2026 CRA Summit Report

October 2026



CRA
Computing Research Association

Authors

James Motes, University of Illinois Urbana-Champaign

Madeline Endres, University of Massachusetts Amherst

Alan Liu, University of Maryland

Pratik Soni, University of Utah

Acknowledgments

CRA would like to thank the 2026 Summit organizers: **Raquel Hill** (Spelman College), **Keith Marzullo** (University of Massachusetts Amherst), **Manuel Pérez-Quiñones** (University of North Carolina at Charlotte), **Rachel Pottinger** (University of British Columbia), and **Jeff Walz** (Acadigm.ai)

CRA gratefully acknowledges the support of the generous sponsors of the 2026 CRA Summit: **Google**, **ACM**, **IEEE-CS**, **Center for Inclusive Computing**, **University of Wisconsin-Madison's College of Computing & Artificial Intelligence**, **University of Texas at Arlington's Department of Computer Science and Engineering**, **University of Washington's Paul G. Allen School of Computer Science & Engineering**, **University of North Carolina at Charlotte's College of Computing and Informatics**, **Drexel University's School of Computer and Information Sciences**, **Kennesaw State University's College of Computing and Software Engineering**, **University of Illinois Urbana-Champaign's School of Information Sciences**, **University of South Florida's Bellini College of Artificial Intelligence, Cybersecurity and Computing**, and **Virginia Commonwealth University's Department of Computer Science**.



This material is based on work supported by the U.S. National Science Foundation (NSF) under Award Nos. 2533723 and 2626796. Any opinions, findings, 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.

Suggested Citation

Motes, J., Endres, M., Liu, A., and Soni, P. (2026). 2026 CRA Summit Report. Washington, D.C.: Computing Research Association (CRA).
https://cra.org/wp-content/uploads/2026/10/2026-CRA-Summit-Report.pdf

Table of Contents

Executive Summary **4**
- AI in Computing Research and Education 4
- Research Strategy and the Future of Computing 5
- Students, Leadership, and Emerging Applications 6

Introduction **7**

Plenary Sessions and Keynotes **7**
- Entertainment & AI: Challenges and Opportunities 7
- Leading Together: CRA Updates, Impact, and Community Recognition 10
- The Transformation of Higher Ed in the Age of Disruption 13
- Advancing National Priorities and Seeding Future Computing Breakthroughs 16
- Wind Against Tide: Computing Research Policy in Turbulent Times 18
- Fireside Chat with Chris Malachowsky, NVIDIA Co-Founder & NVIDIA Fellow 22
- "Doing Computer Science": The Case for Ethics at the Core of Computing as a Discipline 25

Parallel Track Sessions **29**

Track 1: Impact of AI in Computing Research and Education **29**
- When AI Goes In, What Comes Out? Leadership Lessons for Reshaping Computing Programs in this AI Era 29
- AI and the Research Enterprise 33
- AI Leadership in the Academy 36
- What Should the CS Curriculum Be in the GenAI Era? 39

Track 2: Next Breakthroughs in Computing **43**
- Computing for Physical Intelligence/Robotics 43
- Privacy, Trust, and Regulatory Compliance: It's Not Just for Lawyers Anymore 46
- Cybersecurity for the Public Good 48
- Boundless Frontiers: Co-Creating Global Research Futures 51

Track 3: Novel Computing Directions **54**
- CS Enrollment Trends and the Reality of Jobs for CS Graduates 54
- Research in AI and the Future of Arts and Entertainment 57
- Yes, And: Computing Visioning Improv 60
- The Pivotal Role of Computing Research for the Future of Healthcare 61

Track 4: Computing and the Now **66**
- Navigating the "Mess": Strategic Resilience for Computing Research Organizations 66
- Industry Expectations: What Industry Wants from Computing Students 70
- Sustaining Innovation: Navigating Federal Funding in Computing Research 74
- We've Been Here Before: Innovations in Computing 77

Lightning Talks from Emerging Leaders **80**
- Lightning Talks from Emerging Leaders 80
- Lightning Talks from Emerging Leaders 86

Executive Summary

The Computing Research Association (CRA) hosted the **2026 CRA Summit**, bringing leaders from academia, industry, and government together to shape the future of computing research. Plenary talks addressed that overarching theme, with smaller breakout sessions organized into four tracks: Impact of AI in Computing Research and Education; Next Breakthroughs in Computing; Novel Computing Directions; and Computing and the Now. Discussions crossed those boundaries, connecting research opportunities with educational, institutional, and societal choices.

This summary brings together those discussions under three themes introduced by CRA Board Chair James Allan in the Summit opening: **AI in Computing Research and Education**; **Research Strategy and the Future of Computing**; and **Students, Leadership, and Emerging Applications**.

AI in Computing Research and Education

Artificial intelligence (AI) is changing how computing expertise is developed and exercised. The educational challenge is to prepare students to use increasingly capable tools without outsourcing the understanding students need to direct and evaluate their work. Computing foundations, critical thinking, and stronger evidence of learning were recurring priorities across academic and industry perspectives. Views differed sharply over how AI should reshape the curriculum, with some advocating for rapid integration or substantial redesign and others warning that changes driven by current tools and industry pressures could weaken foundational understanding and long-term preparation. The Summit connected these goals with support for instructors, choices about existing requirements, and the integration of ethics and security into technical education.

In research, AI can accelerate discovery and assist peer review. At the same time, easier production of papers is increasing pressure on human reviewers and established ways of recognizing scientific contributions. The challenge is to strengthen verification and accountability while using AI's capabilities productively. Discussions of authorship, credit, and academic rewards identified a growing need to value the significance of research more than its volume. Together, the education and research discussions placed human judgment at the center of how the community should develop expertise and establish the quality of work.

Key Takeaways

1. Computing programs should prepare graduates to formulate problems, understand systems, and evaluate AI-assisted work, combining tool fluency with the foundations, communication, and judgment needed for durable careers.

2. Assessment should establish what learners can understand, explain, and do, combining realistic AI-assisted projects with independent demonstrations of competence rather than treating a polished submission as sufficient evidence of learning.

3. Ethics and security belong within technical computing education, preparing students to

recognize harms, anticipate misuse and attacks, weigh tradeoffs, and act on ethical concerns rather than treating responsible computing as an optional specialization.

4. Integrating AI requires more than new content: limited credit hours force curricular tradeoffs, faculty need training, time, and incentives, and departments must coordinate curriculum changes and AI-use policies while accommodating different course needs.

5. AI can accelerate discovery and assist peer review, but growing research output is straining human evaluation, increasing the need for verification, clear attribution, author accountability, and academic rewards for significance rather than publication volume.

Research Strategy and the Future of Computing

The future of computing depends on sustaining exploration across the field, including work whose eventual applications are not yet clear. AI demonstrates how foundational research can enable new capabilities, while emerging applications raise new questions for computing research. The Summit emphasized opportunities both within and beyond AI, linking long-term intellectual ambition with the institutional choices that determine which ideas can be pursued and who can contribute.

Those institutional choices shape both the conditions for pursuing research and how researchers' contributions are recognized. Shared computing facilities, models, and data can widen participation when accompanied by technical expertise and sustained support. Uncertain awards, changing policies for funding, and barriers to researchers working across borders can interrupt discovery and training. Discussions of researcher evaluation called for hiring, tenure, and promotion decisions to emphasize the significance of contributions rather than publication volume. Partnerships among universities, industry, government, and professional communities offer ways to combine capabilities and address shared challenges. Their value lies in what participants can accomplish together, not simply in assembling larger teams.

Key Takeaways

1. AI builds on advances across computing and creates new research questions; future progress depends on foundational and interdisciplinary exploration both within and beyond AI, including work whose applications are not yet apparent.

2. Shared computing facilities, AI models, data, and technical support can broaden participation in resource-intensive research, reducing the need for each institution to build and maintain those capabilities independently.

3. Research capacity depends on more than funding totals: award predictability, institutional support, and policies governing grant administration, international collaboration, and researchers' ability to work across borders shape whether teams can sustain discovery and train future researchers.

4. Institutions and professional societies should reassess how researchers' contributions

are evaluated, emphasizing significance rather than publication volume in the evidence used for hiring, tenure, and promotion.

5. Sustained partnerships and professional communities enable institutions to combine expertise, develop shared research directions, and address challenges they cannot resolve independently, through genuine joint work rather than simply assembling larger teams.

Students, Leadership, and Emerging Applications

Computing's expanding role across society calls for a broader view of whom the field educates and where graduates can contribute. Declining applications at some institutions and tighter entry-level hiring have sharpened questions about computing careers, but experiences vary across programs and student populations. The Summit connected these pressures with opportunities for interdisciplinary education, new routes into computing, and careers beyond technology companies. The resulting priorities combine better evidence about educational and career outcomes with attention to students' support and practical learning opportunities.

Across healthcare, creative work, robotics, and cybersecurity, computing offers opportunities to extend human capabilities. A recurring challenge is turning technical advances into systems that fit the work people actually do: reducing clinical workload, supporting creative processes, or making protection usable by organizations with limited resources. These opportunities connect interdisciplinary research with responsibilities for safety, privacy, and consent. They also inform the leadership discussion as public confidence requires both honest explanations of computing's benefits and substantive responses to concerns about educational affordability, career preparation, and technology's consequences.

Key Takeaways

1. Computing careers extend well beyond technology companies, but uneven enrollment and hiring trends call for evidence about educational and career outcomes to guide interdisciplinary programs, experiential learning, and routes into or back into computing.

2. Broadening participation in computing depends not only on routes into computing education but also on affordable and accessible tools, human mentoring, and hands-on opportunities that help students with different backgrounds and abilities succeed and contribute.

3. Computing can expand clinical capacity, extend creativity, and improve protection from cyber threats, but technical performance alone is insufficient: useful systems must be designed around human expertise and real-world conditions, with safety, privacy, and consent built into their design.

4. Public confidence in computing and higher education requires demonstrable benefit, responsible conduct, and honest communication, including substantive responses to concerns about affordability, career preparation, and harm — not stronger promotion alone.

Introduction

The Computing Research Association (CRA) hosted the **2026 CRA Summit** on July 21–23, 2026, in Prior Lake, Minnesota. This invitation-only convening of senior leaders from across academia, industry, and government focused on shaping the future of computing research. Through high-level sessions, candid dialogue, and strategic networking, participants stepped back from day-to-day demands to explore what's next for the field. The full event agenda and available presentation slides are on the CRA Summit webpage.

The Future of Computing Research

The overarching theme of the 2026 CRA Summit was **The Future of Computing Research**, organized around four interconnected tracks:

- **Impact of AI in Computing Research and Education**
  How AI is reshaping research agendas, curricula, institutional priorities, and the computing workforce.
- **Next Breakthroughs in Computing**
  Emerging ideas, paradigms, and technologies that could define the next era of computing research.
- **Novel Computing Directions**
  New and unexpected domains where computing research is having transformative impact.
- **Computing and the Now**
  Grounding strategic discussions in current realities, including policy, funding environments, workforce dynamics, and societal context.

Together, these tracks were designed to foster cross-sector dialogue and candid discussion among leaders navigating rapid change.

Plenary Sessions and Keynotes

Entertainment & AI: Challenges and Opportunities

Tuesday, July 21, 2026 | 7:00 pm–8:30 pm CT (dinner program)

Speakers

- **Eric Iverson,** Chief Technology Officer, United Talent Agency
- **Ashok Seshadri,** Senior Vice President, Product – AI, Data Science, and Analytics, United Talent Agency

Key Takeaways

- **Storytelling helps people anticipate and shape the future:** Humans use stories to organize information, rehearse possible trouble, and imagine alternatives to the world as it currently exists. Entertainment and research therefore serve related purposes by helping society "look around corners."
- **AI is changing how entertainment is protected, produced, and monetized:** Its impact extends across cybersecurity, synthetic-media detection, rights and licensing, content analysis, production, localization, digital performers, immersive experiences, and audience analytics.
- **Authenticity, consent, and attention become more valuable as content becomes easier to create:** AI can accelerate and expand creative work, but it also blurs authorship, enables unauthorized uses of people and content, and intensifies competition for an audience's limited time.

