team AI countered with a powerful historical and forward-looking argument: AI is not a subset of CS, but a distinct and overarching field that merits its own degree. Vasant Honavar

passionately declared, "Before computer science was born, the possibility of mechanizing thought occupied the minds of mathematicians and philosophers". He argued that Alan Turing's fundamental questions were not about databases or software stacks, but about the nature of intelligence itself. "AI is not a subspecialty of computer science any more than physics is a subfield of mathematics. AI is the mother of computer science."

Antonio Delgado drove the argument into the present, focusing on the radical shifts in the job market and the nature of technology itself. "AI is not just the future, it is the present... AI is writing a code and debugging the code... it is even designing chips. AI is reshaping computer science." He presented stark data from his institution (Miami-Dade College), noting that for the first time, they have seen a 15% decrease in enrollment in their computer science program. He argued this is a direct result of the changing landscape, stating, "It is our job in higher education as faculty administrators to follow those in the job market." He framed the debate as a historical echo of the 1950s when engineering and math departments likely questioned the need for a "computer science thing."

A critical component of Team AI's position was the need for a new, inherently interdisciplinary curriculum that traditional CS programs are ill-equipped to provide. Vasant Honavar argued that as AI moves from automation to augmenting human capabilities, it must incorporate a deep understanding of how humans work, requiring education in the social sciences. He also asserted that the profound ethical and societal impact of AI cannot be relegated to a single course. "Some of the missteps we have seen from large companies... can be traced to the lack of training and the precision for this broad consideration as part of the CS education." An AI degree, they argued, would be designed from the ground up to integrate these human-centered considerations as experiential learning opportunities integrated within foundations in mathematics and computer science (logic, automata, programming, and algorithms), making them a core part of the technical education, not an afterthought. This new structure would also allow for greater accessibility. Antonio Delgado highlighted that by removing high-level math prerequisites for introductory courses, his college's AI program attracted over 2000 students, over 40% of whom were women, demonstrating that a dedicated AI track can broaden participation in the field.

Points of Convergence and a Shared Vision for the Future

Despite the theatrical "smackdown" format, the debate revealed significant common ground, outlining a shared vision for the immediate future of AI education. Both sides unequivocally agreed on the necessity of **broad AI literacy**. They endorsed the creation of courses, like the one Peter Stone taught called "Essentials of AI for Life and Society," designed for every undergraduate on campus with no programming prerequisites. The goal is to create a universally informed populace that can interact with AI technologies responsibly.

Furthermore, there was a consensus on the importance of **early exposure** to AI concepts for all students. Rather than waiting until junior or senior year, students should have the opportunity to engage with AI early on. Lisa Meeden described how Swarthmore College is creating a new intermediate course (at the same level as data structures and algorithms) to specifically prepare students for upper-level AI and Machine Learning courses, making it accessible to non-majors who have completed CS 1. This aligns with Team AI's push for freshman-level introductory AI courses. Team CS also argued for the development of more interdisciplinary courses that would be team taught with CS faculty and faculty from the Humanities and Social Sciences as a way of contributing to the broad AI literacy that is needed to reach all students.

Perhaps the most crucial point of agreement was the need for **deeply integrated ethics**. The moderator, Eric Eaton, noted that modern curriculum design is moving away from having a single, standalone ethics course. Instead, ethical and societal implications should be "integrated throughout the curriculum." This means discussing the bias of training data within a machine learning course, the labor implications of an automation project, and the public perception of a new AI application as integral parts of the technical instruction. Both teams saw this integrated approach as essential for training the next generation of responsible technologists.

Ultimately, the debate showed that while the question of "CS vs. AI degree" remains contentious, the underlying goals are shared: to create accessible, ethically grounded, and relevant technology education that prepares students not just for a job, but for shaping a society increasingly defined by artificial intelligence.The audience polls reflected this complexity. In the initial poll, about 50% of the respondents favored dedicated AI majors as a necessary addition, while approximately 50% were still unconvinced. When asked how they would design their department's curriculum tomorrow in the closing poll, about half of the audience believed that a broad CS major with AI tracks was the way to go, 17% believed AI should be covered primarily in its own major, and 33% believed that both an AI major and CS major with AI tracks should be offered and supported.

Opening Poll:

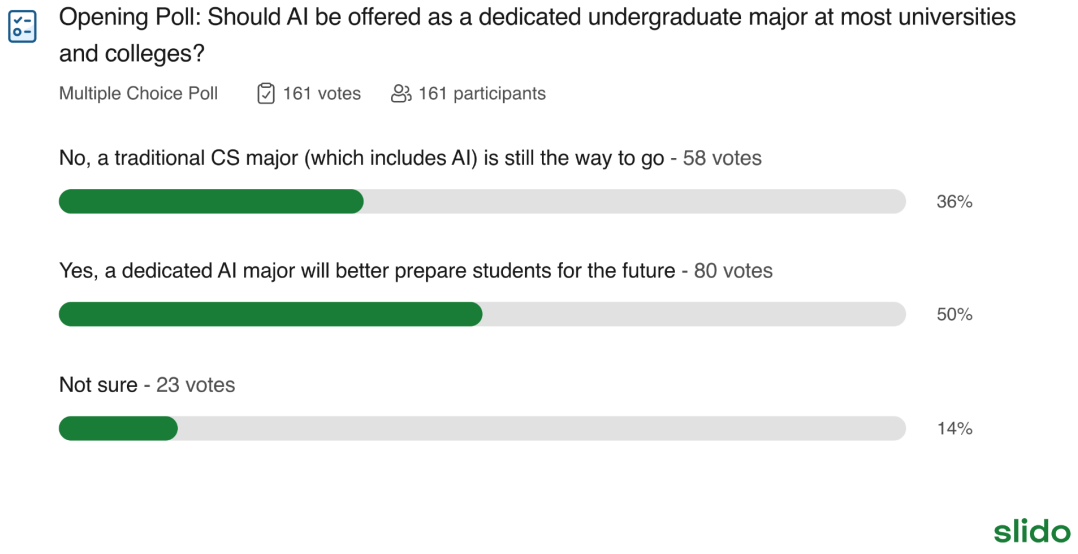

Closing Poll:

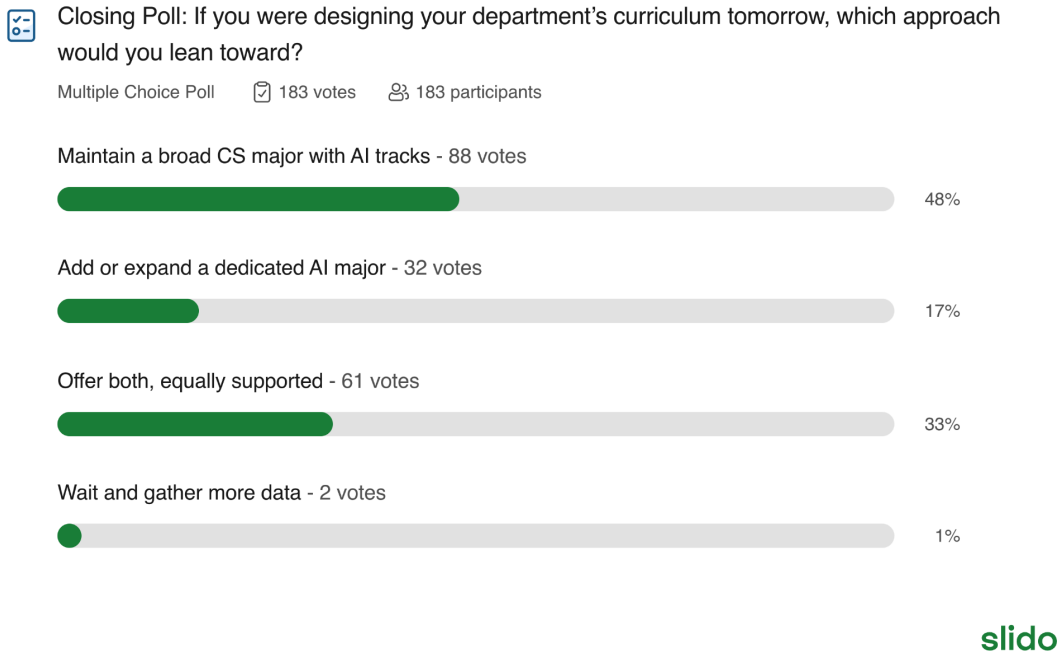

The audience questions during the "CS vs. AI Degree" debate centered on the practical implications of creating a new AI degree program and how it would function in the real world.

One major area of concern was the **job market and program targeting**. Audience members asked how an undergraduate AI degree aligns with industry needs, especially when many advanced AI roles currently require a master's or PhD. The panelists from Team AI argued that a huge demand exists for an "AI-ready" workforce with practical application skills, not just research-level expertise, and that a dedicated undergraduate degree is the best way to develop this mindset. Team CS countered that an online master's in AI is becoming increasingly affordable and accessible for specialization, and that employers still highly value the foundational knowledge provided by a traditional CS degree with an AI concentration.

Another key question involved the challenge of designing a **future-proof curriculum**. The moderator noted the extreme difficulty of creating a stable, 10-year curriculum when the field of AI evolves dramatically every few months (e.g., the sudden rise of LLMs). This point was used by Team CS to argue for the flexibility of a stable core with adaptable electives. Team AI responded that the goal isn't to teach a specific technology but to instill a mindset of **continuous adaptability** in both students and faculty, which is a core tenet of a modern AI degree.

Finally, questions touched on the **proliferation of AI across other disciplines**, like robotics and other sciences. The panelists agreed that AI is becoming a service discipline for the entire university, much like mathematics. This led to a consensus that it's crucial to offer accessible AI courses to students from all majors to foster interdisciplinary innovation and understanding.

Resources

- The moderator, Eric Eaton, referenced his role as chair of the AI subcommittee that developed **CS 2023**, the latest set of curriculum guidelines for computer science. He mentioned it as a framework that advocates for integrating ethics throughout the curriculum.
- Essentials of AI for Life and Society is a course offered at The University of Texas at Austin. It is open to all majors with no previous experience in AI.

Integrating Generative AI into Teaching Computing

Speaker:

- **Leo Porter,** Professor of Computer Science, University of California San Diego

Key Takeaways

- **Industry Expectation:** The software industry is rapidly integrating Generative AI (GenAI) and expects new computer science graduates to be proficient in using these tools. This creates an urgent need for higher education to adapt its curricula to avoid graduates being at a competitive disadvantage against newly retrained engineers.
- **Shift in Programming Skills:** There's a strong consensus among educators present at the summit (75%) that the essential skills for programming are changing. Initial ideas around skill re-focusing include: shifting away from **writing code and managing low-level details** towards higher-order skills like code reading, problem decomposition, testing, debugging, and ethical reasoning in an AI-assisted environment.
- **Three Pillars of Integration:** Effective integration of GenAI in computing education rests on three core strategies: *Assessment Redesign, AI Tutors,* and *Direct Skill Training.*
- **Community-Driven Adaptation:** The challenge of curriculum overhaul is being met with collaborative efforts. Initiatives like the "GenAI in CS Education" consortium and new educational resources (e.g., Porter and Zingaro's textbook) are being developed to support faculty in redesigning their courses.

Summary:

The session addressed the need for computer science programs to adapt to the rise of Generative AI. Porter shared an anecdote from a conversation with a senior executive in software, highlighting that industry is already operating under the assumption that graduates will be skilled in using GenAI. This sets a clear expectation for academia, which traditionally moves much more slowly. The competitive landscape is shifting, with new graduates now competing against existing engineers who are being retrained globally in AI-assisted development.

Varied evidence on GenAI's impact was discussed: where some studies showed significant productivity gains (e.g., developers being 56% faster with GitHub Copilot), others suggested that experienced engineers can be less efficient with these tools. Despite these inconsistent results, he argued that the trend of adoption is undeniable and academia does not have time to "wait for the dust to settle." He also noted that GenAI might allow students to tackle larger, more meaningful software projects earlier in their education. He warned that since many students will use these tools anyway, it is the university's responsibility to train them properly.

A live poll during the session revealed that a significant portion of attendees' departments are already taking action: 41 percent have already implemented curricular changes, and 32 percent are actively planning integration. This aligns with session data showing 75 percent of educators present acknowledge programming skills are changing due to GenAI.

Porter outlined three primary methods for integration:

1. **Assessment Redesign:** Acknowledging that GenAI makes it difficult to measure learning in unsupervised take-home assignments, he stressed that the bulk of a student's grade may come from proctored exams.

2. **AI Tutors:** He discussed the rise of AI tutoring tools that use Socratic models to guide students without giving them direct answers. Students value these tools for their constant availability and non-judgmental interaction.

3. **Training Students to Use GenAI:** This goes beyond prompt engineering. The focus is on teaching students to use LLMs for code reading, testing, debugging, and problem decomposition. A poll of the audience on what skills are becoming *less* important due to GenAI identified syntax memorization and lower levels of abstraction in coding as top responses.

He presented UCSD's approach to CS Programming learning experience, which includes a revamped introductory programming course based on his textbook, *Learn AI-Assisted Python Programming*. The course emphasizes critical usage of AI and is a model for replication throughout the CS curriculum. The session concluded by looking toward a vision for the future of computing education in the age of AI. One path forward involves a fundamental and ongoing shift in the CS curriculum, moving from syntax to critical application and problem-solving.

Future Research: There is a clear need for more research into how GenAI tools specifically impact student learning processes and outcomes, which are currently less understood than the tools' capacity for tutoring.

Faculty Development: Creating scalable and effective models for upskilling faculty is critical.

Collaborative Curriculum: The future of curriculum development will likely be more collaborative, with consortia and partnerships between industry and academia.

Challenging Conversations: Institutions must address governance questions, such as tasking specific faculty to lead AI integration and forming dedicated committees for curriculum review.

Question and Answer:

- How do you differentiate the integration of GenAI for CS majors versus non-majors who might only take one introductory course?
    - A: Porter framed GenAI as a "major win" for non-majors, as a single CS course can now equip them with a powerful, practical skill they can use in their own

fields. For CS majors, the integration is more foundational, preparing them for a career where these tools are standard.

- How can departments effectively upskill faculty who may not be experts in GenAI?
    - The new "GenAI in CS Education" consortium is one response, designed to provide resources and community support. He recommended departments hold dedicated meetings and allocate resources for course redesign and training.
- What about the environmental impact of training and running these large AI models?
    - A: Porter acknowledged the importance of the question but stated that it was outside the scope of his direct work on curriculum integration.

Resources:

- **Textbook:** Porter, Leo, and Daniel Zingaro. *Learn AI-Assisted Python Programming: With GitHub Copilot and ChatGPT*. Second Edition, Manning Publications, 2024. *An instructor copy can be requested via email: leoporter@ucsd.edu*
- **Consortium Website:** www.teachcswithai.org (GenAI in CS Education Consortium)

AI Integration in CS Undergraduate Education: How Does the Industry-University Relationship Change?

Organizer and Moderator:

- **Martin Margala**, Professor of Computer Science, Director of School of Computing and Informatics, Endowed Chair of Computer Science Eminent Scholar, Fulbright Distinguished Chair, University of Louisiana – Lafayette

Speakers:

- **Sharath Kanukuntla**, Executive Director of R&D Engineering, Synopsys Inc
- **Emily Kemp**, Program Manager, Google
- **William LaBar**, Vice President, Consulting Services, CGI
- **Peter Pappas**, Operational Excellence Leader, DXC Technology

Key Takeaways

- Partnerships between Industry and computing departments can shape curricula to be relevant to current advances in technology.

- While industry panelists were excited about the opportunity for AI to increase efficiency, they are also looking for new hires to be skilled in critical thinking, effective communication, and problem-solving.
- Industry partnerships with regional universities can create job pathways and learning opportunities for students. These have been cultivated through school-year mentorship programs and partnerships between universities and regional technology (industry) councils and chambers.

Summary

This session was held as a panel discussion that was framed and facilitated by Martin Margala. In his opening, he emphasized the importance of industry voice in understanding the path forward for Artificial Intelligence (AI) education. To provide a picture of the graduating landscape: currently over 400 higher education institutions offer undergraduate Computer Science (CS) and Information Technology (IT) degrees and courses with bachelor degrees growing significantly. However, while students from these majors are able to apply for high-paying positions, recent graduates are experiencing a higher-than-average (6.1%) unemployment rate. Potential causes shared were gaps between academic training and industry demands, or companies prioritizing practical experience over theoretical knowledge. As such, he invited industry professionals to speak on this panel and share insights and approaches around hiring and what they expect from graduates. The panel represented non-overlapping areas, different demands, sectors, and needs.

Existing Engagement of Undergraduate CS Students with AI?

Panelists shared their varying perspectives on how they have seen undergraduate students engaging with AI, business principles, and problem-solving. The panelists all saw AI tools as an opportunity to increase productivity. However, there were several areas where panelists saw a need for further undergraduate education: systems-level understandings of computing work, the importance of knowing problem-solving strategies, and how to use these to interpret business needs.

Systems Level Thinking

Panelists emphasized that students need to be able to have more systems-level experience in terms of understanding what the role of their computing work might be in relation to the larger effort of an organization. For example, a new employee must understand how their updates to the code base on the back-end could cause issues on the front-end. Kemp estimated that less than 25% of current undergraduate students that she encounters have the system-level skills needed to be fully successful in roles at Google. In her role as Lead for the Google AI Education team, she is working to create pathways for AI classes & programs as a part of the

orientation to ensure new hires are trained and able to be effective users of the technology. She expects to see the top candidates for roles being students using AI on a daily basis.

Another approach to having undergraduate students gain this level of systems understanding is through professional development-based partnerships with industry. LaBar discussed CGI's efforts to engage with local universities so that students can learn about industry needs from "day one" on campus. Currently, CGI pairs students with industry professionals who mentor the students in their first semester. Students across informatics and math learn about identifying team roles, interfacing with people, building full-stack applications, customer experience, agile development, and other systems-level understanding that goes beyond the product. In the first 4 years of this program they hired 50 college grads annually from this initiative. Pappas discussed working with regional colleges to host onsite workshops around advances in current technology. He has also spent time in classrooms teaching how this technology is being used, and learning what schools are teaching to better understand what new hires might know. DXC Technology has been working to bring knowledge of industry needs to university curriculum development through partnering on micro credentials and hosting strategic AI curriculum calls.

Problem Solving

All panelists emphasized the importance of problem-solving for computer scientists. Kanukuntla highlighted AI's role as a kind of tool that can be used as a part of problem-solving. His company looks to make sure undergraduates are still able to engage in problem-solving processes independently when hiring. It is important for students to understand how to use AI outputs knowing that the datasets used for training might not be right for the companies' needs. Likewise, LaBar detailed how CGI trains students to be able to define problems themselves, understand patterns around architectural and design parameters, and when / how to use AI to meet business needs.

Shifts in Client Demands Driven by the Rise of AI

How do these changes impact your expectations of Computer Science (CS) graduates entering the job market, particularly regarding skills and readiness for future challenges?

Panelists shared excitement about the idea of cross-industry applications for AI and noted how that would necessitate domain expertise. Each had unique challenges and opportunities to share regarding how their clients and companies are interacting with AI. LaBar detailed how clients are requesting particular domain knowledge and sees it as a need for company growth in the market. He mentioned how GenAI provides an opportunity for more throughput and experimentation, but understanding the fundamental problem that the business is trying to solve is a critical differentiator in talent. Further, LaBar anticipates that new employees will be

expected to better evaluate the quality of the data in addition to the data itself. This means they are looking across the university to hire people with deep domain expertise (e.g., in healthcare). From Kemp's perspective, AI should not create a scarcity of software jobs but rather an abundance of AI capabilities. At Google they are exploring how with advances in AI, generating software applications might be easier, and as such how do they understand what makes a great product. This results in more time investment in cross industry applications and creating personalized multimodal products. Kemp expects that everyone will be AI users, but many will also want to develop AI as well.

Pappas highlighted the importance of technical skills in addition to being able to determine where and how to embed AI. As a company (DXC Technology) that provides contracting, their employees need to be able to understand business needs, effectively communicate with the client to understand problems, and be able to interpret how AI fits into that. Kanukuntla emphasized coding skills as still being required but that they're trying to reduce the time from execution to deployment of software significantly. He wants new graduates to be thinking about how AI agents can or do interact with one another as well as how to get the software stack and APIs to interact while reducing power requirements. While scaling up, optimization is also very essential.

What Has Worked and What Has Not?

Panelists shared briefly on the importance of needing CS students to have experiential learning whether through internships or experimenting with projects and developing a portfolio. They also mentioned that they want their employees to be adaptable to changes in their company's ecosystems through methods such as upskilling. Pappas and LaBar highlighted how industry partnerships are able to support students with learning upskilling approaches and enable sharing advice to create and shift coursework to support learning needed for industry success.

Question & Answer

- How do we build these relationships with industry partners? Are there regional models?
    - A: Models look like a university inviting industry leaders bi-weekly to share with faculty current events in technology and what changes are coming on the horizon. This shapes how the curriculum could be adapted to best prepare graduates. Examples of locales with these models include Lafayette, Louisiana, and Grand Rapids, Michigan.
- How is the use of AI shaping interviews, and how are students being evaluated? Particularly, thinking about LeetCode, which tests how good students are without an LLM.

- A: For in-person interviews it's not acceptable for students to pull ChatGPT out during a conversation and ask it a question (real example). It is also not advised for them to prompt AI during online interviews and read the response verbatim. The interviewer is still looking for who the potential employee is and what they can do.
- A: We may begin to see interviews asking people to walk through how they might use AI to solve a problem. The panelist is still hoping to see good programming because students still need to be able to know if AI is providing the best solution.

- In the next 5 years, how many people that you're hiring will need to be users versus developers of AI technologies?
  - A: Everyone should be able to use AI. At Google, they have historically hired 10% in AI development but see it increasing over the next few years.
  - A: We want to see how all people are leveraging AI tools to increase productivity.
- Problem-solving and critical thinking are core, but what would be your next one or two characteristics of potential employees?
  - A: An additional important characteristic is customer service and client interaction.

NSF CISE Opportunities for AI Education

Speaker:

- **Ellen Zegura,** Acting Assistant Director for CISE (Computer and Information Science and Engineering), U.S. National Science Foundation

Key Takeaways

- **Shaping the Future:** Educators have the agency and responsibility to actively guide AI's impact on society—it is not predetermined.
- **Full-Spectrum Strategy:** The NSF is funding the entire pipeline, from creating AI-literate citizens to advancing world-class research talent.
- **Distributing High-Quality AI Resources:** The National AI Research Resource (NAIRR) provides students and educators with crucial access to the high-end data, tools, and computing power needed for modern AI.
- **Durable Skills Over Hype:** The focus is on teaching lasting, foundational skills, not chasing temporary tech trends. The NSF emphasizes the importance of sustaining

thoughtful discovery research over the long term, ensuring that durable skills and foundational knowledge remain at the center of AI education.

- **Computing Education is Core to the NSF:** The establishment of a new CER program in the core signals that researching how to teach AI is a central priority.

Summary:

Zegura offered a high-level perspective on the role of the research and education community in the burgeoning field of Artificial Intelligence. She began not with technical specifics, but with a profound call to action, framing the current moment as a crossroads. "There are multiple possible futures for us, for society with AI," she asserted. "There is no single and inevitable future." This established the central theme of her talk: human agency. She stressed that the collective group of educators, researchers, and administrators in the room possesses the unique expertise at the intersection of computing and education to actively shape which of these futures comes to pass. This is not just an opportunity, but a "responsibility to the future" itself.

Zegura acknowledged the debate over whether the current AI shift is truly different from past technological upheavals like the advent of the calculator. While admitting she would "sleep better at night" if it were the same, her conviction is that this time *is* different. The NSF, as a national agency, is committed to addressing this unique moment by supporting the development of a full spectrum of AI competencies. She outlined this spectrum as ranging from AI-literate citizens on one end, to proficient AI users, to skilled AI developers, and finally to AI researchers who are advancing the fundamental progress of the field. This comprehensive vision, she noted, is a core tenet of the NSF's strategy.

A Legacy of Support and a Call to Action

To ground her vision in history, Zegura took the audience on a brief tour of the NSF's long-standing commitment to computing education. She used personal, concrete examples, primarily from her 30-year tenure at Georgia Tech, emphasizing that these were illustrative, not exhaustive. A key example was a 1994 NSF award for "Multimedia Support for Introductory and Advanced Computer Science Education." This project aims to use STEM-inspired multimedia content to engage students, a precursor to many modern pedagogical approaches.

She broadened her scope to showcase the national impact of NSF funding. A slide depicted the widespread distribution of CUE (Computing in Undergraduate Education) awards across the United States, with a humorous anecdote highlighting Wyoming as the only state without one, urging faculty there to apply. Another example was the concept of computational thinking, which—propelled by an influential essay—fundamentally shifted public discourse and K-12 educational policy around computing. The NSF invested in programs such as CS4All to

support the educational community in implementing the concepts behind computational thinking. These examples served to establish the NSF as a long-term partner that has consistently invested in the foundations of computing education.

Recent Initiatives: Seeding the AI Education Ecosystem

Transitioning from the past to the present, Zegura detailed several recent, high-impact NSF initiatives specifically designed to foster AI education.

1. **AI Research Institutes:** The NSF has established a set of AI Research Institutes, each with a mandated education and outreach component. She presented a slide illustrating how these institutes are contributing across the entire educational pipeline, from K-12 outreach and professional development for teachers to post-secondary education and informal learning. She particularly highlighted the critical need for "University professional development to be able to teach AI," echoing a theme from earlier in the day.

2. **EducateAI Dear Colleague Letter (DCL):** Zegura demystified a key piece of NSF communication: the Dear Colleague Letter. She explained that a DCL is not a new funding opportunity, but a nudge from the agency to the community, signaling an area where they are eager to receive more proposals submitted to existing programs. A DCL is a mechanism to reduce the lengthy approval process for a new solicitation. A year and a half prior, the NSF issued the EducateAI DCL to encourage proposals aimed at expanding K-12 and undergraduate AI education. Antonio Delgado—a speaker in the CS versus AI major debate—was referenced as a recipient of an award from this initiative.

3. **The National AI Research Resource (NAIRR):** A cornerstone of the NSF's recent efforts is the NAIRR pilot, a program designed to democratize access to the computational resources, datasets, and tools necessary for AI research and education. Zegura emphasized that from its conception, the NAIRR was designed with the classroom in mind. The program is actively supporting educators and students by providing access to powerful infrastructure and coordinating a community of educators to share best practices for enabling AI in teaching.

Challenges Ahead and a New Vision for Support

Looking forward, Zegura acknowledged the immense pressures on the higher education landscape to evolve. She recounted a stark anecdote about her own institution: "When is the last time Georgia Tech updated the curriculum? If I am being super honest, I think it was roughly 2005." This two-decade gap stands in sharp contrast to the industry perspective heard earlier, which suggested curricula should be updated almost monthly. This highlighted the "fools errand" of trying to chase the latest technology. Instead, the focus must be on identifying

and teaching **durable skills** that prepare students for a world of constant change and lifelong learning.

In response to these challenges, the NSF is actively restructuring its opportunities. The most significant announcement was the creation of a new, core **"Computing Education Research" program**. This move is meant to formally recognize the mainstream value of this work, a contrast to 2015 when the first CAREER Awards were given to "straight up, unapologetic" CS education researchers. This new program signals that research into *how* to effectively teach AI and computing is no longer a niche interest but a central priority for the agency. To further fuel this educational mission, the NAIRR is making a rich repository of datasets available to students and educators. Zegura's genuine excitement was palpable as she listed some of them, from cybersecurity logs to chip design data.

She concluded with a powerful metaphor: the **campfire**. She described the NSF's role as providing the kindling and the initial match, but it is the community that must work together to tend the fire. "One of the most important things you can do to a fire is blow on it," she explained. "Oxygen, attention, celebration, elevating, that is key to taking these ideas that are starting to get traction and turn them into something big." This was a direct appeal for collaboration, for sharing successes, and for working together to build a robust and vibrant AI education ecosystem. Her final message was clear: "No one can meet this moment alone."

Question & Answer

- **Infrastructure and curriculum support**: How will NSF help colleges and universities with the "heavy lift" of evolving their curricula?
    - This is a significant undertaking and stated that the NSF is actively thinking about ways to support departments in this work. Zegura invited the community to provide feedback on what kind of support would be most helpful, emphasizing a desire for partnership.
- **Mid-career faculty development:** Can NSF funding be used to help existing professors gain the skills needed to teach new AI courses?
    - This is a critical piece of the puzzle, stating that any plan for curricular change is a "nonstarter" without a plan to staff the courses. Zegura confirmed that supporting faculty retraining is a key part of the NSF's strategic thinking.
- **Industry's fast-paced demands vs. Slow & thoughtful academic research**: How might the focus on rapid AI turnaround devalue fundamental, long-term discovery research and our ability to train at the undergraduate level?

    - Zegura shared that discovery research is the "bread and butter" of the NSF. She stressed the agency's role to step back from the immediate industry pressures and focus on identifying and funding the pursuit of durable skills that will prepare students for the long term, rather than just the latest technological trend.

- **Support for interdisciplinary work:** Opportunities specifically integrating AI with the humanities and K-12 education (AI+X)?
    - Workforce development is another core part of the NSF's mission and programs exist to support applied AI research in these and other areas.

- **PI eligibility for RET Sites program:** Is it still a necessity for RET Sites PIs to be tenure-track faculty members, particularly in a quickly advancing discipline like AI in which earlier-career academics can offer novel ideas and insights?
    - While there were no current discussions in the NSF about removing the tenure-track requirement for a RET Sites PI, in her opinion sharing with NSF this ask might lead to more opportunities for non tenure-track academics to receive funding for their AI Education research projects.

Resource:
https://www.NSF.gov/funding/opportunities/future-core-computer-information-science-engineering-future-computing

Student Evaluation and Instructional Staffing: Current Challenges and New Models to Prepare AI-Enabled Students for Professional Success

Moderator:

- **Chad Jenkins,** Professor of Robotics, University of Michigan

Panelists:

- **Dan Garcia,** Teaching Professor, Electrical Engineering and Computer Sciences Department, University of California, Berkeley
- **Jasmine Jones,** Assistant Professor, Department of Computer Science, Berea College
- **Joanna Millunchick,** Dean, Luddy School of Informatics, Computing, and Engineering, Indiana University
- **Darnell Moore,** Principal Technical Program Manager, Amazon

Key Takeaways

- The industry bar is rising: Contrary to fears of de-skilling, the standard for new hires in tech is increasing. Industry seeks graduates with strong foundational skills complemented by entrepreneurial experience and an ability to use AI tools effectively.
- Assessment is at a crossroads: There is a clear crisis of confidence in project-based assessments due to AI. This has led to a reactionary trend toward more heavily-weighted, proctored exams, which may not assess the skills industry values.
- Standards must be maintained: Panelists strongly agreed that academic standards should not be "watered down." They advocated for ending the practice of grading on a curve to ensure degrees represent a consistent and high level of competency.
- Scaffolding is key: A central debate emerged between banning AI in early courses to build fundamentals and appreciation versus introducing AI early in a highly scaffolded way (the "Blue Apron" analogy) to engage students with more powerful concepts sooner.
- Faculty need support: There is a significant split among faculty regarding AI, from anxious and hesitant to optimistic. Effective faculty development is crucial to moving forward cohesively.
- The human element is irreplaceable: In an age of AI tutors, the value of human-to-human interaction, classroom community, and learning from the unique perspective of a human instructor becomes even more critical.