Summary

During this keynote, Eric Iverson and Ashok Seshadri of United Talent Agency explored AI's impact on entertainment through the broader role that storytelling plays in human thought. Iverson framed stories as tools for remembering information, anticipating trouble, and imagining futures that researchers and innovators can then work to make possible or prevent. Building on this framing, Iverson and Seshadri described how AI is transforming the entertainment ecosystem, from protecting artists and content to creating, localizing, distributing, and monetizing new forms of media. The session ultimately returned to the human elements that become more important as AI increases the volume and speed of production: originality, consent, authenticity, quality, and the ability to tell a story worthy of an audience's attention.

Storytelling as a Way to Look Around Corners

Iverson presented storytelling as a fundamental human behavior rather than an activity limited to writers, filmmakers, or performers. Stories connect otherwise isolated facts, make information easier to remember, and help people interpret incomplete events by constructing coherent narratives around them (often the source of rumors). In this sense, everyone participates in storytelling, whether through entertainment, research, business, sports, or everyday communication.

Stories are especially useful because they allow people to encounter trouble before it occurs. Iverson described science fiction as a way of imagining threats and possibilities at scale, temporarily setting aside the constraints of the present so that people can consider what might come next. Ideas presented in fiction have inspired technologies, security initiatives, and new directions for research, while fictional crises give audiences a way to practice responding to futures they may never actually face. The connection to the research community was direct: entertainment helps society imagine and recognize future trouble, while researchers develop the knowledge and technologies needed to address it.

AI Across the Entertainment Ecosystem

Iverson and Seshadri emphasized that AI's effects on entertainment reach well beyond generating images, scripts, or videos. AI is becoming part of how the industry protects content and verifies authenticity, particularly as cyberattacks and convincing deepfakes become easier to produce at scale. At the same time, name, image, likeness, and voice rights create opportunities for performers and estates to license identities for advertisements, games, chatbots, and other experiences. Similar agreements could allow rights holders to provide content for model training under negotiated terms rather than surrendering control over how their work is used.

AI is also changing how content is understood and produced. Multimodal systems can analyze scripts, contracts, audio, and video to identify contractual requirements, track brand obligations, and detect compliance issues before production reaches final editing. Generative tools can shorten production timelines, support rapid experimentation, and localize media by adapting both dialogue and the visual performance of the speaker. Digital performers, synthetic characters, short-form media, and immersive environments further expand the kinds of experiences that can be created. These capabilities connect with audience and business data, allowing agencies such as UTA to identify relationships among talent, brands, platforms, and markets across an increasingly complex entertainment landscape.

Human Value in an Age of Abundant Content

The growing role of AI raises a difficult distinction between work that is AI-assisted and work that has effectively been outsourced to AI. Previsualization, scene iteration, translation, and other tools can help human creators explore ideas more quickly and respond to cultural moments while they remain relevant. However, the same capabilities complicate copyright, authorship, attribution, and consent. Synthetic media can reproduce a performer's appearance or voice without authorization, while models trained on creative work may generate related material without compensating or even identifying the original creators. Technical guardrails alone do not resolve who owns a creative contribution or what degree of human involvement is necessary for a work to be considered original.

AI also creates the possibility of far more content than audiences can reasonably consume. Iverson and Seshadri did not suggest that efficiency alone would make creative work valuable. Instead, greater abundance is likely to raise the bar for experiences that justify an audience's finite time. Entertainment is moving simultaneously toward highly immersive formats that ask for deeper investment and shorter formats that minimize the time required from the viewer. Across both directions, compelling storytellers will remain differentiated by their ability to make people care. AI can increase the speed and scale of production, but human judgment, originality, and an understanding of the audience continue to determine which stories rise above the surrounding volume.

Leading Together: CRA Updates, Impact, and Community Recognition

Wednesday, July 22, 2026 | 8:30–9:30 am CT | Presentation slides

Speakers

- **James Allan**, Chair, CRA Board of Directors; Distinguished Professor and Associate Dean for Research and Engagement, Manning College of Information and Computer Sciences, University of Massachusetts Amherst; Director, Center for Intelligent Information Retrieval.
- **Tracy Camp**, Executive Director and CEO, Computing Research Association (CRA)

Key Takeaways

- Computing research advances through collective effort: CRA brings academia, industry, and government together to address challenges that individual institutions cannot resolve alone. Its work depends on community participation, shared expertise, and relationships that extend beyond the Summit.
- Sustaining the field requires investment in people, evidence, and educational resources: CRA's surveys, research training, mentoring, and AI education initiatives help institutions support researchers and adapt to changes in computing.
- Computing's value needs to be communicated beyond the research community: CRA leaders emphasized the importance of communicating with policymakers and the public about computing research and education. They described this as a shared responsibility of CRA, researchers, and their institutions.

Summary

James Allan of the University of Massachusetts Amherst, Chair of CRA's Board of Directors, opened the session before introducing Tracy Camp, CRA's Executive Director and CEO. Camp reviewed CRA's activities across research, education, workforce development, and advocacy, emphasizing the community contributions that make these efforts possible. Her presentation connected opportunities to participate in CRA with the broader goal of strengthening the conditions that support computing research. The session also recognized contributions through CRA's awards program and concluded with audience questions about funding uncertainty, public understanding of computing, and how CRA and its members could respond.

Advancing Computing Research Together

Camp framed CRA's mission around bringing industry, government, and academia together to strengthen computing research. Rather than presenting its activities as separate services, she described complementary ways of supporting the field: identifying future research directions, gathering evidence, developing researchers, building connections, and communicating community priorities. The Computing Community Consortium's Call for Innovative Research Visions offers one opportunity for researchers to contribute ideas and highlight their teams'

work. CRA-Industry's salon gatherings provide another form of engagement, bringing together industry leaders for candid conversations about shared challenges and emerging directions.

Both speakers emphasized that these connections depend on active participation. Allan encouraged attendees to attend sessions outside their usual areas, meet unfamiliar colleagues, and continue conversations during meals, networking activities, and beyond the Summit. Camp reported 441 participants representing 246 institutions and organizations, as well as a record 313 CRA member units, and credited the community's investment of time and expertise for CRA's growing engagement. This community investment was reflected in the volunteers recognized for their contributions across CRA's activities and in awards honoring people advancing the field through leadership, research innovation, service, and mentorship.

- **CRA Distinguished Service Award:** Dan Reed, University of Utah.
- **CRA Undergraduate Research Faculty Mentoring Award:** Lucas Bang, Harvey Mudd College, and Tom Williams, Colorado School of Mines.
- **CRA A. Nico Habermann Award:** Wendy Powley, Queen's University.
- **CRA Anita Borg Early Career Award:** Akshitha Sriraman, Carnegie Mellon University.
- **CRA Skip Ellis Early Career Award:** Elaine Short, Tufts University.
- **CRA Outstanding Undergraduate Researcher Award:** Ahan Mishra, Cornell University; Chau Vu, Pomona College; Haoquan Fang, University of Washington; Luisa Mao, The University of Texas at Austin; Qingyun Yang, Vanderbilt University; Ranya Sharma, University of Chicago; Sophia Tang, University of Pennsylvania; and Victoria Li, Harvard University.

Supporting Researchers and Institutions

Camp connected CRA's educational and mentoring programs with the need to understand how the field is changing. She emphasized the role of the Taulbee Survey and other data-collection efforts in supporting institutional decisions, noting that their usefulness depends on participation from the community. Her presentation also highlighted the Data Buddies Survey as part of CRA's work to inform computing education and workforce decisions. Alongside this evidence-gathering role, she described support for researchers across career stages, including more than 1,000 undergraduates trained through UR2PhD and mentoring opportunities through CRA-WP. These programs also illustrated the limits of available resources. Camp reported more than 500 applications for DREU-AI and DREU-QIS, but funding for only 63 undergraduates. The growth of these initiatives also underscored a continuing gap between demand and the opportunities CRA could support.

AI education was another area where Camp emphasized the value of shared resources and community input. She highlighted classroom-expansion conferences and education fellowships associated with the National Artificial Intelligence Research Resource, alongside the document presented as Defining AI in Computing Education, on which she invited feedback. These activities were intended to help the computing community develop a shared understanding

while supporting institutions' educational efforts. Camp also announced a forthcoming LEVEL UP AI survey of AI degree programs and drew attention to Lowering Barriers to AI Education at Institutions with Limited Resources, a set of NSF LEVEL UP AI workshop recommendations. Together, these examples illustrate CRA's use of convening, reporting, and educational support to help institutions address challenges that they would otherwise face separately.

Communicating Computing's Value

Camp described federal advocacy as another part of supporting the computing research community. Her examples included bringing organizations together in support of NSF, contributing responses to federal requests for information, and hosting policy town halls. She also encouraged participation in CRA's Congressional Visit Day as an opportunity for community members to engage with policymakers. These activities complemented the research and education programs discussed earlier by communicating the needs and contributions of computing research beyond the institutions conducting it.

Camp also raised the question of how the broader public understands AI and its relationship to computing. She emphasized that there is no artificial intelligence without computer science, connecting attention to AI with the research foundations that make it possible. She described efforts to bring other professional societies together in Washington to discuss challenges requiring a broader response. This concern with communicating computing's role continued into the audience discussion, where questions moved beyond engagement with Congress to public understanding of research, uncertainty about computing degrees, and the contributions that individual departments and researchers could make.

Questions and Answers

Q: Who is part of CRA, and how does an individual engage?

The speakers clarified that CRA membership is organizational rather than individual, encompassing academic units, research laboratories, professional societies, and other organizations. Individuals affiliated with member organizations can access membership benefits through those organizations. CRA primarily serves North America but also offers associate membership to organizations elsewhere. Opportunities for individual engagement highlighted during the session included participating in CRA's programs and events, contributing to surveys, providing feedback on shared resources such as the AI education document, and volunteering expertise across CRA's activities.

Q: What evidence is there of the impact of CRA's advocacy activities?

A: Camp pointed to congressional offices contacting CRA for input as an example of its engagement with policymakers. Peter Harsha indicated that the following morning's policy session would provide a fuller discussion of advocacy and its impact (covered later in the report).

Q: How has uncertainty around NSF funding affected CRA?

A: Camp said that Congress had rejected the most severe proposed NSF budget reductions and described the funding picture at the time of the Summit as closer to flat than to a major

cut, while acknowledging uncertainty for individual programs. She also described an opportunity for CRA to help sustain undergraduate research support: with reduced staffing affecting NSF's capacity to administer all the work that exists, NSF CISE asked CRA to help administer an REU stipend program. Camp reported that approximately 450 undergraduates had been supported across that effort. Her response emphasized that CRA stands ready to support NSF when challenges arise in service of the computing community.

Q: Should CRA complement its work with Congress with a broader public-facing campaign about computing research and education?

A: Allan shared that the Board had been discussing broader public communication and ways to explain computing's importance through accessible stories. Camp offered an example of a Northeastern video connecting computing research to safe milk production. She encouraged departments to help faculty explain the everyday significance of their work to the public, whose taxes support research.

Q: Could CRA issue a statement after the Summit about the value of a computer science degree and long-term career opportunities?

A: Allan welcomed the suggestion and invited the questioner to participate. He emphasized that such a statement would need to be a community effort rather than something produced by CRA alone.