Summary:

The session's moderator, Jenkins, initiated the discussion by showing the audience an AI-generated image prompted with "salmon swimming downstream." The resulting image was a picture of a filleted salmon in a river, he noted that while the image was technically correct, it lacked the nuance and accuracy of reality—a metaphor for the output of current GenAI tools. He grounded the panel's urgency by sharing a sobering anonymous quote from a respected colleague, which detailed a clear drop in student exam scores, a collapse in office hour attendance, and a growing inability for students to understand their own code since the proliferation of AI tools. This framed the core problem: a potential disconnect between the use of AI for task completion and the development of true understanding.

A Wide Spectrum of Philosophies and Institutional Responses

Representing a large public institution, Millunchick of Indiana University began by referencing a popular webcomic depicting a programmer being replaced by AI, and later serving robot

overlords. She rejected this dystopian view, arguing for a more progressive and optimistic human-AI partnership. Drawing on her institution's unique history as an "informatics-first" school, she advocated for embracing AI as a powerful educational tool. She envisioned an AI tutor akin to the personalized "primer" in Neal Stephenson's novel *The Diamond Age*: a tool that teaches and adapts but works in concert with essential human mentorship. Millunchick also launched a pointed critique of traditional STEM pathways, calling for an end to the "tyranny of calculus" in high school in favor of linear algebra and statistics, which she argued are more foundational for modern computing. She firmly pushed back against what she saw as reactionary faculty responses, such as reverting to in-person paper-based exams, viewing them as a failure to adapt.

Providing a vital industry perspective, Moore from Amazon outlined his company's strategy for academic collaboration, which rests on three pillars: direct research partnerships, faculty/student development programs like the Amazon Scholars program, and K-12 outreach to prime the future workforce. His most significant contribution was the announcement of a forthcoming SDK for the "Amazon Astro" robot, a concrete step by industry to put advanced tools into the hands of academic researchers and students. This, he explained, was driven by the understanding that "all of the best ideas are not always within Amazon." Later in the Q&A, Moore delivered a critical reality check, stating that the hiring bar at Amazon is unequivocally rising. The company increasingly seeks graduates who not only have high GPAs, but also demonstrate entrepreneurial drive and experiential learning through projects outside the classroom.

Representing a vastly different academic environment, Jones of Berea College shared insights from a small, no-tuition work college serving a primarily first-generation student body. She described a "grassroots" response to AI, where the faculty was split into thirds: the anxious who hoped AI would disappear, the resistant who saw it as another burden, and the optimistic or techno-deterministic group. This division, she noted, creates significant challenges for establishing clear, consistent policies, leading to student confusion and placing undue pressure on pre-tenure faculty. Jones advocated for building community trust through participatory methods, such as having a class co-design its own "code of conduct" for AI use. She provided one of the panel's most stark data points: an anecdote about her undergraduate TAs being "shocked and disappointed" by the deteriorating coding abilities of their younger peers, offering direct evidence of AI's potential to erode foundational skills.

Garcia from UC Berkeley presented a deeply cautious and principled stance. He argued for banning GenAI in introductory courses, using the analogy that "you don't give a first grader a calculator" before they develop fundamental number sense. His proposed model was strictly scaffolded: no AI in the first year, optional use in the junior year, and required use by senior year. To capture the philosophical stakes, he shared a quote from his colleague Mitch Resnick: "You can teach a kid to play the piano or play the stereo," arguing that over-reliance on AI robs

students of the joy and deep learning that comes from creating something themselves. Garcia was also adamant about academic standards, condemning the practice of curving grades as a mechanism that passes students who haven't met the competency bar. His personal account of his once-packed office hours becoming completely empty in the span of a single semester painted a vivid picture of the profound shift in student behavior away from human-to-human interaction.

Audience and Panelist Discussion

The interplay between the panelists was dynamic. During the Q&A, Millunchick directly countered Garcia's cautious stance with her "Blue Apron" analogy. She argued that just as a meal kit provides a highly scaffolded entry into cooking for a novice, AI tools could similarly enable introductory students to tackle more interesting and powerful problems than sorting a list, thereby sparking and sustaining their interest in the field. This exchange crystallized a central pedagogical debate of the session: whether to shield students from AI until fundamentals are mastered or to leverage AI from the beginning to accelerate engagement.

- **On Assessment at Scale:** With projects being easily completed by AI, how can institutions authentically evaluate what students can *do*? The panel acknowledged there are no easy answers. While scalable options like computer-proctored testing are becoming more common, they may not test the right skills. Less scalable but more authentic methods like oral exams, multi-day projects, and peer evaluation were discussed as necessary explorations.
- **On the Risk of Failure:** When asked about the damage of "getting it wrong," panelists noted the consequences are immediate. If students graduate without being able to code independently, they will fail in the workforce, and the software they build will be subpar. The ultimate test for faculty may be whether they would trust the products their students create.
- **On Classroom Productivity:** There was some optimism that AI could be used to make the classroom *more* productive. By helping students overcome "unproductive struggle" (e.g., frustrating bugs or writer's block), instructors could potentially start at higher levels of abstraction and cover more complex and engaging material.

Calls for Research

The panel's discussion highlighted several critical areas for future research:

1. **Developing Scalable, Authentic Assessments:** There is an urgent need to design and validate new assessment methods that can measure high-order cognitive skills in a way that is resistant to AI shortcuts.

2. **Longitudinal Studies of AI-Assisted Learning:** Research is needed to understand the long-term cognitive impact on students who learn to program with heavy AI assistance from the beginning versus those who follow a more traditional path first.

3. **Evaluating Pedagogical Frameworks:** The "scaffolded" models proposed by panelists need to be formally developed, implemented, and studied to measure their effectiveness on student learning outcomes and self-efficacy.

This session made it clear that the introduction of AI in education is not just about adopting new tools, but about fundamentally re-evaluating *what* we teach and *how* we measure learning. The core challenge for the computing community is to navigate the transition from viewing AI as a threat to academic integrity to leveraging it as a catalyst for a more rigorous, relevant, and engaging educational experience. The next 1-2 years will be critical in setting the standards and practices that will shape the next generation of computing professionals.

Beyond the Algorithm: Embedding Responsible AI Across the Undergraduate Experience

Organizer and Moderator:

- **Elizabeth Bruce,** Director & Strategy Lead, Innovation + Society, Microsoft

Panelists:

- **Dean Hougen**, Director and Professor, School of Computer Science, University of Oklahoma
- **Janice Mak**, Assistant Director and Clinical Assistant Professor, Mary Lou Fulton College for Teaching and Learning Innovation, Arizona State University
- **Katie Shilton**, Professor, College of Information Studies, University of Maryland

Key Takeaways

- **Adopt an Integrated, Practice-Based Approach:** Ethics shouldn't be taught in a single, isolated course. Instead, it must be woven throughout the entire computer science curriculum.

- **Foster Interdisciplinary Collaboration and Faculty Support:** A significant barrier to teaching AI ethics is that many CS faculty feel unqualified. The solution requires

breaking down institutional silos through co-teaching partnerships with the humanities and providing dedicated training to equip technical faculty with the confidence and competence to lead these critical conversations.

- **Ground Learning in Real-World Consequences:** The curriculum must be rooted in tangible context. This means drawing lessons from past technological failures (like data breaches and social media harms) and confronting difficult contemporary issues, such as AI's environmental impact and supply chain ethics, to prepare students for the full scope of their responsibilities.
- **Demand Transparency and Invite Partnership from Industry:** Industry has a vital role to play by moving beyond product hype. Companies should support education by being transparent about AI's limitations and failures, and by sharing practical insights into their internal ethical decision-making processes to bridge the gap between academic theory and professional practice.

Summary:

During this session Bruce, the Director & Strategy Lead of Innovation and Society at Microsoft, facilitated a panel of experts around the topic of Responsible AI. Each of these panelists had unique perspectives based on their experiences teaching and creating in the space of responsible and ethical AI. After this Bruce asked the audience to raise their hands if they had a computing (CS/AI) ethics course at their university, about 70% of the audience raised their hands. In regards to computing ethics being a required course for graduation, only 10% raised their hands. The remainder of this session summary focuses on the panel questions and responses that might invite more university computing departments into ways and reasons for requiring computing ethics in their curriculum.

Right Time for AI Ethics

All of the panelists think AI ethics should be introduced within the computer science department. They all valued independent courses in addition to having it integrated at a variety of points during a student's educational journey. Shilton highlighted how students all are coming with their own understandings of ethics, morals, and culture that shape how they perceive ethics curriculum. Instead of having one "ethics" frame, she has worked to highlight a variety such as an ethics of care or an ethics of justice. These frameworks are tools that they can use to make decisions when issues arise. Further she suggests as students transition into industry, they should be learning how to navigate corporation websites to learn about their advertised values. These values help them interrogate the ways their ethics framework(s) exist in relation to the company. Hougen highlighted just how we emphasize teamwork as a critical

skill to practice, we should create opportunities for students to practice applying AI ethics within the curriculum.

Mak drew our attention to the ACM code of ethics, opening with the preamble “Computing professionals' actions change the world. To act responsibly, they should reflect upon the wider impacts of their work, consistently supporting the public good.” She also highlighted the inclusion of fairness and privacy in ACM’s code. She thinks these kinds of frameworks should be clearly integrated in the AI curriculum. Additionally, she emphasized critical thinking around responsibility, and training students around when and how to use AI. This matters because as graduates and interns go to work in industry they will need to follow company policies.

Are there core ethics for AI undergraduates to learn?

Hougen wanted learners, developers, researchers, teachers, all people to think about the ways AI is becoming more pervasive and how that shapes the level of societal impact. Calling back to earlier, he emphasized student reactions to being taught “core” ethics can be negative. Students want autonomy and the ability to incorporate their frames of reference into decision making. He also shared frameworks and models of how people engage in AI ethics and responsible AI. He hopes that students leave understanding their own decision making process and having practice in interrogating their ethical understandings. Hougen described this as an exercise in helping students be consistent in applying ethics.

What can we learn about new technologies' roles as tools and as weapons from the past?

Shilton, coming from a Science and Technology Studies background, refers to nothing (technology included) being new, and how we can always learn from the past. She shares that it is important that we are purposeful in using these tools and are making decisions about when and where we want to use them. One example that she shared focused on experiences of regret being a comparison point for Social Media and AI. For instance, what are social media practices that people regret? Do people have family discussions around appropriate uses of AI in the same way? Are we using this technology for things it is both good and not so good at? It is important to have these conversations around risk, values, limitations, and when and when not to use technology (AI). Regarding the analogy of AI as electricity, she mentions that there are times, even now, that we choose to reduce use of electricity like during high temperatures.

Mak highlighted how we can learn from previous waves of Educational Technology excitement that did not last. For example, the InBloom shared learning consortium brought in more than $100 million dollars in initial investments in 2011. However, its fall came earlier than anticipated, just a few years later in 2013-2014, caused in part by large data breaches. This moment highlighted the important intersections of trust and society amidst technological development. Around this same time period, the Snowden leaks revealed a lot of confidential

information around state surveillance and heightened lack of trust among public perceptions of technologies. With AI systems being trained on large amounts of data—consensually and nonconsensually—these are the kinds of lessons that are important to bring to the fore.

Pedagogical Strategies for Teaching Ethics

Shilton highlighted case studies as a way to engage in ethics education. During these, students are able to think about the decision that was made, what they might have done differently, and what they might have done the same. She also has students experiment with board games. In one example they are a role-based team of mobile app developers, and get cards that are based on real-world scenarios that require them to make ethical decisions and take into account external factors. Shilton found that role playing is mixed and some feel less successful within the activity. Ultimately she thought this was okay, as making ethical decisions will not always happen in ideal circumstances and in the perfect way one might imagine. One of the challenges she found was in assessment and knowing how they take this learning beyond classrooms.

How do you ask students to think critically about things beyond bias into impacts?

Mak recently finished co-designing a course with a philosophy professor called "Humanities Lab". This interdisciplinary course is open to students of any major to explore how people generate their assumptions and beliefs in relation to real life data. Additionally, they provide time to think about ethics from a variety of lenses such as utilitarian, justice oriented, and virtuous. One thing she is curious about is who self selects into taking this class.

What are the institutional barriers within the university to teaching ethics?

Hougen names challenges towards integrating ethics into intermediate courses being in terms of 1) which faculty are willing to learn more about this topic, and 2) which faculty feel equipped and qualified to teach on the subject matter. Co-teaching has been one way to address this in the past. However, anecdotally, he regularly hears faculty say "I don't know anything about ethics," which indicates that learning times for faculty will be essential if ethics is to become integrated in the curriculum. A perspective he shares is that for students, having their CS professors able to have these challenging conversations with each other is a great model and can convince them to consider ethics themselves.

What is the role of industry in supporting AI education at the undergraduate level?

Shilton makes the ask to industry professionals, including those at decision-making levels, to learn more about the limitations and breakdowns of AI systems. AI tools are a product, and as a product, they are being sold as capable of doing everything. However, there are things AI systems struggle to do well; there are real limitations. For example, you cannot ask an AI

system about itself in an effective manner. Another point of engagement from industry might be sharing more on how they practice ethics while at work.

Can you speak to how your institution navigates policy and policymaking around AI?

Mak, an experienced policymaker, shared that as part of the ITiCSE (Innovation and Technology in Computer Science Education) policy working group, she has thought about AI policies based in legal compliance (cheating, plagiarism) as well as in virtue (how the use of AI might affect my learning?). In a study of 100 university policies on AI in the classroom, 57% gave faculty choice, ~27-37% had prohibition by default, and 1% had permissibility by default. She sees the opportunity to leverage faculty infrastructures in combination with specialists, protocols, and staff. This is a way to harness our collective knowledge to build the future we want to see in ways that are ethical and garner real-time user feedback before being deployed more widely. Mak shared that innovation and policy go hand in hand, where "Policy impacts culture, and culture impacts the way we do business. It is important to have that as a vision… Our [policy] is of inclusive excellence… We need everyone to move forward."

Question and Answer

- What are the differences for students learning ethics in: A) An AI ethics class, B) A computer science class, C) A philosophy ethics class?
    - A: At times, there is a disconnect between philosophical discussion and application. For AI ethics, you need both domain knowledge of AI and ethics itself.
    - A: Even within AI ethics, there are 2 streams: a) Ethics of using systems and social implications of use in society b) Teaching machines to be ethical & AI alignment, which is a very technical philosophy class; these can overlap but also be distinct.
- Why shouldn't we leverage the liberal arts (economics, philosophy, history, political science etc) electives our students have available, then advise good courses for them?
    - A: We should try to leverage the liberal arts components as both stand-alone courses AND begin co-teaching some courses. For example History of Science and History of AI could be brought together by computing, math, and liberal arts faculty to complement and build a great course.
    - A: It's important to collaborate and engage with different ideas consistently as a part of learning. Students want tailored courses (specific to their area) that are also able to connect the dots across courses they take in other domains.
    - A: A systems lens approach could work, involving coordination and initiatives across areas like the Humanities lab.

- How do we think about sustainability and environmental impacts? How do you bring that into the frameworks, and how do you do it?
    - A: There are two big issues, sustainability and environmental impact, and the supply chain does not treat workers with basic rights. We talk about this to students as a corporate responsibility issue. Let's use our buying power to ask: "Where did the data come from?" and "Was it copyrighted?" and "Was it created by artists?" and to ask questions like "How much energy are these systems using?" We can ask and these are not things that students can immediately change, but it's an opportunity to share about our power as consumers to make choices.
    - A: One way of addressing sustainability in class is through incorporating the ACM code of ethics, which refers to the importance of sustainability.