The Transformation of Higher Ed in the Age of Disruption

Wednesday, July 22, 2026 | 9:30–10:30 am CT | *Presentation slides*

Introduction

- Keith Marzullo, 2026 CRA Summit Co-Chair; Dean, Manning College of Information and Computer Sciences, University of Massachusetts Amherst

Speaker

- Farnam Jahanian, President, Carnegie Mellon University

Key Takeaways

- **Rebuild public trust through communication and substantive change:** Universities must explain how education and research advance prosperity, security, culture, and the public interest while addressing concerns about affordability, relevance, workforce preparation, and intellectual openness.
- **Renew the federal-university research compact:** Sustained public investment in foundational research remains essential to national competitiveness and security.
- **Treat AI as an accelerant across computing and discovery:** Realizing AI's opportunities responsibly will require continued foundational research, shared

infrastructure, and collaboration among computing, the sciences, social sciences, and humanities.

- **Leverage AI to advance the science of learning:** Generative AI has the potential to realize more adaptive, personalized, and scalable education to improve learning outcomes.
- **Redesign institutions and education for a changing society:** Universities should organize flexibly around major challenges, prepare students for higher-order problem solving with AI, and create pathways for learners beyond traditional four-year degrees.
- **Human judgment and higher-order thinking will only grow in value:** As AI impacts all areas of the economy, the ability to think critically, ask the right questions, exercise sound judgment, and navigate complexity will become even more essential to innovation and leadership.

Summary

Keith Marzullo framed the plenary around leading computing research amid federal policy changes, industry competition, and rapidly evolving AI. Jahanian delivered a brief talk followed by a fireside conversation with James Allan about how higher education—and computing in particular—should respond.

Four Forces Reshaping Higher Education

Jahanian argued that higher education is experiencing more than a routine cycle of disruption. Technological, demographic, political, economic, and cultural changes are occurring simultaneously. First, public trust has eroded. He cited polling from June 2026 by the Lumina Foundation-Gallup on Confidence in Higher Education. This poll indicated that 63 percent of registered voters believe a four-year degree is not worth its cost and that confidence in higher education has fallen by 19 percent over the past decade. Concerns include cost, perceived political agendas, and misalignment between curricula and workforce needs.

Second, a demographic contraction is reducing the traditional college-age population, while immigration and visa uncertainty threaten international enrollment. Jahanian projected a 13 percent national decline in high-school graduates by 2041, with decreases concentrated outside the South (approximately 20 percent in the West, 17 percent in the Northeast, and 16 percent in the Midwest, compared with an increase of 3 percent in the South). Institutions will need to reach different student populations and reconsider established recruitment models.

Third, federal funding uncertainty risks undermining the approximately 76-year compact between government and research universities that helped produce advances from the internet to lifesaving medical treatments. Jahanian emphasized university research's importance to economic competitiveness and national security. Finally, AI is reshaping every sector of the economy while transforming knowledge production, research, and learning.

Renewing the University's Public Mission

Jahanian urged universities to return to first principles: advancing and preserving knowledge,

educating students, and serving society. Institutions must communicate these contributions while addressing access, cost, relevance, intellectual openness, and changing family expectations. Highlighting the messaging problem, he noted that at CMU, students from families earning below $100,000 can attend free of federal student loans, but the public is seemingly unaware of this.

Universities must also operate more nimbly. Recalling an exercise at CMU that asked faculty to imagine building a university from scratch, Jahanian described computing, AI, and data science as a core supporting grand challenges such as health and energy. Faculty should be organized around problems rather than constrained by departmental boundaries. Social scientists and humanists are essential partners in addressing AI's policy, ethical, governance, and human dimensions.

AI, Discovery, and the Future of Learning

Asked whether AI could narrow computing research, Jahanian described it instead as an accelerant enabled by decades of advances in algorithms, architecture, systems, networking, human-computer interaction, robotics, and other foundations. The apparent overnight emergence of generative AI obscures this history. A broad conception of computing is necessary to sustain foundational work and recognize opportunities throughout the research stack.

Jahanian illustrated AI's potential with a remotely accessible laboratory—CMU's AI Science Foundry—combining hundreds of instruments, robotics, computational biology, and machine learning to automate experiments and narrow large search spaces. He remarked, "this is like taking a wet-lab to the cloud." Such systems could help design experiments, integrate data beyond individual human capacity, and shorten the path from idea to scientific impact.

AI also requires reconsidering whom universities educate and what they teach. Jahanian rejected the claim that reduced demand for routine programming makes computing degrees less valuable. Computing is fundamentally about solving difficult problems; critical thinking, creativity, communication, and coordinating AI-enabled tools and agents will become more important. Students in other disciplines must learn to apply AI in their fields, while universities must create upskilling and reskilling pathways for adult and community-college learners. He cited gateway mathematics courses as a barrier, with an estimated 30 to 40 percent of some community-college learners leaving before completing their intended training. To address this, CMU launched a nonprofit initiative called Learnvia through a $55 million gift from the Bill and Melinda Gates Foundation to fix the high failure rates in introductory math courses, enabling a more streamlined pathway and access to education across colleges and universities.

Audience questions pressed whether higher education's credibility problem reflects poor communication or a genuine failure to deliver value. Jahanian answered that both must be confronted: change will occur with or without universities, and institutions must "meet the urgency of the moment." Unanswered questions revealed further curiosity about the evidence CRA should collect, online education, incoming-student preparation, and opportunities for a wider global computing community.

Advancing National Priorities and Seeding Future Computing Breakthroughs

Wednesday, July 22, 2026 | 6:30 pm–8:00 pm CT (dinner program) | Presentation slides

Speaker

- **Irina Dolinskaya**, Directorate Head for Computer and Information Science and Engineering (CISE), U.S. National Science Foundation (NSF)

Key Takeaways

- **NSF and CISE remain committed to supporting fundamental, curiosity-driven computing research**. Future CoRe programs remain in place, and CISE will continue to prioritize CAREER awards to support early-career researchers.
- **The NSF and CISE are undergoing a period of organizational change.** CISE has moved to a new structure centered around six science sessions, with the goal of improving coordination between different sections of the directorate. CISE is also currently understaffed, and is hiring both permanent staff and rotators.
- **CISE is also pursuing targeted programs in areas aligned with current national priorities including AI, quantum, and workforce development.** This effort includes expanding partnerships with industry and across the federal government.

Summary

Speaking from her position as Directorate Head of CISE, Irina Dolinskaya shared a snapshot of the NSF, focusing on what things are changing, what things are staying the same, and what new initiatives are being developed to support the future of computing research.

Dolinskaya opened by reflecting on the current moment and the past year of change. The year began amid the longest government shutdown in U.S. history. Once the government reopened, CISE underwent significant reorganization. Dolinskaya acknowledged that these changes have affected the community but also create opportunities. CISE remains committed to supporting effective, curiosity-driven research.

The CISE Reorganization

In accordance with a broader reorganization at the NSF, CISE is operating under a new organizational structure. Specifically, CISE has replaced its previous division-based structure with six science sections with the goal of improving coordination across the directorate. In addition, there have been significant personnel changes in CISE leadership.

Dolinskaya notes that this reorganization comes at a time when CISE is understaffed. However, the directorate is actively hiring for both administrative and scientific positions (including both permanent staff and rotators). CISE is also creating an assistant program director pathway for junior scientists who want to be program directors in the future. The hope is that this pathway can help junior researchers grow into the program director role.

What Is the Plan for Existing Programs Such as Future CoRe or CAREER?

Dolinskaya emphasized that the goal of CISE has not changed; CISE will continue to support fundamental, curiosity-driving research through programs such as Future CoRe. If anything, the scope of core programs have increased with the addition of Computing Education Research as a new initiative in Future CoRe. CAREER grants also remain a priority at CISE, as they are designed to support early-career faculty and thus help ensure the future of computing researchers and educators.

She also emphasized that CISE plans to continue funding research where industrial impact is not yet clear, highlighting the importance of NSF's past support of foundational research on the technologies driving the current AI boom. One of the main goals of CISE is to fund the research questions of the future, including those that have the potential to become the next AI.

CISE and the Federal Administration's Priorities

At the same time, Dolinskaya emphasized that much of CISE's work is closely aligned with the current federal administration's priorities. For example, CISE already supports research and education that aligns with federal priorities regarding strengthening AI foundations. CISE also supports some Quantum research, though Dolinskaya thinks that this is an area where the directorate can play a more integral role in the future. Dolinskaya also noted that CISE is well positioned to support building a skilled workforce. Advances in AI and quantum require not only new technology, but also students, educators, and researchers who can develop and use these technologies.

What New Programs Are in the Works?

As a result of this alignment with national priorities, CISE is also developing special programs as targeted investments. This includes programs aimed at bringing together infrastructure, researchers, and industry partners with the goal that they will form a mutually reinforcing scientific ecosystem. Dolinskaya briefly described several such initiatives including the NSF Open Multimodal AI Infrastructure to Accelerate Science (NSF OMAI), a partnership with NVIDIA aimed at developing fully open AI models to fuel U.S. scientific innovation. She expects that other such investments will continue to be announced, and encourages the community to keep an eye out for new calls for proposals.

Concluding Thoughts

Dolinskaya concluded by sharing that CISE turns 40 this year! She notes that it has been a phenomenal time over the past 40 years, and that with the help of the computing research community, she hopes we can make the next 40 years just as great.

Resources

- **CISE Directorate**
- **NSF OMAI**

Wind Against Tide: Computing Research Policy in Turbulent Times

Thursday, July 23, 2026 | 8:30–9:30 am CT | Presentation slides

Introduction

- **James Allan**, Chair, CRA Board of Directors; University of Massachusetts Amherst

Speaker

- **Peter Harsha**, COO and Senior Director, Government Affairs; Computing Research Association

Key Takeaways

- **The disruption goes well beyond the budget.** CRA's policy work has shifted from funding toward defending the research ecosystem on immigration, export controls, research areas barred from federal support, eliminated advisory bodies, and a depleted federal science workforce.
- **The administration's goals conflict with its methods, and that is the community's opening.** You cannot lead in AI and quantum while starving the broad research base and talent pipeline that feed them, and computing is well placed to make that case.
- **Congress remains the most effective check.** Coordinated advocacy helped persuade appropriators to reject a proposed cut of more than half of NSF's FY2026 budget. The strategy is now being tested again for FY2027.
- **The rules of research matter as much as the money.** OMB's proposed rule on federal financial assistance would subordinate peer review to political approval, penalize indirect costs, and restrict collaboration, publication, and whole areas of research.
- **Researchers are the most credible messengers.** CRA's advocacy works best when researchers carry the message themselves, through Congressional Visits Day, op-eds, local outreach, service on the Hill, or even running for office.

Summary

James Allan opened the morning with a playful caution about polished technology: the image of his student-era Apple II was generated by AI, and he had never actually used an Apple II. He then introduced Peter Harsha as a longtime guide to Washington who could be entertaining, frightening, and only occasionally comforting.

Harsha, now in his 25th year leading CRA's policy efforts, began by sharing credit with Brian Mosley, CRA's Associate Director for Government Affairs; the Government Affairs Committee, chaired by Alex Aiken of Stanford and vice-chaired by David Danks of the University of Virginia; and the many community volunteers whose voices give CRA's advocacy its credibility. He also reminded the audience that CRA is non-partisan. It supports no candidates, has no PAC, and judges policies only by whether they help or hurt the computing research ecosystem. If the talk

sounded aimed at one side, he said, that was because that side is currently promulgating the policies that concern the community.

He chose a nautical title because he sails, and because he was told he could not use his original one: “Everything Is Awful Pretty Much All of the Time.” A video of boats battering their way through Haulover Inlet illustrated the metaphor. When current and wind run against each other, steep standing waves wear down the crew and make progress hard. “Really, it’s just a really bad place to be,” he said, before inviting the audience to “cast off our dock lines” for a tour of the past 18 months, the reasons behind the disruption, the challenges ahead, and how the community can help.