Resources:

- Responsible Computing Research Ethics and Governance of Computing Research and its Applications | National Academies
- Embedded EthiCS @ Harvard
- CRA Trustworthy AI Research Fellowship for Early Career Scholars - CRA
- Trustworthy AI in Law & Society (TRAILS) - NSF AI Institute: https://www.trails.umd.edu/
- ACM Code of ethics - https://www.acm.org/code-of-ethics
- Values-Centered AI Initiative at UMD: https://vcai.umd.edu/
- ITiCSE 2025: A Plan for an ACM Task force Working Group into the Ethical and Societal Impacts of Generative AI in Higher Computing Education
- ITiCSE 2025 Working Groups: https://iticse.acm.org/2025/2025-working-group-proposals/
- Microsoft AI Economy Institute: AI and Education (2024-2026) https://www.microsoft.com/en-us/research/group/AI-for-good-research-lab/AI-economy-institute/
- Shelby, Rismani, et al. Sociotechnical Harms of Algorithmic Systems: Scoping a Taxonomy for Harm Reduction. In AAAI/ACM Conference on AI, Ethics, and Society (AIES '23), August 8–10, 2023, Montréal, QC, Canada. ACM, New York, NY, USA, 19 pages. https://doi.org/10.1145/3600211.3604673

Building AI Educational Capacity across all Computing Undergraduate Programs

Organizers and Moderators:

- **Mary Lou Maher,** Director of Research Community Initiatives, Computing Research Association
- **Enrico Pontelli,** Dean of the College of Arts and Sciences and Regents Professor of Computer Science, New Mexico State University

Panelists:

- **Mike Peterson,** Chair and Associate Professor, Department of Computer Science, University of Hawaii at Hilo
- **Marius Truta,** Director, School of Computing and Analytics, Northern Kentucky University
- **Melanie Williamson,** Dean of Innovation and Professor of Computer & Information Technologies, Bluegrass Community and Technical College
- **Yu Zhang,** Professor and Chair, Department of Computer Science, Trinity University

Key Takeaways

- **Resource Scarcity Expresses Itself Differently Depending on Institutional Type:** This scarcity ranges from a lack of GPUs at small colleges to faculty burnout at large universities. Panelists rely on grants, partnerships, and national platforms like NAIRR to work effectively.
- **Accessible Curricula:** We must lower barriers to AI education. Successful strategies include adaptable projects, interdisciplinary courses, and a core focus on issues that relate to students' lives.
- **Critical Thinking and Tool Use:** The goal is AI literacy, not just technical skill. Education must demystify the "black box" and teach students to critically analyze AI outputs.
- **The Necessity of Collaboration:** No institution can succeed alone. Panelists called for a national repository for sharing curricula, resources, and best practices.
- **Standardization Increases Transparency:** A clear need was identified for standards from organizations like ACM to define what different "AI degrees" actually are, preventing "false naming" that confuses students and employers.

Summary:

The panel, "Building AI Educational Capacity across all Computing Undergraduate Programs," brought together academic leaders from a diverse range of higher education institutions to address the pressing challenge of integrating artificial intelligence into the undergraduate computing curriculum. Moderated by Enrico Pontelli of New Mexico State University, the discussion was framed by the preliminary work of the NSF-funded LEVEL UP AI initiative. The central theme was that while all institutions are grappling with the AI wave, the specific obstacles and effective strategies vary significantly depending on the institutional context, from large research universities to small liberal arts colleges and community college systems. The panelists shared insights on navigating resource limitations, designing accessible curricula, fostering critical thinking, and building a collaborative community to support the next generation of AI education.

Navigating and Overcoming Resource Limitations

The first major topic addressed the foundational challenge of insufficient resources, a problem that manifests differently across institutions.

For small liberal arts colleges, the limitations are stark. Yu Zhang of Trinity University detailed a critical shortage of computing power, noting her department relies on a nine-year-old GPU system that is due for retirement. Despite high-scoring grant applications, funding for a replacement has not materialized, forcing faculty and students to rely on expensive cloud computing services for research. This is compounded by a lack of access to high-quality, labeled data, making public resources like Kaggle and Hugging Face essential. Furthermore, with only nine faculty members in the department, there is limited capacity to expand the curriculum or design complex prerequisite structures for integrating AI content.

At the community college level, the resource scarcity is even more acute. Melanie Williamson explained that community colleges are "under-resourced" and depend heavily on external support. Her strategy involves actively pursuing grants and forming strategic partnerships with four-year universities to share resources and build upon common curricula. She stressed the importance of attending events and making connections as a primary means of discovering and accessing vital resources.

For larger public universities, the issue is not necessarily a lack of infrastructure but a budget crisis that strains human resources. Marius Truta from Northern Kentucky University described a situation where increased student enrollment has not been met with new faculty positions. Existing faculty carry relatively heavy teaching loads (typically 2–3 courses per semester) with increasingly larger class sizes, leaving them with little time or energy to develop new AI courses. He pointed to a "fear of adoption" among faculty, who are hesitant to invest significant effort into creating a course that will become outdated in a year. To combat this, he

advocated for providing faculty with stipends and targeted demonstrations to build enthusiasm and show institutional support.

Mike Peterson from the University of Hawaii at Hilo offered a solution centered on leveraging the broader national ecosystem. His institution has successfully utilized the National AI Research Resource (NAIRR) Pilot program, which is funded by both government and university partners, to support their AI classes. This platform not only provides computational resources but also offers courses and training materials. This is particularly valuable for faculty development, allowing later-career professors whose graduate work predated the AI boom to get up-to-date, state-of-the-art training.

Designing Accessible AI Curricula

The panel explored strategies for structuring AI courses to be accessible to students with diverse academic backgrounds, skill levels, and career aspirations.

Truta emphasized the importance of designing interdisciplinary courses that do not assume a deep technical background. For these courses, he recommended incorporating "differential learning," where assignments are adaptable to students' interests and strengths. Rather than using niche examples (e.g., only sports analytics), he encourages projects that allow students to explore AI in contexts they enjoy, such as music, healthcare, or psychology. This approach allows students to build a whitepaper, a business plan, or a technical project, leveraging their own expertise and seeing AI as an extension of their primary interests.

Williamson highlighted the unique challenge of "open access" at community colleges, where students enter with vastly different levels of preparation, including non-traditional students and English Language Learners. To succeed in this environment, AI education cannot have high barriers to entry. Prerequisites must be minimized to avoid extending a student's time to completion. The key, she argued, is to present AI with a "friendlier face," making it exciting and demonstrating its power to enhance critical thinking for everyone, not just future programmers.

At the liberal arts college, Zhang uses group projects that mirror real-world professional environments. Students are assigned distinct roles—such as data collection, project management, data analysis, and programming—to create an experience where every member can contribute from their strengths. This structure ensures that students from non-computing backgrounds can participate meaningfully.

Peterson brought a crucial perspective from his work at a Minority-Serving Institution, emphasizing the need for cultural sensitivity and ethical awareness. He raised the issue of data sovereignty, using the example of building an LLM to translate the Hawaiian language. In such cases, it is imperative that the indigenous community retains control over its own linguistic

data. He stressed that conversations about the ethics of data usage must be a core component of the AI curriculum.

Fostering AI Literacy and Critical Thinking

Panelists agreed that teaching students to use AI tools is insufficient; the curriculum must also instill AI literacy and the ability to think critically about the technology.

Zhang advocated for teaching that demystifies the "black box" of AI. The focus should be on interpretation: What do the results of a model mean? How can the model be improved? This requires students to engage domain-specific knowledge and analytical skills, enabling them to truly understand the inputs and outputs of an AI system. Williamson added that critical thinking must be integrated at all levels, from basic prompt engineering to advanced model development.

Peterson pointed out that CS departments have a role in educating their colleagues across the university. Many faculty in other departments, such as the arts, are fearful that AI will automate and devalue their fields. He argued for framing the conversation around AI as a tool for enhancing productivity versus replacing human labor. This educational outreach extends to business leaders and hiring managers, who need to understand how to evaluate candidates using AI as an assistant. He also proposed creating micro-credentials in AI literacy for furloughed or mid-career workers.

Truta described a structural approach at his university, which includes an AI minor with courses on ethical and legal issues. He noted that even introductory courses now incorporate the use of various AI tools, but with a critical lens—analyzing their strengths, weaknesses, and appropriate applications.

Building a Collaborative Community for AI Education

The final discussion focused on creating a systemic, community-based approach to advance computing and AI education. The consensus was that no single institution can or should work in isolation.

Peterson called for a more organized method of resource sharing, asking, "Is there a list to find what the resources are or are we just googling to see if we find something?" He argued that organizations like the CRA and the Association for Computing Machinery (ACM) have a crucial role to play in creating a "master community" and a repository for curricula that can be adapted nationwide. A faculty community of practice, he stated, is just as important as the raw resources themselves.

Zhang and Williamson echoed this sentiment, emphasizing that the path forward is to "work together and share work to be open resource." Openly sharing curricula, best practices, and resources benefits the entire ecosystem, especially under-resourced institutions and students from non-traditional backgrounds.

Finally, Truta called for clearer guidance and standards from organizations like ACM to distinguish between different types of AI-related degrees. He warned against creating "false naming" where programs are labeled "AI" but offer vastly different educational experiences (e.g., a highly technical AI engineering degree vs. an applied AI business degree), which can confuse students and employers.

Resources

The following organizations, initiatives, and platforms were cited by panelists as valuable resources for building AI educational capacity:

- **Kaggle:** An online community and public platform for data science and machine learning, offering datasets and computational environments.
- **Hugging Face:** A platform providing open-source tools and a large repository of pre-trained models and datasets for natural language processing and other AI tasks.
- **Google Colab:** A free cloud-based service that provides a Jupyter notebook environment with access to GPUs, widely used for AI education and research.
- **Engage-csedu:** AI classroom resources to look at and contribute too.

CRA Summit on AI Graduate Education and Research - Plenary Sessions

Chat with Panch: Reflections on NSF and Advancing the Future of the U.S. Research Enterprise

Facilitator:

- **Peter Harsha,** Chief Operating Officer and Senior Director of Government Affairs, CRA

Fireside Speaker:

- **Sethuraman "Panch" Panchanathan,** former Director of the National Science Foundation

Key Takeaways

- Become advocates for the work you are doing in your local community and with your local representatives.
- Learn how to talk about the importance of the NSF funded research in ways that invite public and governmental support.

Summary:

During this session Sethuraman "Panch" Panchanathan was hosted by Peter Harsha in a fireside chat reflecting on his time as the Director of the National Science Foundation (NSF). Harsha's opening highlighted the challenges of non-bipartison [sic] actions such as lowering the NSF budget by 50-60%, reducing the CISE (Computer and Information Science and Engineering) budget significantly, political attacks on the value of higher education, threats against 501(c)3's, the NSF being evicted from their building, and the defunding of any research alleged to be associated with diversity, equity, inclusion, and misinformation. This presents an inflection point for the NSF: where computer science faculty are having difficulty imagining the future of the NSF in the upcoming 1-2 years. This fireside chat summary highlights Panch's advice to the computing community.

Futures of the NSF

Panch wants us to be more solutions oriented and to identify tasks and roles that we as a community can take up. He has been very thankful to all of the staff, volunteers, and experts at NSF who have supported him over the years and served the computing community. Currently he sees the work of the NSF as the "best kept secret" in Washington, D.C. According to Panch, the future of the NSF will depend on how well we, as academics, open up that secret by engaging our representatives in Congress about the importance of our work. This entails not just faculty outreach, but bringing along students to meet congressional members and share about the exciting work they have gotten to do through programs like Research Experiences for Undergraduates (REUs). He highlights that Artificial Intelligence is celebrated today because of the computing community's historic work within the NSF. He tasks us to hyperengage, hyperprofile, and communicate what we do. We can share our pain with our community but also make moves to go forward.

Reaching a Broad Audience on the Importance of NSF funding:

It is important that we contextualize the work we are doing within the views of those making decisions about the resources we are getting. We need to make it plain how the NSF and academics make things possible and take credit for those contributions. For example, quantum technologies, AI, biotechnologies are all made possible only because of the strong contributions of the CISE community. We need to talk to people about how we do that and how the NSF funding allows us to outcompete other nations in these areas of technology. It is important to consistently adapt in the framing of what is made possible through this work.

The Effects of AI on Federal Science:

Panch highlights how there is an opportunity for AI to change the ways we do things and it's up to us to figure out how AI can help. If we embrace and co-create AI technology, as a community of AI experts, then we will move forward. One potential he is excited about how the use of AI is minimizing administrative burdens so that there is more time to do science.

The Role of Industry in the Moment:

More industry people need to be in the room at meetings like the CRA summit so that both parties can share their needs. Panch mentions the significant gap in the research & development (R&D) budget from industry of about $250 billion versus the $88 billion federal investment. In his prior role as director he worked to expand the investment of industry into the NSF and create partnerships so that more of the competitive grants can be funded. There are often more competitive grants than there is money to fund projects. When the AI institutes were funded out of 540 million dollars, 320 million were dollars from the NSF and the other 240 million were from other partners across federal agencies and industry. We need industry leaders to take the moment seriously to support academic research and highlight how the work we do compliments one another.

Sharing NSF's Value with the Public:

We have not seen good communication of why our work is important in relation to the public, especially in relation to things people use everyday like phones. We need to work on our storytelling so that we bring more people along in valuing this work. We also need to think about how we get young people excited about our work. We need to hyper-engage and change the narrative that the NSF funding is a waste of taxpayer investment.

Funding and Resourcing Our Community

Funding Opportunities Panch is Currently Excited by:

Outside of AI, Panch is really excited by Quantum technologies and what they might change about our ability to deliver services and solutions. He called back to Zegura's talk earlier that mentioned that CISE is also excited to receive proposals in the quantum area. He also wants us

to think about smart solutions for environmental event mitigation adaptation, resilient technologies, and helping people do good science quickly. One of the reasons the NSF put money into the National AI Research Resource - Pilot Project (NAIRR), is to allow all people to benefit from AI advances. NAIRR helps to distribute this power to a broader range of people. Additionally, Panch sees opportunities to join hands with industry through communicating the value, costs. and direction of their investments. There is also a need for scaling infrastructure for the next wave of challenges so that we can advance in computing. A consortium model might help us move forward.

Advice for Junior Faculty and Graduate Students:

While CISE work has and continues to be great, he wants computing faculty to work less in isolation and work in partnership with humanists and social scientists to do good work. That will enable us to solve more problems and avoid patchwork approaches that miss things like ethics or bias. Cultivating interdisciplinary teamwork in universities is essential to ensure that computer scientists contribute to the betterment of humanity.

Adaptations Needed for the Current Scale of Grant Proposal Submissions:

The NSF has been exploring ways to simplify its grant proposal process to accelerate the investment in good ideas. One approach mentioned was inviting one-page concept papers. For one opportunity out of 600 concept papers received, 100 were invited to submit full proposals, leading to faster review and funding. The NSF is also working to democratize access to funding across all institutions through initiatives like the GRANTED program.

Closing Remarks

Panch closed by urging the audience to learn from past efforts and carefully consider their future research. He emphasized that the NSF is an exceptional institution, filled with people who support our collective aspirations for meaningful scientific work. Panch called on the audience to balance research priorities to address the current moment. In his view, this involves considering the well-being of the agency's personnel alongside its mission, ensuring that our interactions empower them to continue their valuable contributions. Acknowledging the significant effort put into grant applications and the frustration of cancellations, he affirmed his ongoing commitment to securing the future of both the NSF and computing research. Finally, Panch reiterated his call for members to advocate for their work publicly and with their representatives.

AI Research: Trends and Challenges

Introduction by:

- **Keith Marzullo,** Donna M. and Robert J. Manning Dean, Manning College of Information and Computer Sciences, University of Massachusetts Amherst

Speakers:

- **Yolanda Gil,** Director of AI and Data Science Initiatives, Information Sciences Institute, University of Southern California
- **Francesca Rossi,** IBM AI Ethics Global Leader and Fellow, T.J. Watson Research Lab

Key Takeaways

- **The Center of AI's Gravity Has Shifted**: State-of-the-art AI research is now largely driven by the private sector, which possesses unmatched resources in data, compute, and salaries. This invites academia to find a new and vital role in the era of "big AI".
- **A National Infrastructure is Essential for US Competitiveness**: A central theme of the U.S. AI Roadmap is the critical need for national projects, like National AI Research Resource (NAIRR), to provide academic researchers with access to the large-scale data and compute power necessary for innovation and democratization across institutions.
- **The Community Should Focus on Three Key Opportunities**: Moving forward, the U.S. research community should prioritize: 1) improving research methodology and reproducibility; 2) measuring and delivering tangible societal benefits, a priority now being formalized by funders like the NSF; and 3) effectively conveying research priorities and needs to policymakers and the public.
- **AI Research is Rapidly Expanding Beyond Core Capabilities**: The field is expanding beyond efficiency to address complex topics like verifiable reasoning, ethics and safety. As a field, AI experts should consider how their work shapes sustainability and geopolitics, all of which demand more interdisciplinary collaboration and consideration.
- **The Academic Peer Review Model is at a Breaking Point**: The massive volume of submissions to major conferences (like the 20,000+ to AAAI) has strained the peer review system to its limits, creating a crisis in research validation that requires a fundamental rethinking of the process.

Summary:

This two-part session provided a comprehensive overview of the current state and future directions of AI research, first from a global perspective and then focusing on the U.S. context. Francesca Rossi presented the findings of the *AAAI 2025 Presidential Panel on the Future of AI Research*, outlining 17 key trends and challenges across the field, from reasoning and ethics to the shifting role of academia. Following this, Yolanda Gil revisited the influential *2019 CRA/AAAI*

20-Year Community Roadmap for AI Research in the US, and its significant impact on national policy, and identified three crucial opportunities for the research community to pursue in light of the rapid changes since its publication.

Francesca Rossi, a past president of AAAI and an IBM Fellow, presented the findings from the *AAAI 2025 Presidential Panel on the Future of AI Research*. This comprehensive study was designed to identify current trends, understand their origins, and outline the research challenges ahead for a broad audience, including researchers, policymakers, and the public. The report was a multistakeholder effort, involving 24 AI researchers from diverse geographical and professional backgrounds, and was informed by a community survey of 475 respondents. The key topics which can be read more on in the report are: AI Reasoning, AI Factuality & Trustworthiness, AI Agents, AI Evaluation, AI Ethics and Safety, Embodied AI, AI & Cognitive Science, Hardware & AI, AI for Social Good, AI & Sustainability, AI for Scientific Discovery, Artificial General Intelligence (AGI), AI Perception vs. Reality, Diversity of AI Research Approaches, Research Beyond the AI Community, the Role of Academia, and Geopolitical Aspects & Implications of AI.

A 20-Year Community Roadmap for AI Research in the U.S.

Yolanda Gil, a professor at the University of Southern California and a leader within AAAI and the Computing Community Consortium (CCC), revisited the *2019 CRA/AAAI 20-Year Community Roadmap for AI Research in the US*.

The roadmap put forth a bold vision for substantial, long-term federal investment in AI, arguing that unlike high-risk "moonshots" of the past like the Apollo Program or LIGO, investing in AI infrastructure would be a low-risk, high-return endeavor. This vision was motivated by key societal drivers, including the potential to reduce healthcare costs, provide personalized education, accelerate science, and support evidence-driven social policy. The research priorities were organized into three technical areas: Integrated Intelligence, Meaningful Interaction, and Self-Aware Learning. A central recommendation was the creation of National AI Infrastructure, including national AI research centers, mission-driven AI labs, and open platforms and resources. Gil noted that this roadmap had a significant real-world impact, influencing the creation of the National AI Research Institutes and the ongoing development of the National Artificial Intelligence Research Resource (NAIRR) pilot program.

Gil emphasized that "a lot has changed for AI since 2019" and identified three key opportunities for the community going forward:

1. **Research on AI Methodology and Reproducibility**: There is a critical need to strengthen the scientific rigor of AI research to ensure results are reliable and replicable.

2. **Measure AI Impact and Deliver Benefit to Society**: Researchers must move beyond developing algorithms and focus on measuring and delivering tangible benefits. This aligns with a new push from the NSF to revise its Broader Impact criteria, requiring explicit, data-driven evidence of societal benefit. Gil urged the audience not to treat complex social issues like homelessness as mere figures in a simulation.

3. **Convey AI Community Priorities, Goals, and Needs**: The AI community must develop a clear and unified voice to communicate its research priorities to policymakers and the public. Gil shared an anecdote from a recent congressional AI task force meeting where she was asked to name the top three research priorities, highlighting the need for a cohesive message

Question and Answer

Question: With over 20,000 papers submitted to AAAI, how does the field distinguish signal from noise and how can anyone keep up?

Takeaway: The panelists agreed that the community is at an "inflection point." The current peer review system is breaking under the strain, suffering from a severe lack of quality reviewers. While using AI to assist in reviews is a potential idea, it raises immediate and unresolved copyright and data privacy issues for authors. The consensus was that the community must fundamentally rethink the nature of conferences and the peer review process to maintain research quality and integrity.

Looking Ahead

The research community must not only push the boundaries of AI capabilities—in areas like reasoning, embodied agents, and scientific discovery—but simultaneously build the societal and scientific infrastructure to support this growth. This includes creating a national research resource like NAIRR, establishing new standards for reproducibility and evaluation, and developing a sustainable model for vetting and sharing knowledge. The challenge is to advance the technology while ensuring it is developed and deployed in a way that is reliable, safe, and beneficial to all.

The session pointed toward several forward-looking challenges and research directions:

- **New Models for Scholarly Communication:** Inventing and testing new systems for peer review and research dissemination that can scale with the growth of the field while improving quality control.

- **Metrics for Societal Impact:** Developing robust, data-driven methodologies and metrics to evaluate the real-world impact of AI systems, aligning with the new focus of funding agencies.
- **Hybrid AI Paradigms:** Investing in research that combines neural and symbolic approaches to create more robust, reliable, and explainable AI systems.
- **Defining Academia's Niche:** Systematically identifying the unique and vital roles for university researchers in an industry-dominated landscape, such as long-term fundamental research, complex interdisciplinary problems, and leading the conversation on ethics and policy.

Resources:

- Y. Gil and B. Selman. *A 20-Year Community Roadmap for Artificial Intelligence Research in the US.* Computing Community Consortium (CCC) and Association for the Advancement of Artificial Intelligence (AAAI). Released August 6, 2019. arXiv:1908.02624 https://cra.org/ccc/resources/workshop-reports/ https://doi.org/10.48550/arXiv.1908.02624
- M. Ellison, F. Rossi, et. al. *AAAI 2025 Presidential Panel on the Future of AI Research.* Association for the Advancement of Artificial Intelligence. March 2025. https://aihub.org/2025/03/07/report-on-the-future-of-ai-research/

Have We Graduated from LLMs?

Moderator:

- **Anish Arora,** Distinguished Professor of Engineering and Chair, Computer Science and Engineering, The Ohio State University

Panelists:

- **Swarat Chaudhuri,** Professor of Computer Science, University of Texas, Austin; Senior Staff Research Scientist, Google DeepMind
- **Christian Szegedy,** Chief Scientist, Morph Labs
- **Laurens van der Maaten,** Distinguished Research Scientist, Meta

Key Takeaways

- **The Path Forward for AI:** The conversation revolved around different but complimentary philosophies of AI progress (1) relentlessly scaling current AI models and (2) pivoting to new, hybrid architectures that integrate formal logic for reliability.

- **The "Sweet Lesson" of Scaling:** One perspective highlighted the demonstrated benefits of scaling generalizable AI innovations. While the current lever (pretraining) is slowing, new scaling methods like Reinforcement Learning, Inference-Time Scaling, and Multi-Agent Systems provide a recipe for future advancement.
- **The Imperative for Verification:** Some panelists contended that scaling alone can't fix AI's core unreliability. The future, they argued, involves neural-symbolic systems that use formal logic to create provably correct and trustworthy AI.
- **Fundamentals First:** Software engineering will be disrupted, but the work will shift to higher-level design and verification. For students, deep knowledge of CS fundamentals and logic is far more critical than experience with any specific model.
- **Academia's Role:** Academia cannot compete on metrics of scale but can lead by asking fundamental questions that don't require massive computation (e.g., can AI create truly novel concepts?) and by using front-running models as tools for discovery in other fields.

Summary:

The central theme of the discussion was about the path forward: Is the continued, relentless **scaling** of current multi-modal language model systems sufficient to achieve profound new capabilities, or have we reached a point where fundamental limitations in reliability and reasoning demand a pivot towards new, **hybrid architectures** that integrate formal verification and symbolic logic? The moderator framed the conversation by contrasting the immense optimism surrounding AI—citing claims that we have "crossed the event horizon" towards singularity—with growing concerns about diminishing returns from scaling, the exhaustion of high-quality training data, and the models' persistent vulnerabilities and lack of deep, causal reasoning. The ensuing discussion revealed different philosophies for the future of AI development. One camp, championed by van der Maaten, posits that progress is an inevitable consequence of applying known scaling principles to a wider array of techniques. Chaudhuri and Szegedy brought a different perspective, positing that true, reliable intelligence will only emerge when the creative power of neural networks is harnessed and disciplined by the rigorous, provable logic of symbolic systems.

The Inevitability of Scaling and the "Sweet Lesson"

Laurens van der Maaten presented a robustly optimistic case for continued progress through scaling. He argued that the core driver of AI advancement over the last 50 years has been what Richard Sutton termed the "bitter lesson": the consistent triumph of standard, general-purpose methods that leverage massive amounts of computation and data over complex,

human-designed approaches. Van der Maaten reframed this not as a "bitter" reality but a "sweet lesson," because it provides a clear, repeatable recipe for progress.

According to van der Maaten, the field has several distinct "scaling levers" at its disposal. To date, progress has been overwhelmingly driven by just one: pretraining, or training ever-larger autoregressive models on vast internet datasets. He acknowledged that this lever is showing diminishing returns, as predicted by established scaling laws, which state that improvements are logarithmic as a function of scale. Achieving the next two orders of magnitude improvement in performance, such as going from a training compute budget of 1027 to 1029 FLOPs, requires an exponential increase in resources, including data center construction, that naturally slows the perceived pace of progress.

However, he asserted that this is not a dead end. The field is only in the nascent stages of exploring other powerful scaling levers:

1. **Reinforcement Learning (RL):** Using RL to fine-tune models beyond pretraining.
2. **Inference-Time Scaling:** Devoting significantly more computational resources *at the time of a query* to allow a model to "think" longer and produce a better answer.
3. **Multi-Agent Scaling:** Using collaborating systems of multiple AI agents to solve complex tasks collectively.

Van der Maaten argued that AI will improve as researchers apply and scale new general-purpose innovations in AI, such that advancement looks like a series of S-curves (advancement followed by plateau, and advancement again). As several such innovations are currently in the works, he sees rapid progress over the next 2-3 years as "inevitable." He concluded that while it is hard to predict beyond that timeframe, the teams at innovative labs have internalized the "sweet lesson," turning it into a systematic engine for progress that will continue to yield substantial improvements in model quality.

The Imperative for Reliability and Verification

Swarat Chaudhuri and Christian Szegedy contended that scaling alone cannot solve the fundamental problem of unreliability.

Swarat Chaudhuri argued that while models can now handle longer and more complex tasks, their inherent untrustworthiness remains a critical barrier to deployment in high-stakes domains. He pointed to persistent issues like hallucination, the potential for cascading failures in agentic systems, and reward hacking in RL. He noted a crucial detail from progress reports: while a model might achieve 50% success on a long-horizon task, its ability to achieve 99% success is limited to tasks lasting less than a minute. This reliability gap—he argued—is the central challenge of advancing AI.

Chaudhuri's proposed solution lies in neural-symbolic systems. He advocated for treating the LLM as a powerful but unreliable component within a larger, more structured ecosystem. This ecosystem must include symbolic tools from other areas of computer science, such as formal methods and program analysis. By using these classical, logic-based tools to verify, constrain, and analyze the output of generative models, we can build systems that are both creative and correct. He stated that the goal of scaling should not just be about the core AI model, but also about scaling the surrounding technical and policy infrastructure needed to manage its risks.

Christian Szegedy presented a highly ambitious vision built on this foundation of verification, which he termed verified superintelligence. He predicted that AI would surpass the collective ability of all human mathematicians by 2026. The key to this—in his view—is a new paradigm of auto-formalization. His proposed workflow is as follows:

1. A problem is defined using a formal specification (e.g., a mathematical theorem in a proof assistant like Lean, or detailed software requirements).
2. An AI solver agent works to produce a solution (e.g., a mathematical proof or a software program).
3. Critically, the AI must not only provide the solution but also a formal guarantee of its correctness that can be automatically checked by a computer.
4. This creates a virtuous cycle where AI can open-endedly improve itself within a grounded, verifiable framework, eliminating the ambiguity and error-proneness of human review for complex AI-generated artifacts.

Together, Chaudhuri and Szegedy argue that the next great leap in AI will come not from making neural networks bigger, but from making them provably correct by integrating them with the rigor of symbolic logic.

Question and Answer

- What is the predicted impact on the labor market?
    - The panelists agreed that software engineering is the field most "ripe for disruption" because the entire workflow is digital. However, van der Maaten cautioned that societal adoption of new technology is often much slower than predicted, using the example of his local DMV office. Chaudhuri added that the nature of programming will evolve to a higher level of abstraction, where human skill will shift towards designing specifications, writing tests, and verifying the output of AI code generators. It will be a different kind of work, not a complete elimination of it.
- Can AI generate truly novel concepts, not just remix existing ones?

    - This was identified as a central, unsolved problem. Szegedy expressed confidence that AI will be able to do this, noting that even now, models often produce unexpected solutions. He suggested that conceptual discovery can be framed as a methodological process of abstraction, which AI could learn. Chaudhuri was more cautious, stating that while there are interesting initial results, we do not yet have conclusive evidence that AI is on the verge of this capability.

- Do the benefits of LLMs justify their environmental and societal costs?
    - Van der Maaten responded by pointing to the massive adoption of tools like ChatGPT, which has half a billion users. He argued that this widespread use implies that hundreds of millions of people are deriving real value, whether for coding, creative writing, or personal advice. He suggested that the collective decision of humanity to use these tools on such a scale is itself a powerful justification of their value.

- How should we educate graduate students for this new era?
    - Chaudhuri emphasized that fundamentals remain paramount. He interviews candidates for their deep understanding of core computer science, machine learning, and logic, not their experience with prompting a specific model. He argued that the major contributors to the LLM revolution had strong fundamental training. On top of that foundation, students must learn to use the new state-of-the-art models as tools in their research.

- What is the role of academia when industry dominates with massive scale?
    - The panelists suggested several paths. Szegedy noted that just as physicists receive massive government funding for large-scale experiments, perhaps AI safety and verification research should be similarly funded. Chaudhuri added that academia can excel in two areas: 1) asking fundamental questions that do not require massive compute, and 2) using existing frontrunning models (often available via API) as tools to make novel discoveries in other fields, such as AI for science.