A Pendulum Torn From the Ceiling

After some 30 years in science policy, Harsha said, he had prided himself on knowing how the pendulum swings. This time, it feels as if the pendulum has been ripped from the ceiling and is rolling around the floor. A timeline of the past 18 months set the administration’s actions, from executive orders and a grant freeze to NSF staff cuts, canceled awards, and a government shutdown, against CRA’s responses, including statements, joint letters with peer societies, a call to action, and an amicus brief supporting international students.

As a result, CRA has spent much of its time playing defense. Where funding once dominated its agenda, it now devotes far more effort to immigration and export controls; research areas, such as mis- and disinformation research, that are effectively barred from federal support; the quality of advice reaching the government after the elimination of the National Science Board and NSF’s advisory committees; and steep losses in the federal science workforce.

One Machine, Not Three

Why dismantle a research ecosystem that has been the envy of the world for 75 years? Before offering his answer (“Dark Peter mode”), Harsha established his bona fides: a conservative upbringing, Hillsdale College, work for Republican members of Congress and campaigns, and service on the Republican majority staff of the House Science Committee. Today’s GOP, he noted, is unrecognizable from the one he once supported, and he no longer associates himself with Hillsdale or the party.

He then described a theory, not his own, that Project 2025 states openly: that higher education, the press, and the federal science agencies are one machine, not three. Ideas born in federally funded university labs are normalized by friendly coverage and turned into rules by unelected agency staff, and the money flows back to universities to run the cycle again. The document calls federal spending “the secret lifeblood of the Great Awokening” and says indirect cost reimbursements “cross-subsidize leftist agendas.” In that view, the money is the machine, and cutting it is how to “defang and defund” the institutions they oppose.

The theory has real teeth, Harsha said, but it does not tell the whole story. A strike on ideology would not gut hurricane forecasting, theoretical computer science, or defense basic research, and Project 2025 itself is loud about AI, quantum, and competing with China. Other actors, including OMB Director Russ Vought, the President, and Stephen Miller, are also pulling levers, often against their own goals. They want U.S. leadership in AI and quantum, but they

discourage the world's best minds from coming here and assume that funding only those fields will get them there.

An image of a tree with a single AI branch being watered while the rest withered captured the contradiction. That, Harsha argued, is the community's opening: proving, on the administration's own terms, that you cannot lead in AI or quantum while starving the roots that feed them. Because those technologies depend so heavily on computing, the computing community is well placed to make that case.

Audience questions pressed on this mismatch, including a proposal to direct grants to principal investigators rather than universities and why education, a major national export, was being weakened through visa and hiring restrictions.

Congress as the Firewall

The President's FY2026 request proposed cutting NSF's budget by more than half, to levels not seen since 1999. The Government Affairs Committee chose to focus on Congress, not because it had shown much backbone, but because the dysfunction of the appropriations process could give members cover to find it. With a Schoolhouse Rock-style illustration, Harsha explained that Congress last passed all 12 appropriations bills under regular order in 1994. It now relies on continuing resolutions and omnibus or "minibus" packages. That is governance by brinkmanship, but it lets lawmakers fund science in a single up-or-down vote without visibly bucking the President.

CRA urged appropriators to reject the request, arguing that it would make the administration's own AI and quantum goals unachievable and cut programs Congress had long supported on a bipartisan basis. It issued a rare Action Alert, joined other computing societies in speaking with one voice, organized a statement from 86 industry researchers (many others were barred from signing by their companies), and held a record Congressional Visits Day with 50 participants and more than 130 meetings. Alongside the broader science community, the effort worked. The final FY2026 appropriation set NSF at $8.75 billion, 3.4 percent below the prior year, with the research account held level and education cut 20 percent. Declaring victory over a flat budget is a little sad, Harsha said, but compared to the alternative it was a huge vote of support.

The strategy is being tested again. The FY2027 request repeats the 55 percent cut to NSF and would eliminate the Social, Behavioral and Economic Sciences (SBE) directorate, a cut justified in the budget's "Cuts to Woke Programs" document. Harsha credited Brian Mosley with persuading 38 other STEM organizations, none of which receive SBE funding, to urge appropriators to reject the cut, a remarkable result in a parochial town. At the time of the Summit, only the House had approved its Commerce, Justice, Science bill, rejecting many NSF cuts but staying silent on SBE. The midterms and early talk of a December continuing resolution complicated the outlook.

Asked about long-term recovery, Harsha stressed mitigation first, then rebuilding from the federal-university compact that has historically supported basic research. Research pipelines and international talent can be weakened quickly but take much longer to restore.

Rules That Shape Research

Harsha described an even bigger threat: OMB's roughly 400-page proposed rule on federal financial assistance, published in late May. Until then, the administration had largely relied on executive orders, which cannot override existing regulations. This rulemaking would change the rules of grantmaking themselves in ways CRA believes would be catastrophic. Among other things, the proposed rule would:

- require pre-issuance approval of all discretionary awards by a senior political appointee, demoting peer review to an "advisory" role and introducing ideological filters into merit review;
- treat lower indirect cost rates as a preference factor in award selection, mischaracterizing F&A as administrative waste rather than essential research infrastructure;
- chill legitimate scientific exchange through overbroad foreign-collaboration restrictions;
- impose prior-approval requirements for conference and publication costs that run counter to the norms of computing research;
- create legal uncertainty for research on topics such as algorithmic fairness and bias mitigation; and
- reclassify the Uniform Guidance from guidance to regulation.

OMB allowed 45 days for comment with no extensions and aimed to implement the rule on October 1, 2026. Because co-signed comments would count as a single response, CRA filed its own and urged affiliates and individual researchers to do the same, while also organizing a joint release with seven peer societies. By the July 13 deadline, OMB had received 496,775 comments.

Looking ahead from July, Harsha expected OMB to miss October 1, since it must respond to substantive comments and submit the rule to Congress. A congressional disapproval resolution would face a veto, making litigation more likely, possibly with CRA's participation. The November midterms could change the calculus.

In the Q&A, an audience member noted that peer review has always been formally advisory. Harsha agreed, but distinguished between an agency retaining final responsibility and a process that effectively says, "do not listen to the review at all." Other questions covered publication and peer-review costs, the proposed FAIR model for indirect costs, and CRA's engagement with NIDILRR.

More Stiff Winds

Harsha also flagged other concerns. NSF has no Director, and nominee Jim O'Neill would be the first without an advanced degree. Science and Nature reported that about $1 billion in FY2026 NSF funds had not been distributed for awards, and Nature reported a $500 million clawback from CISE, MPS, and Engineering, which NSF has denied. CRA will work to get accurate information to the community as soon as it can establish what is actually happening. Still, winds shift, and November could bring a sea change.

How the Community Can Help

Harsha closed with ways to help. One of the most impactful ways to contribute is through Congressional Visits Day, held this year on September 22–23, where CRA trains participants and schedules educational, not adversarial, meetings with their own representatives.Researchers can also join the Friends of the Government Affairs Committee discussion list or take part in any CRA activity, nearly all of which are shaped by decisions made in Washington.

Beyond CRA, researchers can write op-eds and speak to local groups. Harsha recalled that Chairman Jim Sensenbrenner used to ask visiting researchers whether they had spoken at their local Rotary or Kiwanis club, since he had to be able to explain his support to those audiences. Researchers can also spend a sabbatical on Capitol Hill or at a think tank, or, for the real over-achievers, run for office.

He wrapped up with the hope that the community can hold on until the seas calm and it is back to "fair winds and following seas."

Final questions asked how to enlist technology companies with ties to the administration, coordinate university government-relations offices, arrange local meetings for those unable to travel, and pair evidence with stories that make the stakes visible. When a junior researcher thanked the advocacy team, Harsha replied that CRA cares deeply about protecting the system that allows research and researchers to thrive.

Fireside Chat with Chris Malachowsky, NVIDIA Co-Founder & NVIDIA Fellow

Thursday, July 23, 2026 | 9:30–10:30 am CT

Speaker

- **Chris Malachowsky**, Co-Founder and Fellow, NVIDIA

Moderator

- **Vivek Sarkar,** Dean, College of Computing, Georgia Tech

Key Takeaways

- **The benefits of AI depend both on how AI advances are deployed and also how they are communicated to the public.** Malachowsky argued that concerns regarding data centers, environmental impact, and jobs are important to consider when deploying AI. He also emphasized the importance of sharing the concrete societal benefits of AI, rather than advertising AI as an end goal in itself.
- **Malachowsky believes that with sufficient focus on education and access, AI has the potential to act as a great equalizer**. Malachowsky advocated for efforts to generally improve AI literacy and education.

- **Universities can act as centers of state-level AI infrastructure and education.** Malachowsky envisions a future with shared computing infrastructure, inter-disciplinary education, and public-private partnerships. He is working towards this goal with his *50-State Plan*.

Summary

In this session, Chris Malachowsky, a co-founder of NVIDIA answered a series of questions posed by both the audience and moderator Vivek Sarkar. The conversation spanned a variety of topics, from recollections about the founding of NVIDIA to perspectives on AI and its future in relation to society, the workforce, and education. Below, we provide a summary of the conversation, grouped by conversation topic.

Malachowsky's Story: Education, the Founding of NVIDIA, and Becoming a Leader in AI

Malachowsky began by describing how he got into engineering. When he started undergraduate at the University of Florida, he was planning on going into medicine. However, near the end of college, he decided to do engineering instead. Post graduation, he moved to California and worked at HP and then at Sun Microsystems. While at Sun, he met Jensen Huang and Curtis Priem. The three together would go on to found NVIDIA. Malachowsky next recounted how NVIDIA grew out of a series of conversations held at a local Denny's between the three cofounders where they would drink bottomless coffee for hours in the backroom. At its start, NVIDIA was excited by the possibility of hardware support for 3D graphics, with the initial idea being that they were going to try and build a game console on the chip. Malachowsky noted that this approach completely violated the rules of business, saying that the founders were initially enamored by building what was good for the customer. However, they quickly learned that in practice, they needed to sell their chips to computer manufacturers that were making decisions on existing benchmarks and reviews. Thus, to survive, NVIDIA decided to build what the market was asking for and just try and out-engineer everyone, which ultimately they were able to do.

Malachowsky next talked about how NVIDIA evolved from a computer graphics hardware company to a leader in AI hardware. First, he indicated that when starting, NVIDIA chose a place in the marketplace that had infinite head room. This is in part because NVIDIA is in a domain that is "disgustingly parallel." Each pixel really has nothing to do with the one next to it, and this allowed them to really just focus on parallelization and trying to make their products as parallel as possible.

Early on, graduate students realized because of this parallelization, they could use NVIDIA graphics cards to do their computing research work. NVIDIA realized this was important, and decided to make it easier to accomplish, leading to the development of CUDA. This effort led to more users and use cases, and allowed them to build the software and middleware muscles for allowing the GPU to reach its full potential. One of these use cases ended up being AI. He is proud that NVIDIA had the insight and foresight to expand their scope in this way.

Malachowsky's Perspectives on AI: Opportunities and Critique

Speaking in a personal capacity rather than representing NVIDIA's official positions,

Malachowsky also discussed his views on AI's societal implications. He acknowledged that some of the criticisms of large data centers are justified, particularly in cases where center infrastructure was built without sufficient attention to impacts on the community or the environment. At the same time, he argued that many conversations about AI collapse different technologies into a single category of "AI" or "Data Centers" which is often treating the problems with too broad of a brush, and thus obscuring the positive societal impacts of these emerging technologies.

More broadly, in his experience leading conversations with policymakers and community leaders with the words AI or research is often ineffective. Instead, Malachowsky argued that advocates for computing research should focus more on stories about specific scientific discoveries or advances in health or education. These concrete outcomes are often more compelling. That is, when arguing for the importance of AI infrastructure, in his view it is more important to talk about workforce development or AI literacy. Such activities will only use a fraction of these machines as the machines are ultimately for research, but this rhetoric helps everyone want the technology.