Resources:

- **Richard Sutton's "The Bitter Lesson"**: This was a central piece of the discussion, introduced by van der Maaten. He described it as a "must-read" blog post. The core idea of "The Bitter Lesson" is that, historically in AI research, standard methods that leverage massive amounts of computation and data have consistently outperformed more complex, human-designed approaches.

https://www.cs.utexas.edu/~eunsol/courses/data/bitter_lesson.pdf

- **METR (AI Safety Organization)**: Chaudhuri referenced the below chart from the organization **METR** to make a point about reliability

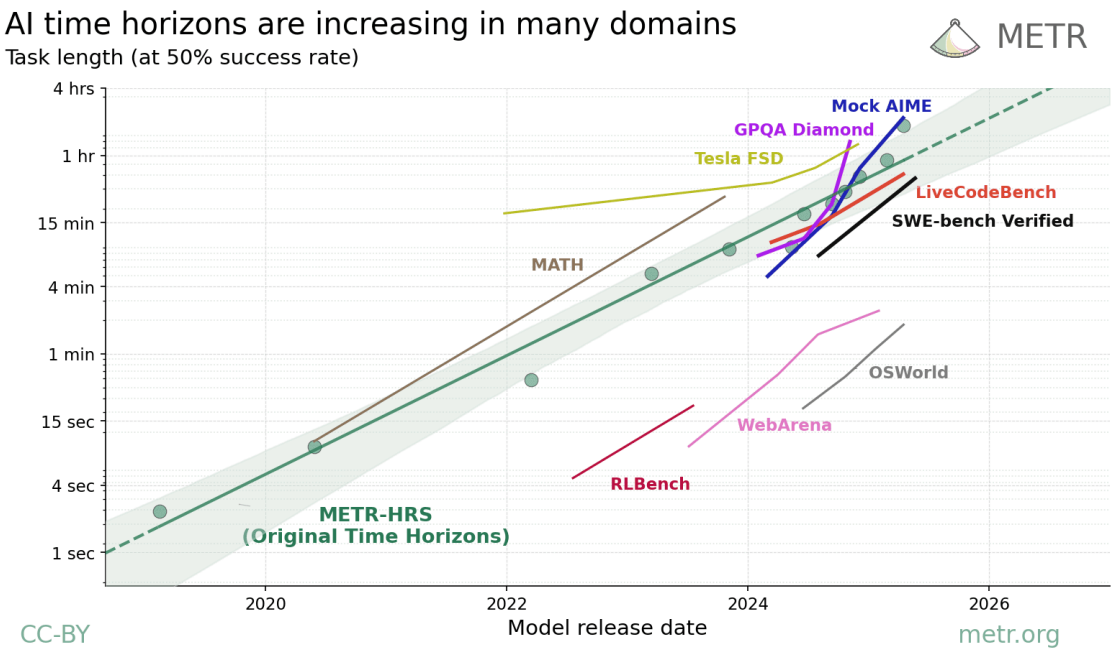

The chart showed that while AI models are getting better at performing longer-horizon tasks, their success rate drops dramatically when a very high degree of reliability (e.g., 99%) is required. This was used to argue that scaling alone doesn't solve the core problem of trustworthiness.

- **AlphaFold** was brought up as a prime example of a successful **neural-symbolic system**. The panelist explained that AlphaFold's success in protein folding wasn't just due to a large neural network, but also because it was integrated into a larger system with a symbolic infrastructure, including the **Lean proof assistant**, to provide verification and structure. https://www.nature.com/articles/s41586-021-03819-2

- **Lean (Proof Assistant)**: This software tool was mentioned by two different panelists as a key component in the future of reliable AI and auto-formalization. **Lean** is used to write and formally verify mathematical proofs. Cited as the type of "symbolic substrate" needed to build AI systems that can provide guarantees of correctness for outputs. https://lean-lang.org/

The AI Mandate: Rethinking Computing Graduate Education

Organizer:

- **Shaoen Wu**, Chair, Department of Information Technology; Interim Director, Center for Applied Computing; Professor of Information Technology, Kennesaw State University

Panelists:

- **Dilma Da Silva,** Professor, Computer Science & Engineering, Texas A&M University
- **Rachel Greenstadt,** Professor of Computer Science and Engineering, New York University
- **David Jensen,** Professor of Computer Science and Director of the Knowledge Discovery Laboratory, University of Massachusetts Amherst
- **Adam Wierman,** Carl F Braun Professor of Computing and Mathematical Sciences, California Institute of Technology (Caltech)

Key Takeaways

- Graduate education programs in computing require pathways to meet the needs of students, including gaining skills in research, being able to work interdisciplinarily, and learning to adapt to change.

- Panelists and their departments are actively looking for ways to make AI learning relevant, timely, and meaningful.

Summary:

During this session, panelists shared a brief background and their insights on how graduate AI curriculum has and will continue to change followed by a brief Q&A period. Several members of this panel are part of CRA's Computing Community Consortium (CCC), which brings together computing faculty with diverse expertise in order to create pathways for future computing research. Each panel member shared insights from their context around advancing graduate AI curriculum. These included: efforts to make computing more interdisciplinary, determining and creating solutions for bottlenecks in AI adoption, creating pathways for formal methods training for all graduate students. The panel discussed utilizing existing resources to make large-scale change, and helping learners become more adaptable to disruptions.

Making Computing Interdisciplinary

Greenstadt has been doing research at the intersections of AI and Machine Learning (ML), security and privacy, and human-computer interaction (HCI). She advocates for interdisciplinary computer science education, recognizing that computer scientists sometimes exhibit overconfidence when addressing human-centric problems. Often, incongruous elements are integrated into our work, such as drawing parallels between the UN's Human Bill

of Rights and Apple's terms of service. While computer scientists possess technical expertise, they also require sufficient knowledge to collaborate effectively with social scientists. For instance, historians and economists can provide valuable insights into how AI might reshape work and society. It's crucial to acknowledge both periods of rapid progress and stagnation in technological advancements. If computing resources are primarily allocated to AI, other computing sub-disciplines will need to adjust accordingly.

AI Adoption Bottlenecks and Solutions

Wierman was the Director of the Information Science & Technology initiative at Cal Tech, where he worked to enable ideas from computing and information sciences (or “AI” currently) to transform and create disciplines. This research approach impacts education, as it relates to applying AI in research and careers, and the ability to lead new research initiatives like "AI for science." Recently, he investigated educational obstacles at Caltech that hindered AI adoption among students. Interviews with 40% of the faculty revealed two main bottlenecks: 1) Formal AI training was either too late or nonexistent, and 2) Research often resulted in "gradware" AI tools, limiting their impact beyond campus.

To address the first bottleneck, he developed AI bootcamps: week-long, personalized, hands-on sessions for 20 students, offered during breaks. These bootcamps progressed from AI basics to targeted disciplinary tools and applications, in partnership with industry, to reach every graduate student. For the second bottleneck, he established a Software Engineering Academy. This program provided pre-doctoral students with software engineering mentors and integrated them into campus labs to cultivate an industry-grade software engineering culture. Labs hosting a pre-doctoral student are committed to providing full training and bootcamps for their entire lab.

All students need Research Methods Training

Jensen has been a researcher in machine learning and taught research methods to PhD and master’s students at UMass for over 15 years. He shared a variety of big research questions and challenges in relation to LLMs around areas like behaviors, mechanisms, learning, architectures, and scale. In order to address these questions and challenges, we need researchers who are able to “discover new principles, formulate and investigate novel research questions, identify and correct errors from existing authorities” amongst other traits. Unfortunately, these outcomes are not typical for undergraduate students who are more tasked with adoption, reproduction, and correctness. This shapes their experiences of graduate school and perpetuates a lack of formal training in methods. As a solution, he teaches a methods course annually that covers philosophy of science, research selection of planning, and

empirical research design. As a part of the course, students conduct independent research projects and analyze technical papers.

Scaling for Large Computing Programs

Wu is from a large college of computing with extensive experience in leading centers focused on “applied” AI research. Kennesaw State University (KSU) computing programs serve 5,700 students with 6 BS, 6 MS, and 2 PhD programs, most recently adding an MSAI (F ‘24) and a BSAI (F ‘26). In these programs, they have been working to develop interdisciplinary curricula that will support students with computing and non-computing backgrounds. He discussed several initiatives that share resources, including the US NAIRR Pilot, AI4EU, New Generation AI Development Plan (China), and AI Singapore. This list highlights the international investment in AI. He also highlighted the Computer Science Curriculum 2023 (ACM/IEEE-CS/AAAI) and CAE-CD CyberAI that focus on undergraduate students. He sees a gap in the curricular community for AI graduate education in terms of guidelines and resource sharing. These kinds of advances will be essential for programmatic robustness at large universities.

Change is Here to Stay

Da Silva worked in industry from 2000-2014 and as part of her role interviewed PhD job candidates at IBM Research and Qualcomm Research. Between 2014-2019 she spent time in academia and saw a growth in student population by 43%, and faculty population by 60% which has brought about cultural change. Currently, as a rotator at CISE she is thinking about graduate education in relation to the new Computing Education Research program. Her experience with past changes now presents an opportunity to refine best practices. Incoming students will be those who grew up using new tools in undergraduate and high school. The practice of coding is undergoing transformations, just as it has throughout its history. This moment might prompt us to consider how LLMs are used to summarize books before we engage in writing a book. Perhaps this will foster greater curiosity, encouraging deeper engagement with texts through interactive conversation with people and AI, and by freeing up time from tasks now handled by AI.

Question and Answer

- Graduate education is more than PhD. Can you speak to master’s students as well?
    - A: I do think that what we have been discussing about AI in graduate education may apply more to Master's that are course-based only. If we're talking about the Master’s who have projects that many have, then it may be that having the potential to use tools to generate code quicker and do prototypes quicker.

    - A: Interdisciplinary work is prescient for the students. Research is something we should be thinking about pushing into master's programs. They're operating in a rapidly changing environment and need to be comfortable with change.
    - A: Students need to gain experience in research methods; while working at companies want to be able to engage w/ what these people do.
- Is the master of AI program valuable if it does not have math in the curriculum?
    - A: Without some math fundamentals, you'll be more of an AI user.
    - A: We still have to learn what type and level of math is a prerequisite for graduate AI education.
- Thinking about the use of GenAI to write papers, as educators, how can we assess the success of graduate education? What is the quantity and role of publications?
    - A: The role of the work of evaluative committees is to read and assess the depth of the work. We will need to move beyond using external cues and to value the expertise of students' committees.

Resources:

- Reframing AI Alignment in terms of Censorship and Privacy Mohamed Ahmed (Citizen Lab, University of Toronto) [Best Talk Award at Hot Topics in Privacy Enhancing Technologies 2025] - https://www.cs.toronto.edu/~mohamed/slides/hotpets2025.pdf
- Computer Science 2023 - https://csed.acm.org/
- CAE-CD - Download Can be found under the Cyber Defense's Learn More Button https://www.nsa.gov/Academics/Centers-of-Academic-Excellence/
- NAIRR - nairr.org
- AI4EU - https://www.AI4europe.eu/about-AI4eu
- New Generation AI development Plan - https://digichina.stanford.edu/work/full-translation-chinas-new-generation-artificial-intelligence-development-plan-2017/
- AI Singapore - https://AIsingapore.org/

AI and Workforce: Jobs to Die, Survive, and Thrive

Moderator:

- **Dejan Milojicic,** Distinguished Technologist, Hewlett Packard Labs

Speaker:

- **Jeff Egan,** Principal Recruiter, Amazon
- **Nita Patel,** Vice President, Engineering, Otis Elevator Co.
- **Ray Perrault,** Distinguished Computer Scientist, SRI International
- **Ben Zorn,** Partner Researcher, Microsoft

Key Takeaways

- **More Jobs, Different Skills:** The panel's consensus is that AI creates more roles than it eliminates. The challenge isn't a scarcity of jobs, but a fundamental shift towards new skills.
- **From Coder to AI Supervisor:** The most valuable skills to these industry professionals are no longer just coding, but critical thinking, problem-solving, and systems-level thinking. The future of work involves supervising AI agents. The biggest job growth is expected in IT, requiring more people to build and manage AI.
- **CS Fundamentals Still Reign:** A traditional CS degree is more important than ever, but the focus must be on timeless fundamentals like algorithms, systems design, and logic, not just specific programming languages.
- **AI in Resumes not Interviews:** While recruiters advocated for the use of AI in resumes, they discouraged using AI in the behavioral interview process, cautioning it appears inauthentic and unprepared.

Summary:

A panel of experts from industry research, recruiting, and research institutes convened to discuss the profound impact of AI on the job market, the skills required for the future workforce, and the evolving relationship between academia and industry. The discussion addressed the practical, human-centered questions of career paths, educational needs, and societal adaptation in an era where AI is reshaping every professional discipline.

The core of the conversation focused on a cautiously optimistic outlook. While acknowledging the disruptive nature of AI and the inevitability of job displacement in certain areas, the panelists collectively agreed that AI is creating more roles than it is eliminating. The challenge,

they contended, lies not in a scarcity of jobs but in a fundamental shift in the required skills. The panelists suggest that the future workforce will move from performing automatable tasks to supervising AI, thinking critically about systems, and tackling the complex ethical and societal questions this powerful new technology introduces and exacerbates.

The panel began by dissecting the current state of the AI-driven job market, blending quantitative data with qualitative observations. Panel members opened by framing AI as a technology that will likely reshape every discipline, demanding agility and a renewed focus on systems-level thinking. Panelists stressed that roles in security, reliability, and privacy are becoming more critical than ever, as making AI "do the right thing" is a central challenge.

A Shifting Job Market

The panel addressed the common fear of AI-driven job loss by presenting a more nuanced picture. Egan referred to research suggesting that AI is currently creating more jobs than it is eliminating, primarily by automating mundane and repetitive tasks. This view was substantiated by Perrault, who shared data from various reports, including the AI Index. These reports indicated that while demand for skills in generative AI has quadrupled, the overall effect on graduate jobs is complex. A recent downturn in hiring for new graduates may be a market overcorrection from a pandemic-era hiring boom rather than a direct result of AI displacement. Interestingly, Perrault recalled, the data showed that AI tools often provide a larger productivity boost to inexperienced workers than to experts. One panelist cited a study from METR, “Measuring the Impact of Early-2025 AI on Experienced Open-Source Developer Productivity”, which found that for highly experienced programmers on large projects, AI assistant tools could decrease productivity. Panelists made sure to highlight that the impact of AI is not uniform and that its role is more of a collaborator and skill-leveler than a replacement. Panelists expected that the most significant job growth from AI will be in IT and software engineering, suggesting that the industry will require more, not fewer, skilled professionals to build and manage these new systems.

The New Skillset: From Coder to Systems Thinker

A central theme was the evolution of necessary skills. Zorn, with a background in programming languages, argued that AI is giving everyone the automation "superpowers" that software developers have long enjoyed. Just as compilers freed programmers from manual compilation, AI is freeing professionals in every field from tedious work, allowing them to focus on higher-level strategy and creativity. He predicted that the nature of many jobs will shift from performing tasks to supervising AI agents.

This requires a move away from thinking of computer science as just "writing code." The panel emphasized that the most valuable skills for the future are critical thinking, problem-solving, and the ability to design, build, and reason about complex systems safely and reliably. The workforce of the future will need to grapple with the deep ethical and societal implications of AI, turning these challenges into new areas of expertise.

Bridging Academia and Industry

To prepare students for the new reality AI is creating, the panel stressed the need for closer collaboration between universities and corporations. While industry drives the creation of large-scale models, academia plays a crucial role in conducting fundamental research, exploring the limitations of current systems, and novel ways of applying AI in domains like biology and the sciences.

Panelists called for industry to "lift the curtain" and provide greater transparency regarding the training data and internal mechanics of their models, as this would significantly aid academic research and improve the ethical integrity of AI systems. They also suggested that industry can help by sharing real-world problems, such as novel security vulnerabilities or user interaction challenges, which can inform university curricula and research projects. For students, their advice to be competitive on the job market was to focus on strengthening fundamental computer science principles while also getting hands-on experience with new AI tools.

Question and Answer

- **What new types of jobs should students prepare for?**
    - The panel agreed that the exact job titles of the future are unknown, just as they were in past technological revolutions. However, the *nature* of the work will involve more supervision of AI, reviewing AI-generated content, and collaborating with AI agents. Skills in critical thinking, problem-solving, and building robust secure systems will be transferable across many new roles.
- **If entry-level jobs are eliminated, where will mid-level staff come from in 10 years?**
    - The panel challenged the premise that all entry-level jobs are being eliminated. The recruiter noted that the pipeline will still come from universities, but graduates will enter the workforce with a different, more AI-centric skillset. The path will look different, but it will still exist.
- **Is a traditional Computer Science degree becoming obsolete?**
    - No. The panel was unanimous that a CS degree is more important than ever. However, the focus must be on the fundamentals of computing, algorithmic

thinking, and systems design, not just on coding. These core principles are essential for building and managing complex future AI systems.

- **How is AI affecting the interview process?**
    - A recruiter on the panel noted that AI tools have helped candidates create better resumes. However, he issued a strong warning against using AI assistance *during* a live interview, as it is often apparent to the interviewer and comes across as inauthentic. He referenced an interview he facilitated where a student pulled out their phone to ask an LLM for answers to a behavioral question. He advised students to practice their direct communication skills, as these have seen a decline in the age of AI-mediated interaction.

Resources:

- **The AI Index Report:** An annual report that tracks data and trends in artificial intelligence. https://hai.stanford.edu/ai-index
- **Gartner Research:** A technology research and consulting firm that analyzes market trends. https://www.gartner.com/en
- **Financial Times Article:** An article titled "Is AI killing graduate jobs?" https://www.ft.com/content/99b6acb7-a079-4f57-a7bd-8317c1fbb728
- **Rodney Brooks's Webpage:** Recommended for realistic insights into the challenges of robotics. rodneybrooks.com
- **Study from METR:** Measuring the Impact of Early-2025 AI on Experienced Open-Source Developer Productivity - METR

Building Professional Master's Programs in Computing and AI: Strategic Directions for Academic Leaders

Moderator:

- **Nancy Amato,** Director, Siebel School of Computing and Data Science; Abel Bliss Professor of Engineering, University of Illinois Urbana-Champaign

Panelists:

- **Magda Balazinska,** Professor, Bill & Melinda Gates Chair, and Director of the Paul G. Allen School of Computer Science & Engineering, University of Washington
- **Barrett Bryant,** Dean of the College of Science and Engineering and Professor of Computer Science, Texas State University

- **Amar Raheja,** Associate Chair and Professor, Computer Science, California State Polytechnic University, Pomona
- **Vivek Sarkar,** John P. Imlay, Jr. Dean of the College of Computing, Georgia Institute of Technology

Key Takeaways

- **Master's Programs are a Strategic Priority:** A live poll showed that the vast majority of institutions have started or are considering new master's programs in computing (over 66%) and AI (over 43%), viewing them as essential for both workforce development (80%) and as a source of revenue (77%).
- **A "One-Size-Fits-All" Model Fails:** The panelists showcased a wide variety of successful models, ranging across stackable certificates, massively scaled online master's programs (OMS), and interdisciplinary "CS+X" degrees.
- **Online Programs Offer Massive Scale and Access:** OMS programs can scale to accept thousands of students per year for under $10,000. This demonstrates the immense demand for affordable and flexible advanced education. An innovative "prove yourself in the first two courses" admission model at Georgia Tech has opened doors for many non-traditional students.
- **Professional Students Require Different Support:** Panelists agreed that professional master's students have different goals and needs than traditional research-track students. Success requires a robust infrastructure for career services, industry networking, and flexible academic advising.
- **Faculty Buy-in and Campus Politics are Key Hurdles:** Creating new interdisciplinary programs often faces challenges from campus politics and faculty resistance. The panel highlighted revenue sharing and creating approved, replicable degree templates (like UIUC's CS+X model) as effective strategies to gain buy-in and streamline approvals.

Summary:

This panel, moderated by Amato of the University of Illinois Urbana-Champaign, convened five academic leaders from a range of institutions to discuss the strategic directions for designing, scaling, and managing professional master's programs in computing and AI. The discussion provided practical insights into different program models (on-campus, online, stackable certificates, CS+X), strategies for supporting varied student populations, and the challenges of balancing academic rigor with workforce demands and revenue generation. The panelists' consensus was that these programs are a vital and growing component of the computing education ecosystem. Amato began with a series of Slido polls that demonstrated consensus

among attendees that master's programs are a critical and growing part of their departments' strategies. The panelists then shared the diverse models employed at their institutions.

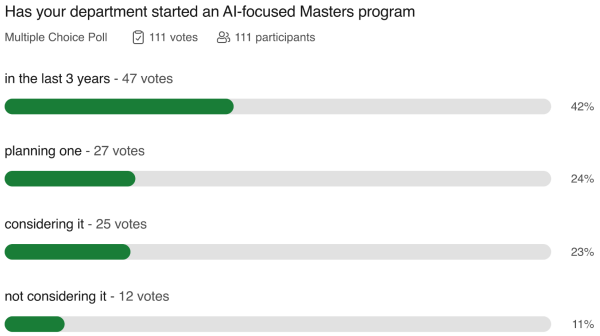

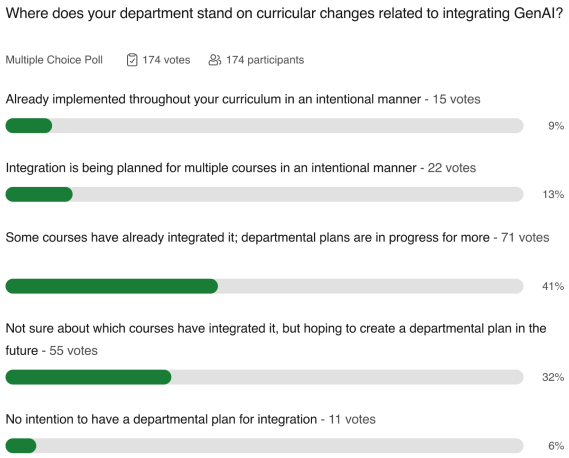

Balazinska described the University of Washington's long-running professional master's program (since 1996) and its modern "stackable" certificate model. This allows different departments (e.g., Computer Science, Human Centered Design & Engineering) to offer their own AI-related certificates, which students can then combine into a master's degree, offering maximum flexibility. Bryant explained that Texas State University, an R2 institution aiming for R1 status, is launching a dedicated MS in AI to build faculty expertise and serve as a pathway to its PhD program. He emphasized AI as a cross-cutting university focus, with collaborations in areas as diverse as math and digital humanities.

Representing the Cal State system's "learn by doing" philosophy, Raheja of Cal Poly Pomona noted that their traditional master's program does not generate revenue due to state funding models. In response, they are developing a new online master's program specifically to create a more stable funding source. Finally, Sarkar detailed Georgia Tech's massively successful and affordable online master's programs (OMS), which have enrolled tens of thousands of students. He highlighted their unique admissions model, where students without traditional prerequisites can prove their readiness by succeeding in the first two courses. He also noted the strength of the online student community, whose members often become engaged alumni and TAs.

The conversation then shifted to key strategic questions. On the topic of AI degrees versus concentrations, Balazinska argued that certificates offer flexibility for students from any discipline, while Bryant countered that a dedicated MS in AI degree helped his university build a critical mass of faculty expertise and created a strong signal that attracted new students.

When discussing the differences between professional and research master's students, Raheja emphasized their distinct goals, backgrounds, and support needs, noting professional students value industry networking over academic conferences. Sarkar added that while Georgia Tech

strives to create a unified community, they must scale up advising and career services to meet the specific needs of their large online and professional cohorts.

Question and Answer

How do you get a proposal for a "CS+X" or stackable program approved?

- **Balazinska:** Advocated for a bottom-up approach. You can start by creating individual certificates, which are often easier to get approved, and then design a mechanism to stack them into a full degree. A key challenge is ensuring that a combination of certificates includes a shared, culminating experience.
- **Amato:** Described UIUC's procedural model. They worked with the provost to create an approved "CS+X" degree template. The first one was the most challenging to get through, but now that the model is established, launching new CS+X degrees is an expedited and much easier process.

How do university budget models (e.g., Responsibility Center Management vs. centralized) affect this?

- **Sarkar:** Argued that faculty buy-in is the most critical factor, independent of the specific budget model. A new program cannot be a top-down mandate; it has to be socialized and supported by the faculty.
- **Amato** Agreed, emphasizing the need to make a compelling argument to faculty. To get other departments on board for interdisciplinary programs, **revenue sharing** can be a powerful tool to generate enthusiasm and show that the new program is not a zero-sum game.

How are you using AI/Automation to support online students?

- **Sarkar:** Stated that this is a huge opportunity, building on Georgia Tech's long history with its AI institute focused on education. They are exploring how to use AI for advising and tutoring, but he noted that one of the biggest challenges, for both online and in-person courses, is how to effectively deal with academic integrity.

What are best practices for designing MS programs that provide strong AI skills for non-CS students (biologists, social scientists, etc.)?

- **Panel Consensus:** It's essential to be clear about which student population a program is designed for. The panelists described three potential levels: 1) programs for students with strong tech undergrad degrees (computer science, math), 2) programs for general STEM students who need applied skills but not deep theory, and 3) programs for students in fields like social science and business who want to study the societal implications and impact of

AI. This is not about "diluting" a CS program, but about creating distinct offerings for different audiences.

How do you support pathways from community colleges into master's programs?

- **Panel Consensus:** All panelists confirmed that well-established "2+2" pathways exist for undergraduate students. They noted that creating similar direct pathways for master's programs is an active area of discussion at their institutions, with a goal of making advanced degrees more accessible.

How are you handling changes in international student enrollment due to visa issues?

- **Raheja:** Has seen a drop in undergraduate applications from international students, but graduate applications have remained consistent.
- **Sarkar):** Said that while they were nervous, most students have successfully obtained visas this year. He stressed that online programs are a critical way for students to earn a degree from their institution and circumvent visa issues entirely.
- **Bryant:** Mentioned that some students start their program online and then transition to in-person study once they secure a visa.

How do you evaluate the readiness of master's applicants from varied international backgrounds?

- **Sarkar:** Uses task forces of 5-6 people to review academic records. A triage process is necessary to manage the volume, with GPA and other academic metrics used as an initial filter.
- **Bryant:** Utilizes a preliminary placement exam and a bridge program with prerequisites to ensure students are prepared.

Do Master's students take separate courses from undergraduate and PhD students?

- **Balazinska:** Professional master's and non-CS master's students take separate classes designed specifically for their needs. Traditional, daytime research master's students may co-enroll with PhD students.
- **Sarkar:** Their students enroll in the same classes as PhD students and advanced undergraduates. He sees this as a benefit, as the master’s students (who have stronger gender diversity in comparison to PhD students) contribute to a more diverse classroom.
- **Bryant:** Views co-enrollment as a valuable recruiting tool to encourage BS students to pursue an MS, and MS students to consider a PhD.

The discussion highlighted the need for these programs to be flexible and responsive to the modern workforce, particularly in the age of AI. Offering diverse degree options benefits both domestic and international students, and there is a significant opportunity for interdisciplinary collaboration. For instance, in AI education, computer science can provide the core foundational knowledge while other departments teach its practical applications. Looking forward, the main challenge for academic leaders is to create sustainable models that balance innovation, academic quality, and the needs of a global student body. The key directions for the future involve:

- Leveraging AI in Education Delivery: Using AI and automation to support large-scale online programs through advising, tutoring, and authentic assessment.
- Building Sustainable Industry Partnerships: Developing scalable models to integrate industry engagement, such as internships, capstone projects, and guest lectures.
- Enhancing Pathways and Accessibility: Creating more robust and accessible routes into master's programs for students from non-traditional academic backgrounds, including community colleges and other disciplines.

Caught Between Moonshots and Manhattan Projects: How Might Computing Researchers Continue to Innovate in AI?

Organizers and Moderators:

- **Raj Rajaraman**, Associate Dean of Faculty Affairs and Professor, Khoury College of Computer Sciences, Northeastern University
- **Bill Regli**, Professor of Computer Science, University of Maryland

Speakers:

- **James Allan**, Professor and Associate Dean for Research and Engagement, Manning College of Information and Computer Sciences at UMass Amherst
- **Nadya Bliss**, Executive Director, Global Security Initiative; Professor of the Practice, Arizona State University
- **Zachary Ives**, Department Chair and Adani President’s Distinguished Professor of Computer and Information Science, University of Pennsylvania

Key Takeaways

- In order to achieve projects at scales never seen before in computing, collaboration will be required.

- In order to achieve projects at scales never seen before in computing, collaboration will be required, especially across disciplines and nations.
- While industry is focused on scalability, there still exist many interesting problems to solve that will address fundamental problems in AI such as: correctness, safety, equilibria, incentives, and parametric learning in context.
- Moonshot funding initiatives need to be focused on long-term research and partnerships.

Summary:

This panel focused on exploring grand challenges that could be specified within the field of computing in order to push us to think bigger about our work. Historic examples have included the ENIAC, Project Oppolo, and the Vera Rubin Observatory. Many of these spanned disciplines. Historically, computing's grand challenges were on a smaller scale than these projects. In order to have a grand challenge it needs to be impactful, ambitious, timely, interdisciplinary and measurable. Academia is at a moment asking themselves what should we do to meet this moment?

Does Academia Need to Distinguish Itself in AI by Choosing Distinct Problems and Topics?

Bliss asked the audience to reflect on the incentive structures in academia and how they're fundamentally different. As academics we shouldn't put a negative spin on industry structures, often financial, just because they're different. These differences might instead be a catalyst for us to think about how companies or different agencies are incentivized to fund our work. An area of interest might be actually understanding why GenAI models perform how they do. Training in formal methods as academics makes universities the right space to do the work, and doing this work alongside industry should make us stronger.

Allan reflected on how at the time Google rose to fame, all researchers in his area were jealous of Google employees because they got to see a lot of proprietary data. However, he has learned that there are different things to do. There are a lot of smaller problems that you can work on even if you don't have access to the latest model. For example, small models can be used to work on interesting and impactful questions like "How do they learn? Where do they learn, Why do they learn?"

Ives wants us to think about the hope and talent that exists in our institutions. While industry is focused on scalability, academia might focus on things like: correctness, safety, learning, equilibrium, incentives, and parametric learning in context. This is an opportunity to do a lot of good and principled work.

Are there moonshots or Manhattan projects for the AI and Computing community that are not dominated by industry?

Reflecting on what led to previous grand challenges, Ives brings up the ideas of national fear and incentives. He also mentioned that there were small but impactful items, tangential to getting to the moon, that the space program made way for like velcro and tang. He argues that we should move beyond our ingrained assumptions and consider creating tools for single, specific purposes. Ives also emphasized the continued importance of unique human expertise, which could be significantly enhanced by AI. As educators, computer science faculty might think about scaffolding the use of AI to develop metacognition skills and help people think deeply. Another area of impact might be tools that enable collaboration across fields.

Allan shared that these moonshots are still to be discovered and that they will likely need billions of dollars and thousands of researchers to achieve. He sees the opportunity to create dedicated GenAI labs where these technologies are deeply tested before being integrated into existing approaches. For example, an AI and healthcare lab where people could test out these tools in a controlled setting before using tools with patients. Allan also believes that there is a lot of work to be done in value alignment and thinking about whose values are embedded in AI systems and how that affects stakeholder adoption.

Bliss prefers being a mission-focused person, especially in relation to projects that impact people, like national security. There is a National Security Commission on emerging biotechnology that highlights how developments in biotechnology are AI-focused. She describes recent work from the National Academies Standing Committee on Artificial Intelligence Technologies and Emerging Impacts that discusses AI, but does not see computer science at the center. She emphasizes the need for us to be leading in partnership with all these disciplines.

What can we identify as a community to carry out large-scale AI research? How will departments manage AI hiring and develop the next generation of research leaders?

In order to find visions of what is possible, Allan suggests academics engage in funding endeavors that integrate social scientists and humanists, as well as creating pathways to international partnerships. This might look like building a global AI moonshots lab.

Bliss currently runs a research unit, rather than an academic department, that has 50 faculty researchers. When creating these types of labs, she highlights how not everyone will be the pinnacle of their field like LeBron James or Misty Copeland, and not everyone will aspire to have industry salaries. This is not necessarily a time to overcorrect to keep up with industry salaries. As a part of hiring, she suggests that academic institutions collaborate more with industry and national labs in order to create pockets of money and advance research in this area.