Several questions from the audience asked if AI will really reduce inequality or instead allow already well-positioned people and groups to pull further ahead. Malachowsky argued that AI has significant potential to act as an equalizer by increasing access to expertise and knowledge.

However, he stressed that this is not guaranteed; it is only possible if people have equal access to the technology and have the education and skills to use it. To accomplish this goal, he advocated for the importance of open models and a competitive AI ecosystem with many participants to spur innovation. He also noted that if AI makes expertise more commodified, aspects of a liberal arts education such as critical thinking and judgment may even become more important. Malachowsky also cautioned against businesses viewing AI only as a way to replace workers and lower costs, describing that approach as a "race to mediocrity." Instead, he argued that businesses should consider how to expand what they do rather than giving everything to an agent. Overall, he indicated that much remains to be learned, but he does not subscribe to the line of thinking that this is the end of jobs. Instead, he thinks there will be more jobs, but they will look different.

Education, AI Infrastructure, and the 50-State Plan

Throughout the conversation, Malachowsky argued strongly for teaching AI to everyone, regardless of if they were studying computer science. He acknowledged that many people are anxious about AI increasing inequality. From his perspective, for AI to realize its potential as an equalizer, people across disciplines and backgrounds must be able to use it, evaluate it, and apply it. This will require broadly teaching people judgment and critical thinking.
As an example of how his vision could be realized, Malachowsky discussed his experience and partnership with the University of Florida (UF). In 2021, Malachowsky provided the university with what was then among the most powerful computing infrastructures at a university. He did so because when asked what he would use it for, the president of the university said he'd teach AI to students across the curriculum. UF has successfully accomplished this goal. This may be in part due to UF being a land-grant state institution that is willing to share its

infrastructure. This year, when Malachowsky gave a commencement speech at UF, it was the first class where all four years had AI infused in all of their classes. Overall, Malachowsky thinks it is important to educate people regarding AI because he views this as essential for supporting the jobs of the future.

In addition, he thinks broadly teaching AI to the non-computing students is really going to go a long way to defusing anti-tech sentiments.

Drawing on his experience at UF, Malachowsky developed his 50-State Plan. His goal is for each state to develop shared AI computing and educational infrastructure, typically centered around its public university system. This initiative will require public-private partnerships among universities, local benefactors, and state governments. He believes it has the potential to boost state-level economic growth for individual states and the United States as a whole. At the time of the conversation, Malachowsky indicated that as part of the initiative he had already spoken with representatives of 30 states and visited 20 in person. He has also made NVIDIA colleagues available to meet with governors and school leaders. Overall, he can't think of anything he could do with his money more effective than helping the students coming out of state schools, as this will boost the economy of that state.

When asked if the CRA should also start engaging more directly with state governments rather than focusing on advocacy at the federal level, Malachowsky said it probably should. He argued that computing capacity should be viewed as an aspect of state infrastructure, and indicated that he thought that researchers are well positioned to explain how such investments can support research, education, workforce development, and regional economic growth.

Resources

- NVIDIA's 50 State Plan: https://resources.nvidia.com/en-us-her-industry-web-page/gtc26-s82238

"Doing Computer Science": The Case for Ethics at the Core of Computing as a Discipline

Thursday, July 23, 2026 | 6:30 pm–8:00 pm CT (dinner program) | Presentation slides

Introduction

- **Manuel Pérez-Quiñones**, 2026 CRA Summit Co-Chair, University of North Carolina at Charlotte

Speaker

- **Casey Fiesler**, Associate Professor of Information Science and William R. Payden Endowed Professor, University of Colorado Boulder

Key Takeaways

- **Two crises, one root cause:** The public backlash against AI and the tough job market facing graduates both trace back to how the field defines computing and what it leaves out.
- **Ethics is part of doing computer science:** In 1997, a letter in Communications of the ACM argued that discussing the social impact of computing "is not doing computer science." Fiesler made the case that it is, and that treating ethics a s a specialty shortchanges students.
- **Do both, but integrate if you can only pick one:** Standalone courses bring depth and expertise. Weaving ethics into technical courses reaches every student, keeps ethics in context, and does not hurt technical learning.
- **The barriers are structural, not a lack of interest:** Few instructors consider ethics irrelevant to computing. What they lack is room in the course, authority to change it, resources, incentives, and a community of colleagues.
- **Teach students to act, not just to notice:** Graduates are getting better at spotting ethical problems, but they often lack the power and the organizational skills, like persuasion and navigating office politics, to do something about them.

Summary

James Allan, chair of the CRA Board of Directors, opened the Summit's final session with a theme that had run through three days of conversation: computing shapes society, and the choices leaders make matter. Summit co-chair Manuel Pérez-Quiñones then introduced Casey Fiesler, one of the field's best-known voices on technology ethics and policy, with a large following on social media. Fiesler joked about her reputation as the "buzzkill ethicist" and called a room full of department chairs and research leaders her "ideal captive audience." She also pointed out that "ethics" had recently shown up in *The New York Times* crossword as the answer to a clue about what AI programmers should care about. Her argument was simple: ethics belongs at the core of computing as a discipline, not at its edges.

Two Reckonings

Fiesler started with two reckonings facing the field. The first is public backlash. Americans are more negative about AI than people in many other countries, and she described this "techlash" as a response to perception that the tech industry that has not been the responsible steward the public wants. It is not that people are scared or do not understand the technology; many understand it very well. Even computing students are sometimes strongly against using generative AI, and a common reason they give is that they want to keep using their own brains. Some also cite, for example, environmental concerns. The second reckoning is a jobs crisis. Students are being told that the jobs they trained for may no longer exist because AI can write code now, so the backlash is hardly surprising. Fiesler argued that the two problems are connected, and that both trace back to what we think computing is.

Is Ethics "Doing Computer Science"?

To make that point, she walked through the history of computing ethics education. In 1994, Batya Friedman and Peter Kahn asked why, if social and ethical concerns are an integral part of designing computer systems, they are so often missing from computer science education. Two years later, a proposal to make ethics a "tenth strand" of the CS curriculum drew a sharp reply in a 1997 letter to Communications of the ACM. Just as debating the morality of nuclear weapons is not doing physics, the letter argued, discussing the social and ethical impact of computing is not doing computer science. Fiesler said that letter probably would not be written today, but the field still has a long way to go. She tied this attitude to how the field treats its more human-centered areas, pointing to the essay "The Othering of HCI" and to research highlighting that HCI has the lowest subfield "prestige" in computer science.

Her interest in computer science education was sparked in part by Cambridge Analytica, and she quoted an observation in an op ed that software engineers "continue to treat safety and ethics as specialties rather than as the foundation of all design." Ethics, she argued, is not only a specialization. However, research studies have shown that ethics is often absent from technical courses such as AI and machine learning. If we teach students to do computer science and ethics is not part of it, she said, we are doing them a disservice. She also stressed that "ethics" does not have to mean moral philosophy; it is shorthand for a wide range of social and human concerns.

What the Research Shows

Fiesler then summarized what research says about teaching computing ethics. Standalone courses can be taught by experts with deep knowledge, whether from computer science or other fields, and they open the door to interdisciplinary collaboration. Integrating ethics into technical courses has its own advantages. It keeps ethics tied to the technical context, and also ensures that everyone who takes a single computing class hears about ethics, since "you can do a lot of damage with a little bit of code." It also boosts engagement and makes assignments feel more real. And research suggests that integrating ethics does not detract from technical learning and may even improve it. Ideally, departments do both. If they can only do one, integration is often the better choice.

Instructors are largely on board. Surveys of computing instructors found strong support for both standalone and integrated ethics education, with many agreeing that teaching ethics does not take away from technical material. Very few considered ethics irrelevant to computing, although some saw it as less relevant to particular computing topics. The real obstacles are structural: no room in the course, a set curriculum instructors are not allowed to change, a lack of resources and expertise, and concern that the work will not count in their department or toward tenure. Community makes a real difference. Research suggests instructors are more likely to include ethics when they have colleagues to talk it through with.

From Spotting Problems to Taking Action

The evidence beyond the classroom is encouraging but incomplete. Research has shown that students with ethics training are more likely to report encountering ethical issues at work, suggesting greater ethical sensitivity. Still, ethics is often an afterthought, such as a single lecture on the last day of a course, with no assessment. Many students doubt that knowing about ethics will help them get a job, and some worry that bringing it up in an interview could even hurt. And while students are getting better at spotting problems, a growing body of research also suggests that computing professionals often lack the power to act on them. Making responsible computing happen inside a company depends on persuasion, relationships, and organizational politics, which are skills the field rarely teaches.

That brought Fiesler back to her two reckonings. The techlash is not just an external PR problem, since many of our own students believe the tech industry does not care about ethics. And the parts of computing that AI cannot do, the human parts, are the same parts departments have traditionally treated as optional.

What Departments Should Do

Fiesler closed with four recommendations. First, make ethics core from day one. Second, teach students to act when they see unethical behavior. Third, fix what you reward: hire faculty whose research is in computing ethics and give them tenure, and also offer support and incentives for other faculty who teach the topic. Fourth, build community so that people in the department are talking about these issues together. She closed by arguing that doing computer science well means doing everything we can to make technology the best it can be, and that we should model for students what it means to be good at this work.

Audience Questions

One attendee, who had never thought about including ethics in their own classes, asked how to bring it into related fields such as electrical and computer engineering. Fiesler said the same conversation is happening across many disciplines. Another attendee asked how to handle pushback against bringing social and ethical considerations into research. She acknowledged that these issues are more fraught today than they have been in the past, and that being able to talk through these issues can be helpful. She suggested building extended communities of practice, and also writing about ethical decisions in research papers to help create norms.

Asked about teaching resources, Fiesler pointed to the Responsible Computing Challenge, which has funded this kind of work for years, and its playbook, as well as EngageCSEdu, ACM's collection of open CS education materials. Open resources, she said, are one of the best things we can offer each other. On funding, she noted how much an NSF Dear Colleague Letter on cybersecurity education had spurred both research and open resources. In response to another question she also pointed out that many people who teach ethics are adjuncts, which makes the work harder to sustain. Asked for examples of companies doing the right thing, she mentioned OpenAI's early bias-mitigation techniques. Responding to an analogy about safety standards on construction sites, she agreed that reasonable people will disagree,

which is central to doing ethics. Students need to learn to make trade-offs, explain them, and think ahead about "ethical debt," though many students today do not even know how to recognize it. One attendee noted the irony of saving the ethics keynote for the very end of the Summit, much as ethics often ends up in the last lecture of a course.

Conclusion

Fiesler's keynote tied the Summit's many threads to a single question: what it means to do computing well. Her answer was that the social and ethical side of computing is not an add-on but part of the work itself, and that departments can act on this now through their curricula, their hiring and reward systems, and the communities they build. James Allan closed the Summit by thanking its sponsors and volunteers and reminding attendees that the future of computing research will be shaped not by any single development but by what they do after they leave. The 2027 CRA Summit will be held in Seattle, Washington.

Resources

- Internet Rules Lab: Responsible Computing resources
- EngageCSEdu (ACM open CS education resources)
- The Othering of HCI (Michael Correll)
- Subfield prestige and gender inequality among U.S. computing faculty (*Communications of the ACM*, 2022)

Parallel Track Sessions

Track 1: Impact of AI in Computing Research and Education

When AI Goes In, What Comes Out? Leadership Lessons for Reshaping Computing Programs in this AI Era

Wednesday, July 22, 2026 | 11:00 am–12:30 pm CT

Moderators

- **Jon Bell,** Northeastern University
- **Borja Sotomayor,** University of Chicago

Panelists

- **Carla Brodley,** Northeastern University, Center for Inclusive Computing
- **Dan Grossman,** University of Washington
- **Stephanie Ludi,** University of North Texas
- **Steve Swanson,** University of California San Diego

Key Takeaways

- **Rebalance rather than simply add:** AI is entering programs through new majors, courses, certificates, and graduate pathways, but credit hours are fixed. New AI degrees are being built largely by shifting elective buckets toward AI, while non-AI foundations remain essential.
- **Treat curriculum change as a cultural process:** Changing a curriculum is as much social and political work as technical work. Leaders should say early which direction they want to go, bring every stakeholder in before decisions are made, and assume good faith in disagreements. Once the right answer for students is clear, they should be willing to hold the line.
- **Standardize the core, offer a menu above it:** A single AI policy across all courses does not work. But, intro courses need a common AI policy and common assessment across all sections; this means the same assignments, exams (or a bank of questions from which to pull exam questions), grading, and policies — otherwise student performance is too dependent on which instructor they had for the course. Upper-division instructors can choose their AI policy from a short menu of approved policies, which leaves room for those courses to experiment.
- **Treat AI access as an equity issue:** Premium AI subscriptions can cost $100 or more per month. Campus-licensed tools and self-hosted models, even if not frontier systems, keep students who can pay from gaining an advantage.
- **Redesign assessment rather than retreat:** Proctored work, computer-based testing facilities, frequent quizzes, and oral exams offer ways to protect integrity while innovating on how learning is measured.
- **Protect the human side of learning:** Shrinking TA budgets and AI-mediated help are changing how students get support. Struggling students are often socially disengaged, a problem AI can deepen, so TA training and engagement data matter more; particularly important is to teach TAs how to do code walkthroughs.

Summary

Jonathan Bell and Borja Sotomayor framed the session around a zero-sum problem: computing programs are adding AI tools, AI literacy, and new assessment approaches, but

credit hours are fixed. Rather than debating pedagogy or how much programming a 2027 graduate needs — questions left to the companion session, "What Should the CS Curriculum Be in the GenAI Era?" — the panel focused on how departmental leaders manage AI operationally. Carla Brodley, Dan Grossman, Stephanie Ludi, and Steve Swanson each opened with a concrete change from their institution before the discussion turned to four questions: updating curricula and degrees, guiding faculty, staffing and infrastructure, and downstream effects on the career and research pipeline.

Rebalancing Curricula and Launching AI Degrees

The opening examples showed how differently institutions are absorbing AI. Brodley reported that schools working with the Center for Inclusive Computing (CIC) are folding generative AI into the introductory sequence and creating "Math for AI" courses to shorten the chain of prerequisites before technical AI coursework.

Grossman described new courses at the University of Washington (UW), including human-AI interaction, systems for machine learning, and a 100-level AI principles course, along with graduate programs that assume less CS background. Grossman proposed building an "embarrassment list" of what a department should be embarrassed for a graduate not to know as a thought exercise that convinces people the full list has for many years been too long for all of it to be in required courses. Once you accept that "embarrassed list" items might only appear in elective courses (i.e., not all students will engage with them) you can create a rich curriculum that meets the varied needs and interests of a broad student population. UW's BS in CS is deliberately flat and flexible, with a recent move toward optional specializations that still leave room for electives.

At the University of North Texas (UNT), Ludi added AI-integrated assignments and deliverables to the undergraduate capstone, addressing both process and product, and worked with her department to create undergraduate certificates and a new degree. Swanson's change at UC San Diego (UCSD) was the most contentious: a strong shift toward heavily proctored assessment and computer-based testing, adopted with "a great deal of faculty hand-wringing."

Grossman cautioned that programs must still teach non-AI foundations and prepare students to be informed citizens, even as new AI degrees reach students from outside CS. Demand, however, remained uncertain. UCSD had not yet seen strong demand for its new AI major, which Swanson attributed partly to negative media coverage of AI and partly to student judgment: many still see the core CS major as the better choice. Brodley expected enrollment to follow with time and marketing, citing CIC's landscape survey showing roughly 40 percent more AI majors and minors [CIC reported 57 AI majors and 48 minors in Dec. 2025 (CRN); majors reached 107 by September 2026; see the September 2026 snapshot]. Asked what an AI bachelor's degree contains, panelists described incremental rather than wholesale change: UCSD added two AI fundamentals courses and ethics coursework while shifting elective buckets toward AI, and UNT added an experiential learning requirement. Ludi also noted that certificates work well, but departments should show students their value, for example that the certificate appears on the transcript alongside the degree.

Guidance Without Uniform Mandates

No panelist endorsed a single AI policy for every course.

Grossman described UW's approach of defining 4-5 policies and letting instructors choose among them, while developing some courses that actively encourage AI use. Brodley drew the line at the course level: instructors may keep their own lecture styles, but all sections of a course should share the same policies, assignments, exams, and grading. Her test for excessive variance was pointed: *can a department tell from students' preparation who taught their CS1 or CS2 section?* Ludi noted that some faculty are more interested in AI policy for their research labs than for their classrooms, and that classroom guidance ranges from encouraging to discouraging AI use. Swanson described UCSD's faculty working groups, which treat AI use as acceptable in principle.

The process mattered as much as the policy. Grossman argued that curriculum change is a social, political and cultural process as much as it is about content and advised bringing all faculty into discussions early and assuming good faith in disagreements. Brodley urged leaders to put students first and to emphasize the education mission in research-focused departments. She also noted wide variation in how programs interpret ABET criteria and encouraged faculty to review recent requirement changes, including mathematics credits [proposed CS program criteria changes clarify math requirements (CSAB); current 2026–27 ABET computing criteria].

Staffing, Access, and Assessment

Budget pressure shaped the discussion of infrastructure. Ludi described fewer TAs, more online office hours, and continuously evolving TA training as the relationship between TAs and students changes; campus-wide Copilot access and a Texas computing cluster help, but faculty are still seeking more. Grossman added that struggling students are often socially disengaged — a problem that predates AI but that AI worsens, and called for explicit efforts to (re)build culture among students, TAs, and instructors. Brodley said CIC is studying changes in TA usage and has identified that many institutions are working explicitly on culture: repurposing TA hours previously used for office hours to new recitations and labs. Swanson raised a structural issue: TA positions often subsidize PhD programs, a role he argued should be filled by bridge funding.

Access to AI tools raised a parallel equity concern. Grossman questioned how departments can avoid advantaging students who pay $100 or more per month for premium subscriptions, and favored self-hosted local models even when they are not frontier systems. Asked later how this approach plays out, he said it gives students new options.

Swanson presented UCSD's computer-based testing facility as a response to academic-integrity concerns: converting underused lab space into a testing facility with Docker-based environments is expensive up front but efficient to run, thanks to automated grading and reduced staffing. Grossman framed the moment as an opportunity to innovate rather than "reinvent 1965," through more frequent quizzes and experimental assignments. An audience member suggested that oral exams could also build social connection and interview skills. Swanson reported that they already work well in UCSD's large classes, while Brodley cautioned that TAs need training to assess students equitably.

Downstream Effects

Grossman expected introductory courses to attract fewer students and worried that majors may discount faculty advice when the world is changing so quickly. The research pipeline looked healthier: PhD applicants remain strong.

Conclusion

The panel treated AI less as a topic to add than as a pressure that forces departments to decide what is essential. Its most transferable strategies were menu-based faculty policies, consistency across multiple sections of courses, equitable tool access, secure and inventive assessment, and renewed attention to TA training and student community. Open questions remain about demand for standalone AI degrees and how departments will fund the staffing and infrastructure these changes require.

Resources

- CRA-E white paper: AI in Computing Education: Principles and Practical Guidance for Program Leaders (July 2026), co-authored by session organizers Jonathan Bell and Borja Sotomayor
- CIC map of AI programs in U.S. and Canadian universities: CICMAP.AI
- CIC three-question survey on changes in TA usage see the Next Steps slide in the session slides

AI and the Research Enterprise

Thursday, July 23, 2026 | 11:00 am–12:30 pm CT | Presentation slides

Moderators

- **Vandana Janeja,** University of Maryland, Baltimore County
- **Manish Parashar,** University of Utah

Panelists

- **Stephen Smith,** Carnegie Mellon University (AAAI)
- **Ramin Zabih,** Cornell University (arXiv)
- **Ellen Zegura,** National Science Foundation
- **Ben Zorn,** Microsoft Research

Key Takeaways

- **Scale is outrunning human capacity:** Submissions to venues such as AAAI and arXiv are growing faster than the community's ability to review, moderate, and read them, and a rising share of low-quality work compounds the strain.
- **Use AI to support the systems it is stressing:** AAAI added an AI review to every paper sent to full review, and arXiv now relies on automated screening to hold roughly 10 percent of submissions. Panelists treated these tools as supplements to human judgment, not replacements.
- **Design research artifacts for two audiences:** As agents increasingly read, synthesize, and produce research, papers and other artifacts will need both human-facing and agent-facing forms.
- **Treat access to compute as a national equity issue:** Access to models and computing resources remains highly uneven across the country. NAIRR and NSF infrastructure investments aim to broaden it, but other federal priorities may not share that goal.
- **Reward impact over volume:** Tenure and promotion decisions should prioritize impact, even though impact is hard to measure and credit for ideas will be harder to assign as ideas spread faster.

Summary

Vandana Janeja and Manish Parashar, co-chairs of the CRA/CCC AI Research Ecosystem (AIRE) Task Force, opened by describing research ecosystems as structures held up by supporting "pillars": funding landscapes, publication and reviewing systems, academic reward structures, and translational processes. AI is stressing these pillars while creating new ones.

AIRE is gathering community input through an online discussion forum and an in-person workshop to identify those stresses and develop responses for researchers and for leaders across academia, industry, and government. Emerging community pillars include peer review, infrastructure, interdisciplinarity, and public-private partnerships. The organizers noted that industry now plays a dominant role in AI research, but academia remains essential for alternative approaches, novel algorithms and architectures, new applications, and the future AI workforce. The panel then brought together perspectives from a scientific society, a preprint server, a federal funding agency, and an industry research lab.

A Research Ecosystem Under the Strain of Scale

Stephen Smith, president of AAAI, described a conference system straining under growth: AAAI-26 received roughly 30,000 submissions. In response, AAAI experimented with using large language models to identify candidate desk rejections and, in a pilot program, added one clearly labeled AI review to each of the 22,977 papers sent to full review to supplement the human reviews. The AI reviews were informational and visible to authors. Looking further

ahead, Smith expected AI to take on more incremental research, leading to a resizing of the research community and changes in how success is defined and measured.

Ramin Zabih, faculty director of arXiv, reported similar pressure. The archive hosts more than three million articles, received more than 284,000 submissions in 2025, and set a monthly record of 32,040 submissions in June 2026. It is seeing many new and unusually productive authors, many single-author submissions, and many more problematic papers, particularly in CS. Roughly 10 percent of submissions are now placed on hold, primarily through automated screening, and arXiv has introduced measures that raise the bar for some kinds of submissions. Zabih also flagged misconduct such as citation rings, and noted that arXiv's openness makes it a valuable source of training data.

Ben Zorn called the current shift perhaps the biggest change he had seen in his career. As research output and expertise grow exponentially, the gap widens between scale and the human capacity to respond, and much future content may be generated by AI. He argued that research has long assumed human authors and human readers; agents are not bound by those limits, so every research artifact may need both a human-facing and an agent-facing element.

Ellen Zegura endorsed the idea of human- and machine-facing research artifacts and emphasized the importance of human-AI teaming.

Access, Infrastructure, and Federal Priorities

Ellen Zegura, who has served as acting assistant director of NSF's CISE directorate, focused on access. Current access to models and computing resources is extremely uneven across the country, and she argued that NSF's role is to create opportunities everywhere. She pointed to the National AI Research Resource (NAIRR) as one mechanism for broadening access and described NSF's moves toward AI infrastructure and sustainable access within a federal landscape that places heavy emphasis on AI for science. She cautioned, however, that the goals of the Genesis Mission may not align with the goal of more even access. Zabih added that arXiv is in contact with companies about resources and sees opportunities for AI to support tasks such as summarizing papers' main points.