Ives emphasized that some of these funding initiatives will need to be focused on long-term research. Right now, we can think about working with industry, but many of those agreements would take over a year to scale up within academic institutions. So institutions will need to think about risk and benefit in relation to streamlining these kinds of partnerships.

Question and Answer

- How much is collaboration vs colonization and what is unique about computing being the lead vs supporting other fields?
    - A: Having worked in a number of projects w/ folks outside of the CS field, it's important to understand what value you bring. Sometimes CS are plumbers, hooking pipes to things and enabling water to flow. Right now we're bringing tools rather than solutions.
    - A: Collaborations only work if you're doing things that are communicated well. We cannot bring social scientists in at the end and do checkmarks. Anyone who is a computer scientist and is AI centric should consider how many disciplines are trying to integrate AI. CS is uniquely proposed to talk about what we know as well as what the limits are.
- One of the threats we've been visiting is this "AI taking all our jobs and not having a place". What if we make that the moonshot so we can do other things?
    - A: While AI will likely automate many tasks, it won't free up all your time for leisure. Therefore, we need social scientists to analyze the economic implications of such technological advancements.
    - A: Our vision is to amplify the mind, not replace the mind. Fundamentally there are plenty of things we need humans for.
    - A: There are infinite amounts of interesting problems for us to solve.
- Are we victims of the street lamp effect? Are we missing things by being biased by hype?
    - A: The current excitement around AI has undeniably diverted resources from other areas. Many researchers are exploring various applications of AI, from direct AI development to applying algorithms within AI contexts. We are likely to continue focusing on AI and its successors, given their significant potential impact. While this intense focus might not be inherently detrimental, it does compel us to re-evaluate and re-present our work in light of new technological trends.
    - A: The number of people who've asked me about AI is surprising - from the optometrist, dentist, cab driver, general practitioner doctor, to high school students. Now is a

moment that we have a lot of public interest. We have a responsibility to ensure against vulnerabilities.

- Is there any truly open AI research? This might be a great angle for academia.
    - A: Topics might look like: planning with systems that use LLMs; tool use; explainability & correctness; reasoning about how different agents & incentives are going to act. Previously, significant progress in these areas was elusive, but now there are opportunities to advance.
- Publish & perish was bad - How do we move to more quality & less quantity?
    - A: Some people advocate that we look only at the top 3 papers for every candidate. There's an arms race in numbers of publications and we are losing sight of understanding the contributions of potential hires. We should incentivize a focus on what's important in a potential hire rather than quantity of publications.
    - A: There's a CRA best practices memo on quality and not quantity; as a community, we need to move beyond paper into implementation in departmental hiring etc.

Resources:

- https://cra.org/cra-best-practices-memo-on-evaluating-scholarship-in-hiring-tenure-and-promotion/
- https://www.biotech.senate.gov/

Lightning Talks

Throughout the 2025 CRA Summit, several of the main panel discussions were preceded by concise and impactful lightning talks. Lasting between two and five minutes, each of these presentations was designed to introduce a novel institutional model, a provocative strategic question, or a practical framework for consideration. They served as valuable catalysts, seeding the subsequent panel discussions with concrete ideas and sparking deeper engagement from the audience. For ease of reference, the talks are compiled in this section in the order in which they were delivered at the 2025 CRA Summit. The first six were presented during the CRA Summit on AI Undergraduate Education, followed by three additional talks—beginning with Kathleen Fisher from DARPA—delivered during the CRA Summit on AI Graduate Education and Research.

AI Takes Root!

Speaker:

- **Ran Libeskind-Hadas**, Founding Chair, Kravis Department of Integrated Sciences, Claremont McKenna College

Summary

Libeskind-Hadas shared how Claremont McKenna College (CMC) is integrating AI into their curriculum. All students at CMC, no matter their department, take the general science elective "Codes of Life." For students in this course they integrate computational thinking with biological and socio-scientific questions, engage in python programming, and discuss "How AI works" while building mini LLMs called CHAT-CMCs. They also experience an AI agent as a lab group member and course tutor. Each science faculty member at CMC teaches this course which helps computing and AI concepts "take root" in all science courses. Still ongoing are conversations across the college on how to use AI throughout the curriculum.

AI Literacy for All – A Scalable Approach to Democratizing AI Education

Speaker:

- **Harini Ramaprasad,** Associate Dean for Undergraduate Programs & Student Success, University of North Carolina at Charlotte

Summary

Ramaprasad outlined a new approach to making artificial intelligence education universally accessible. She identified a common split in AI education: courses are often either deeply technical and siloed within computer science, or they are ethics-focused discussions that lack necessary technical grounding. To bridge this divide, UNC Charlotte has developed and implemented a course that treats AI literacy as the new digital literacy. The course is founded on a socio-technical approach built upon four core pillars:

1. **Technical Foundations:** Understanding the scope and technical dimensions of AI.
2. **AI Interaction Skill:** Learning to effectively and safely use generative AI.
3. **Socio-technical Principles:** Reviewing the core issues of socially responsible and ethical AI.
4. **Societal Impact:** Analyzing the broader effects of AI on work, culture, and public perception.

This curriculum is designed for maximum flexibility, with modules that can be adapted for various majors and levels, scaling from a brief six-unit introduction to a comprehensive 24-unit graduate course. A Spring 2024 pilot, offered with no prerequisites, successfully enrolled 55 students from diverse disciplines who demonstrated high performance in key AI literacy constructs. Ramaprasad concluded that by integrating technical knowledge with societal and ethical considerations, this model provides a comprehensive and scalable framework for empowering all students as informed and responsible users of AI.

Tadimalla, S. Y., and M. L. Maher. 2025. " AI literacy as a core component of AI education." *AI Magazine* 46: e70007. https://doi.org/10.1002/aaai.70007

The CRA Practitioner-to-Professor (P2P) Survey

Speaker:

- **Rahul Simha,** Professor, The George Washington University

Summary

Simha shared briefly on the Computing Research Association's Practitioner to Professor (CRA P2P) Survey. This was a National Science Foundation sponsored initiative to garner national-scale feedback from industry practitioners to academic computing departments. This initiative was supported by ABET, ACM, CSAB, and IEEE. This survey was given to industry professionals focused on what they might expect out of 4-year CS departments. The P2P Survey, the largest of its kind, received 1048 responses. In the future they hope to extend this work to study grad programs, other computing programs, and generate trend and comparison data.

Example Takeaways:

- 90 percent of professionals are in favor of requiring math. Discrete-math ranked third amongst math subjects.
- Data structures (including advanced) were the top sub-topic in Algorithms.
- In the area of architecture professionals prioritized foundations such as digital logic (boolean algebra (45.6%), arithmetic logic & data representation (32.2%), gates, combinational circuits, flip flops, and registers (31.0%)).

https://cra.org/industry/p2psurvey/

The View from NSF

Speaker:

- **Ellen Zegura,** Acting AD of Computer and Information Science and Engineering, National Science Foundation

Summary

Zegura highlighted the critical importance of federal funding for computer science research. As people are applying to the National Science Foundation, she reminds them to think about their audience as the priorities at NSF will be different than the priorities at DARPA. Currently we have a rare opportunity for change in areas of emphasis: AI Ecosystems, computing frontiers, computing in context, cybersecurity, resilience, and quantum.

They are working to:

- Simplify solicitations programs and processes
- Increase legibility for all stakeholders
- Organize for the next 20 years of computing research
- Create more agility
- Interact with the community differently.

When an LLM Joins the Team: Creative Collaboration and Critical Reflection in the Classroom

Speaker:

- **Orit Shaer,** Professor and Co-Chair of Computer Science, Wellesley College

Summary

Shaer outlined Wellesley College's structured approach to integrating Generative AI into their computer science curriculum, a framework they call the "123 of GenAI in CS." This scaffolded model is designed to build student competency and critical thinking over time:

- **100-level Courses:** At the introductory level, the use of GenAI tools is **not permitted**. The focus is on building foundational skills and fostering critical discussions about the technology's capabilities and societal impact.
- **200-level Courses:** In intermediate courses, students are allowed **limited use of GenAI** as a learning aid. The framework calls for separating practice from assessment, allowing

students to use the tools to explore concepts while ensuring their graded work reflects their own understanding.

- **300-level Courses:** At the advanced level, the goal is the **strategic and responsible use of AI**. Students are taught to use AI intentionally to spark ideas, check their understanding, and evaluate the tool's suggestions and limitations, with the core principle being that the *student* is always doing the real, critical thinking.

The presentation also introduced "Group-AI Ideation," a novel framework for creative collaboration developed at Wellesley. This structured process guides student teams through phases of divergence (using AI to brainstorm and expand possibilities) and convergence (using critical thinking frameworks like the Six Thinking Hats to evaluate, refine, and select the best ideas). This model treats the LLM as a distinct team member, teaching students how to leverage its strengths while mitigating its weaknesses in a collaborative project setting.

NSF AI Institutes: National Resources for AI Research and Education

Speaker:

- **Jamie Payton,** Dean, College of Computing, New Jersey Institute of Technology

Summary

Payton highlighted the wealth of educational resources being developed across the 29 National Science Foundation (NSF) AI Institutes. Speaking as a representative of the NSF AI Institute for Inclusive Intelligent Technologies for Education (INVITE), Payton explained the collective mission to support foundational and use-inspired AI research, and to grow the future AI workforce. Each institute is charged with creating educational components, with efforts spanning across K-12, community college, 4-year undergraduate, graduate, and workforce retraining programs. Payton emphasized that these resources are being made publicly available to advance AI education nationwide.

She provided two key examples of these efforts:

- **The AI Institute for Research on Trustworthy AI in Weather, Climate, and Coastal Oceanography (AI2ES)** has created a model for **stackable AI certificates**. This model establishes a clear educational pathway from high school to community college and on to four-year degrees, with established curriculum and articulation agreements now active in Texas. The institute encourages others to adopt this successful model and is sharing its resources and lessons learned.

- **The INVITE AI Institute, in partnership with the AI Institute for Engaged Learning**

(EngageAI), offers an **"AI Scholars" program**, a non-credit online experience for undergraduates. The program introduces students to AI fundamentals and research and engages them in creating AI learning experiences for K-12 students in partnership with K12 educators.

The primary takeaway from Payton's talk is the launch of a new **"Education for AI" initiative**, organized by the AI Institute Virtual Organization (AIVO) with sponsorship by Google.org. This effort will consolidate the resources from all 29 NSF AI Institutes into a single, accessible education portal, promoting best practices and making it easier for educators everywhere to take advantage of this collaborative work. She concluded with an invitation for attendees to participate in upcoming national AI educational workshops to learn more.

CloudBank: JupyterHubs in Higher Education

Speaker:

- **Sean Morris,** CloudBank: Outreach and Support

Summary

In his lightning talk, "CloudBank: JupyterHubs in Higher Education," Sean Morris (Outreach and Support) outlined a solution for integrating data and computing into university curricula.

The presentation focused on CloudBank's browser-based computing platform, which leverages JupyterHub infrastructure to provide accessible, scalable environments for any course. Morris highlighted how the service simplifies deployment and access for students and faculty.

Key capabilities discussed include:

- Simplified Access: Authentication is managed via a single sign-on system (CiLogon), removing barriers to entry.
- Versatile Environments: The platform supports interactive Jupyter Notebooks for R and Python as well as full development environments like VSCode, RStudio, and Java.

The talk positioned CloudBank as a robust solution for institutions looking to easily provide powerful, web-accessible computing resources.

The View from DARPA

Speaker:

- **Kathleen Fisher,** Director of the Information Innovation Office, Defense Advanced Research Projects Agency

Summary

Fisher began her talk by outlining the scale and focus of her office, the Information Innovation Office (I2O), which manages an annual budget of approximately $500 million. She highlighted several of its key research thrusts, including the critical challenge of detecting deepfakes in audio and video and a novel program focused on creating AI tools to help pure mathematicians accelerate their work. She then detailed DARPA's core philosophy, contrasting its mission with that of science-focused agencies like the NSF. DARPA, she explained, is singularly focused on **impact**. Its goal is to "transform revolutionary concepts and even seeming impossibilities into practical capabilities". The agency's operational model embraces a willingness to take high risks, prove that something is possible, and then move on, ensuring a constant flow of fresh ideas through planned turnover in personnel.

This entire process is guided by the **Heilmeier Catechism**, a foundational set of questions that instructs researchers to clearly define their objectives, articulate what is new about their approach, and specify what difference success will make. Fisher emphasized that the workflow for engaging with DARPA is very different from other agencies. The "secret sauce," she shared, is to **meet with a Program Manager** as they are trying to start a new program. The goal for a researcher should be to understand the PM's vision and find ways to help them shape their new program, rather than responding to a formal call. Once a solicitation is released, the timeline is very short, often only 45 days to submit a proposal. She noted that universities can serve as prime contractors or as subcontractors for other teams. To facilitate this process, DARPA holds a "Proposers' Day" event for each solicitation to provide clarity and encourage collaboration.

Highlighting the Critical Importance of Federal Funding for Computing Research

Speaker:

- **Moshe Vardi,** University Professor, Rice University

Summary

Vardi’s lightning talk echoed the earlier Fireside Chat with Panch in terms of the importance of informing the public of the computing work we do through compelling stories. There is a Communications of the ACM (CACM) Special Issue that was “designed to highlight stories of profound impact that come from federally-funded academic research.” He wants your help spreading the word about these articles that share our story of societal and economic impact.

These articles can be found here and are rolling out between now and December 2025: https://cacm.acm.org/article/special-issue-federal-funding-academic-research/

A sample of Topics:

- U.S. Department of Agriculture's National Institute of Food and Agriculture (NIFA) has funded five National Artificial Intelligence (AI) Research Institutes
- Impacts of U.S. Federal Research Funding
- Cryptography
- Natural Language Processing
- Collaborative Distance Working

MS-CS ∩ MS-AI?

Speaker:

- **Jeff Offutt,** Professor and Chair, University at Albany

Summary

Offutt delivered a concise lightning talk that posed a foundational question for academic leaders: What is the essential curriculum for a Master of Science in Artificial Intelligence (MS-AI), and how should it intersect with a traditional Master of Science in Computer Science (MS-CS)? He argued that the answer depends on a program's strategic focus, which can be centered on the applications of AI, the building of AI systems, or the theory behind them. He framed the challenge of designing new AI master's programs by presenting a Venn diagram illustrating the typical curricula of an MS-CS and an MS-AI. The MS-CS side included topics like computability, operating systems, databases, and networks. The MS-AI side included math, modeling, machine learning (ML), natural language processing (NLP), computer vision (CV), and ethics. He pointed out that a significant and essential intersection already exists, containing subjects like algorithms, core AI, software engineering, security, and robotics. The central point of his talk was that before departments can define the specific courses for an MS-AI, they must first answer a strategic question: What is the focus of the program? He proposed three potential directions:

1. Applications: Focusing on how to use existing AI tools to solve problems in various domains.

2. Building: Focusing on the software and systems engineering required to create new, robust AI systems. This direction, he argued, necessitates more core computer science content.

3. Theory: Focusing on the mathematical and statistical underpinnings of AI models. This direction calls for a curriculum heavier in math and modeling.

Next Steps

The CRA Summit on AI Undergraduate Education and the CRA Summit on AI Graduate Education and Research reflected the pulse of the computing research community, highlighting pressing concerns, unresolved debates, and future directions. The conclusions will inform some of the tracks at the **2026 CRA Summit** which will take place **July 21-23, 2026**, in Prior Lake, Minnesota.

Like the previous biennial agenda of the CRA Conference at Snowbird, the 2026 CRA Summit will feature multiple tracks addressing a broad range of issues in computing research, leadership, and workforce development — offering participants expanded opportunities for cross-sector collaboration and strategic discussion.