Rethinking Rewards and the Role of Humans

Audience questions turned to what these changes mean for junior faculty. Zegura argued that tenure decisions, in the best case, rest on impact; although impact is hard to measure, junior faculty should be encouraged to prioritize it. Smith agreed and maintained that conferences remain meaningful places for early-career researchers to meet colleagues. Zorn offered a more speculative view: as agents take on more research work, future meetings may include agents interacting with other agents.

On how reward and promotion structures should evolve, Zegura placed the question within a chaotic higher education landscape marked by disruption and declining public confidence in higher education as a path to a career. Zorn identified two emerging problems: crediting the originality of ideas when ideas flow much faster, and the ethical questions that arise when agents solve problems without human involvement. Asked about the role of humans amid a

chaotic reviewing process, Zegura answered clearly: the joy of being an educator lies in conveying an idea to another human, not to an agent, and researchers deserve to enjoy the work they do. Smith framed the challenge as one of governance, calling for guardrails and attention to ethical issues.

Conclusion

The panel closed on a note of cautious optimism. Zabih pointed to the accelerating pace of scientific discovery, and Zorn described a revolution that brings disruption but also technology that can improve the human condition; Smith agreed. Zegura called herself pragmatically optimistic. AI is lowering barriers to innovation and creativity and placing ethics at the center of the field, a topic students are eager to take up, yet she remains proud of resisting the hype. Taken together, the discussion framed AI as both a stressor and a potential support for the research enterprise, making community efforts such as AIRE essential for rethinking review, access, credit, and reward.

Resources

- CRA/CCC AI Research Ecosystem (AIRE) Task Force, including its community discussion forum and interest form
- AI-Assisted Peer Review at Scale: The AAAI-26 AI Review Pilot

AI Leadership in the Academy

Thursday, July 23, 2026 | 2:00 pm–3:30 pm CT

Moderators

- **Manish Parashar,** University of Utah
- **Amarda Shehu,** George Mason University

Panelists

- **Vipin Kumar,** University of Minnesota
- **Sudeep Sarkar,** University of South Florida
- **Shishir Shah,** University of Oklahoma

Key Takeaways

- **Coordinate without smothering:** AI programs, certificates, and courses are springing up all over campus. Universities are creating chief AI officers, task forces, and working groups to connect these efforts without shutting them down.

- **Make AI literacy everyone's job:** Beyond AI degrees, institutions are building AI into majors and first-year courses so that every student understands AI as more than ChatGPT.
- **Don't make students pay for tools:** Campuses are negotiating enterprise access with major vendors, using on-campus and state clusters, and testing whether open models are good enough for most coursework.
- **Protect research quality as AI speeds it up:** AI is accelerating writing, proposals, and discovery, but panelists warned that AI-assisted work may converge on a narrower set of ideas and that clear rules are needed on what is allowed.
- **Build trust before scaling up:** Privacy worries about university-provided tools, IT concerns about agents, and unclear data-use rules all call for clear guidelines and transparent practices.

Summary

Manish Parashar and Amarda Shehu framed the panel around a problem facing senior university leaders: AI is reshaping research, education, and community partnerships all at once, and someone has to coordinate that change across academic units and decide where the money goes. The panelists all hold formal AI leadership roles. Shehu is George Mason University's vice president and chief AI officer. Shishir Shah is the University of Oklahoma's first chief AI officer and the new director of its School of Computer Science. Sudeep Sarkar, a computer vision researcher, served as launch dean of the University of South Florida's Bellini College of artificial intelligence, Cybersecurity and Computing. Vipin Kumar directs the University of Minnesota's Data Science Initiative. The moderators organized their questions around education, research, resources, coordination, and accountability.

Coordinating AI Across Campus

Shah described his job as pulling AI efforts together across the university: helping research units adapt, working with the legal office, and making sure people without a technical background can benefit from AI. Sarkar argued that the field needs new metrics for evaluating different kinds of schools, and noted that departments across USF face many of the same challenges. Asked whether the wave of new AI programs, certificates, and concentrations should be decentralized or coordinated, he described embedding AI into majors across the university rather than keeping it inside computing.

AI Literacy for Every Student

Shah said Oklahoma is tackling AI literacy across schools and programs, starting with a suite of first-year experience courses. Many students, as he noted, do not realize that AI is a much broader field than ChatGPT and the other tools they use every day. Since faculty and students are short on time, the university is also making online courses easy to access for additional training.

Kumar stressed that different disciplines need different kinds of AI education; social scientists, for example, see AI very differently than computer scientists do. He added that evaluating student work has become a major problem, because generative AI produces plausible answers without knowing whether they are right. Asked what an AI minor requires, Sarkar described a sequence that includes AI fundamentals, application courses such as no-code AI, and a more in-depth introduction to AI.

Tools and Infrastructure

Every institution has had to figure out how to provide AI tools. Shehu said George Mason does not want students to pay $20 to $100 a month for subscriptions, so a 70-person task force worked on tool access, and Microsoft and Google tools are now available campus-wide. Cybersecurity remains a central concern, along with the challenge of keeping up. Sarkar described a similar path at USF: on-campus clusters, Microsoft cloud services, and Copilot access, with staff adopting the tools faster than faculty outside AI. USF is also working with AWS on sandbox environments and uses low-cost clusters provided by the state. Kumar reported campus-wide Gemini access at Minnesota, backed by university clusters. Shah offered a data-driven check on costs: Oklahoma's analysis of course-related queries found that open models are good enough for most needs, and the university is continuing to study how students use them.

Research at AI Speed

Panelists agreed that AI has changed how research gets done, and that everything, including publishing, now moves faster. Sarkar cited research showing that AI-assisted work tends to cluster around similar ideas, while work done without AI is more diverse [this Nature study], and urged researchers to use AI carefully. Kumar noted that AI now helps write papers and proposals, which raises the question of what should and should not be allowed. He sees tremendous opportunity in AI for science, but said the successes so far are isolated. Shah described internal tools that help research teams come together and the university's interest in building larger teams, while asking whether anyone actually reads the growing flood of output.

Trust, Privacy, and Governance

One audience member said they would not use university-provided AI tools because administrators might be able to read their chat histories. Shah acknowledged that the university probably could, but said it does not have the storage to keep them; Shehu added that no one wants to hold on to chat histories. Agents raise further questions. Shah said his IT department is nervous about agents even as faculty use them for research, and Sarkar reported that OpenClaw-style agents are already producing interesting findings. Oklahoma is responding with university guidelines drafted by a working group, clearer rules about what information can be collected in research settings, and internal seed funding that requires detailed information about data and how it will be used.

Conclusion

Asked how they would judge success five years from now, the panelists were candid about the uncertainty. Shah admitted that universities are a step behind what is happening, and Shehu said leaders are still working out what their own role should be while continuing to support the many roles AI touches. The session made clear that AI leadership in the academy is still being invented. For now, the work centers on connecting scattered efforts, giving everyone fair access to tools, protecting research quality, and building the trust that campus-wide adoption will require.

Resources

- Artificial intelligence tools expand scientists' impact but contract science's focus (Nature, 2026)

What Should the CS Curriculum Be in the GenAI Era?

Thursday, July 23, 2026 | 4:00 pm–5:30 pm CT

Moderator

- **Manuel Pérez-Quiñones**, University of North Carolina Charlotte

Panelists

- **Steven Gordon**, The Ohio State University
- **Leo Porter**, University of California San Diego
- **Titus Winters**, Adobe

Key Takeaways

- **From Code Production to Software Judgment:** generative AI reduces the relative importance of producing routine code and increases the importance of understanding problems, reading and evaluating code, testing, design, architecture, and judging software quality.
- **Integrate Generative AI Without Outsourcing Learning:** Students should learn to develop software with generative AI, but curricula and assessments must still verify that they can understand, critique, and work independently of the tools.
- **Curriculum Change Requires Subtraction, Not Just Addition:** The existing curriculum cannot simply absorb more generative AI content. Courses and requirements must be reconsidered, consolidated, or removed, although the panelists differed on whether this should be a radical rebuild or an incremental redesign.

- **Faculty and Institutional Adaptation Are the Immediate Bottleneck:** Accreditation does not appear to prevent change, but faculty expertise, incentives, disciplinary identity, and departmental governance may. Waiting for stable tools or conclusive research risks allowing other departments and educational providers to define AI-assisted computing education first.

Summary

During this session, Manuel Pérez-Quiñones of the University of North Carolina Charlotte moderated a panel on how computer science education should respond to generative AI. Leo Porter drew on his experience redesigning introductory programming courses and building a faculty community around generative AI integration, Steven Gordon presented findings from an ACM education advisory task force survey, and Titus Winters offered an industry perspective on changes in software development. Although the session was framed around whether to integrate generative AI into a largely familiar curriculum or rebuild the curriculum around a different developer role, the panelists broadly agreed that routine code production is becoming less central while problem formulation, code comprehension, testing, design, architecture, and software judgment are becoming more important. They then considered how to integrate generative AI without allowing students to outsource learning, how assessment should combine controlled demonstrations with authentic open-ended work, and what existing material must be reduced to make room for new priorities. Porter favored rapid experimentation within existing courses, while Winters argued for more fundamental reconstruction, and Gordon emphasized changing assessment practices and supporting faculty development. Across these differences, the panel agreed that waiting for the tools or evidence to stabilize is itself risky because industry practice, student expectations, and programs outside CS are changing faster than conventional curriculum processes.

From Code Production to Software Judgment

The panel largely agreed that generative AI is changing what it means to be competent in software development, even though students still need to understand and write code. Leo Porter described a redesigned CS1 course that allows students to use generative AI to build substantially larger projects in areas such as data science, image processing, and games. Rather than evaluating only whether the submitted program works, the course asks students to explain why they selected their approach, how they decomposed the problem, how the functions fit together, and how one function works in detail. Titus Winters framed the same shift from an industry perspective: as AI becomes increasingly capable of producing functional code, the important human contributions move toward understanding the problem domain, evaluating possible solutions, designing systems, testing behavior, and recognizing failures. Steven Gordon reinforced this direction through survey findings showing that many programs are already asking students to analyze and critique code.

Winters described the broader educational goal as teaching “taste”: the ability to recognize whether software is well designed, maintainable, efficient, and appropriate for its context. Students can develop this judgment through sustained exposure to good examples, comparisons between alternative solutions, code review, benchmarking, and discussion of why one approach is better than another. Porter connected these activities to the actual work of

new developers, who are often asked to enter a large existing codebase, understand it, locate a problem, and make a targeted change rather than write a small program from scratch. The panel did not argue that algorithms, complexity, or other foundations should disappear. Instead, Winters argued that students should gain a conceptual understanding of when ideas such as asymptotic complexity matter, when other factors such as memory behavior dominate, and when a particular technique or tool should be applied. A university CS degree would therefore remain distinguished by its breadth, depth, critical thinking, and systems-level understanding rather than by fluency with any one programming tool.

Integrating Generative AI Without Outsourcing Learning

Porter emphasized that integrating generative AI into the curriculum does not mean allowing students to outsource the learning objectives of a course. In his CS1 model, students can use generative AI while developing projects, but the underlying programming and software-development skills are assessed without access to the tools. The required project walkthroughs provide an additional opportunity to confirm that students understand the code they submit. Responding to concerns that beginners who rely heavily on AI may learn very little, Porter distinguished between using AI to support development and using it to avoid learning. His proposed safeguard was direct assessment in a controlled environment: students who cannot demonstrate the required skills independently should not pass. He also identified metacognition as a new learning objective, including the ability to recognize when an AI system is leading a project in the wrong direction and decide whether to inspect the code, reformulate the problem, or continue working through the agent.

The discussion also exposed a tension between scalable assessment and authentic software work. Winters argued that open-ended exercises such as code review can evaluate whether students recognize functional bugs, architectural violations, quality problems, and questionable design choices while also teaching them how to communicate critical feedback professionally. He cautioned that mechanical rubrics and conventional autograders are poorly suited to evaluating qualities such as taste, elegance, and maintainability. Pérez-Quiñones offered a partial counterpoint, noting that autograders can still evaluate useful properties such as test coverage and style. Taken together, the discussion suggested that no single assessment model is sufficient. Controlled individual assessments can verify that students possess foundational knowledge, while larger projects, explanations, reviews, and comparative evaluations are needed to measure whether they can apply that knowledge to realistic AI-assisted software development.

Curriculum Change Requires Subtraction, Not Just Addition

Winters argued that departments cannot respond to generative AI by adding new material on top of an already crowded curriculum. In his view, the traditional CS degree devotes too much of its limited instructional time to producing code and completing tightly specified tasks that AI can now perform well. This creates an opportunity to reconsider topics and courses that remain primarily because they have always been taught. He offered circuit switching and the breadth of traditional data-structures instruction as provocative examples of material that might be

reduced to create room for testing, design, architecture, code comprehension, teamwork, and domain knowledge. Pérez-Quiñones similarly observed that many courses are organized around systems computer scientists needed to build decades ago and suggested that some low-level details distributed across several courses might be consolidated. Porter summarized the contrast succinctly: industry usually asks universities to add more content, while Winters was explicitly asking them to cut.

Porter agreed that substantial change is necessary but disagreed with Winters' more radical "burn everything down" framing. He warned that comprehensive redesign efforts are often assigned to committees and then make little progress. His alternative was to let motivated faculty begin changing existing courses, demonstrate workable models, and build broader adoption from those examples. This produced one of the panel's clearest differences in approach: Winters advocated stripping the curriculum down and rebuilding it around present needs, while Porter favored immediate experimentation within existing institutional structures. Both approaches would preserve important foundations but change how they are taught. Algorithms courses, for example, could place less emphasis on implementing every technique by hand and more on understanding what problems a technique addresses and when it should be used. Porter also suggested reducing the number of required courses, and potentially some mathematics requirements, to create space for practical, human-centered, and domain-oriented computing courses. ABET requirements were not viewed as a fundamental obstacle, although Winters argued that accreditation guidance should evolve more frequently in response to changes in professional practice.

Faculty and Institutional Adaptation Are the Immediate Bottleneck

The panel repeatedly returned to the difficulty of changing faculty practices and departmental structures. Porter noted that revising courses is often not a high faculty priority and that many instructors do not yet feel qualified to teach software development with generative AI. Gordon's survey work similarly revealed a wide range of attitudes, including faculty who refuse to permit the tools under any circumstances. Porter warned that a prolonged "wait and see" response carries its own risks. Students may encounter curricula centered on tasks that AI can already perform and conclude that CS teaches obsolete or impractical skills. Pedagogical research is also moving more slowly than the technology: by the time a study is designed, completed, and published, the tools it examines may already have changed substantially. Meanwhile, programs in business, engineering, data science, and other disciplines are beginning to teach AI-assisted programming themselves, potentially defining this area before CS departments decide how to respond.

The discussion nevertheless identified several mechanisms for moving forward. Porter described the CS and generative AI Consortium as a source of adaptable course materials and a community of practice, with more than 2,000 educators participating in its events and more than 300 faculty joining its online community. Gordon emphasized the need for formal professional development rather than expecting faculty to retrain themselves individually. When departments encounter sustained resistance, Gordon argued that they may eventually need to proceed without unanimous agreement. Porter offered a more collegial version of that strategy: identify courses that resistant faculty can continue to teach effectively, move some traditional material into electives, and prevent a small group from exercising veto power over the larger

curriculum. Pérez-Quiñones suggested using an external facilitator to reduce internal conflict. Winters framed the deeper challenge as one of professional identity. Faculty and practitioners who define themselves by expertise in a particular language or system may experience AI as a threat; defining themselves instead as problem solvers who can learn and adapt makes it easier to approach the transition collectively.

Track 2: Next Breakthroughs in Computing

Computing for Physical Intelligence/Robotics

Wednesday, July 22, 2026 | 11:00 am–12:30 pm CT

Moderator

- **Martial Hebert,** Carnegie Mellon University

Panelists

- **Maria Gini,** University of Minnesota
- **Basti Scherer,** Carnegie Mellon University
- **Peter Stone**, University of Texas at Austin

Key Takeaways

- **Robotics is a systems discipline:** Intelligence in the physical world requires sensing, actuation, real-time computing, learning, and decision-making under uncertainty, bringing together many areas of research and education.
- **Robotics is a "gift" to CS education:** Robotics gives students a concrete and engaging setting in which to connect topics that are often taught separately.
- **Infrastructure and access currently limit educational scale:** Meaningful robotics education requires experience with physical hardware, but robotic platforms are expensive, difficult to maintain, and rapidly changing, making it difficult to provide hands-on opportunities at scale.

Summary

During this session, Martial Hebert facilitated a panel on robotics and Physical AI. Maria Gini, Sebastian "Basti" Scherer, and Peter Stone drew on their experiences in research, teaching, and industry to examine what distinguishes intelligence in the physical world and what that distinction means for computing education. Through examples and demonstrations, the panelists emphasized that deploying a robot requires the integration of perception, actuation, learning, real-time computing, and decision-making under uncertainty. The discussion then

turned to robotics as a "gift" to computer science education: a motivating application that connects topics often taught separately. The panel also considered how curricula should balance disciplinary depth with interdisciplinary breadth, why students need hands-on experience with hardware, and how the cost and rapid evolution of robotic platforms limit educational scale.

Robotics as a Systems Discipline

Much of the panelists' opening remarks and the subsequent discussion focused on embodiment as the defining characteristic of robotics and Physical AI. Operating in the physical world introduces uncertainty, timing constraints, and system-integration challenges that do not arise in the same way for purely digital technologies. These challenges have shaped robotics research for decades and now also shape how robotics should be taught.

Gini organized the foundations of Physical AI into four areas: sensing and perception, robotics and actuation systems, edge computing, and foundation models. Sensing and actuation allow an intelligent system to perceive and affect its environment, while edge computing supports the real-time decisions required for physical operation. Foundation models can support generalization and decision-making, but they must operate as part of a larger physical system. This led Gini to pose a broader educational question: If teaching robotics requires teaching AI, should teaching AI also require some understanding of robotics and physical intelligence?

Scherer expanded on the role of foundation models in robotics by drawing on his work at Field AI. He described foundation models as having "gone physical" and characterized deployment as a new data engine. Language-based AI models can draw on enormous quantities of data from the internet, but robotics does not yet have an equivalent collection of data about physical interaction. Deploying robots can help generate that data while revealing the conditions under which systems succeed, fail, or need to improve.

At the same time, Scherer emphasized that a foundation model is only one component of a robotic system. Increasing model or dataset scale does not eliminate the uncertainty of operating in the real world. Physical AI therefore requires more than programming or data science; it also depends on an understanding of estimation, control, mathematics, system behavior, and the domain in which a robot will operate. Students and researchers must be able to reason critically about how these components work together rather than evaluating a model or algorithm in isolation.

Opportunities for CS Education

These system-level demands also explain why Scherer described robotics as a "gift" to computer science education, a sentiment echoed by Stone. Robotics gives students an exciting, concrete setting in which to connect artificial intelligence, perception, systems, security, human interaction, and real-time computing. Building a system that must operate in the physical world requires students to practice testing, systems engineering, and critical thinking while confronting the consequences of decisions made by individual components. Stone also pointed to robotics clubs and competitions as environments in which many

students first develop technical, teamwork, and leadership skills, illustrating the broad appeal of robotics as an entry point into computing.

The panelists emphasized that robotics curricula must balance interdisciplinary breadth with depth in a traditional discipline. Stone described the approach of Texas Robotics, which offers an undergraduate minor and honors program rather than a standalone robotics major. Students remain grounded in a home discipline while gaining cross-departmental robotics experience, including work with physical systems. Carnegie Mellon similarly incorporates both breadth and depth into its robotics programs, while the University of Minnesota has expanded the hands-on component of its master's program through courses using manipulators and the acquisition of new humanoid platforms. Across these models, the panelists emphasized that meaningful robotics education requires experience with hardware, even though providing that experience becomes increasingly difficult as enrollment grows.

Challenges and Limitations in Robotics Education

The same physical embodiment that gives robotics its educational value also creates its central scaling challenge. The need for hands-on experience has traditionally kept robotics courses small, even at universities with substantial resources. Institutions must acquire enough platforms for students to use, maintain those systems, and replace them as robotics hardware and software evolve.

Simulation can help expand access and give students a place to develop initial skills before working with hardware. Stone nevertheless cautioned against simulators with the popular quote by roboticist Rodney Brooks: "simulations are doomed to succeed." Simulated environments cannot reproduce all of the uncertainty, failures, and unexpected interactions that students encounter when deploying a physical system. Simulation is therefore a valuable starting point, but it cannot fully replace hands-on experience.

Scherer suggested that smartphones may offer a lower-cost entry point. Modern phones already contain cameras, other sensors, communication capabilities, and substantial computing power. He pointed to a program that turns a cellphone into a Robot Operating System (ROS) node, illustrating how an existing device could serve as the sensing and computing component of a robotics platform when paired with actuation.

Access also depends on curriculum and faculty expertise. The panel discussed the value of shared materials that would allow computer science faculty without a robotics background to teach introductory courses. Audience members pointed to curriculum resources being developed by Chad Jenkins and collaborators at the University of Michigan as one possible foundation for this work. CRA could help broaden access by distributing adaptable curricula and supporting Physical AI and robotics opportunities through existing undergraduate research programs. Such resources would need to emphasize durable principles and remain flexible enough to accommodate platforms that change much more quickly than university curricula.

Conclusion

Taken together, the panel presented robotics as both a rapidly evolving area of computing research and a powerful context for computing education. Its systems nature requires the

integration of learning, sensing, actuation, real-time computing, safety, and decision-making under uncertainty, while giving students a motivating way to understand how these topics fit together. Yet the embodiment that makes robotics educationally valuable also makes it difficult to scale. Expanding access will require curricula that preserve disciplinary depth while creating interdisciplinary pathways, shared materials that faculty without robotics expertise can adopt, and practical approaches for giving more students meaningful experience with physical systems.

Privacy, Trust, and Regulatory Compliance: It's Not Just for Lawyers Anymore

Thursday, July 23, 2026 | 11:00 am–12:30 pm CT

Moderator

- **Hank Korth,** Professor, Lehigh University

Panelists

- **Ed Felten,** Co-founder and Chief Scientist, Offchain Labs
- **Ian Miers,** Assistant Professor of Computer Science, University of Maryland

Key Takeaways

- **Move from trusting organizations to verifying claims:** Zero-knowledge proofs can allow people and businesses to demonstrate identity, authorization, or regulatory compliance without surrendering the underlying private data.
- **Treat compliance as a living system, not a fixed checklist:** Real requirements are often stateful, risk-based, partly confidential, and continually adapted in response to fraud. Cryptography can enforce a policy, but deciding what policy should be enforced remains a legal, institutional, and technical challenge.
- **Minimize disclosure without overlooking abuse:** Age verification and proof of personhood demonstrate the promise of privacy-preserving identity, but stolen credentials, fake accounts, and determined attackers require protections beyond zero-knowledge proofs alone.

Summary

Hank Korth opened the panel by arguing that technology has reached a point where privacy can increasingly be achieved through proofs rather than litigation or trust in a third-party custodian. He introduced Ed Felten and Ian Miers and guided a conversation about the opportunities created by zero-knowledge proofs, the complications of real-world compliance, and the responsibilities of computing researchers as these systems move into practice.

From Trust to Proof

The panel contrasted past, present, and possible future approaches to privacy. Traditionally, individuals hand their information to banks, companies, governments, or online platforms and rely on laws, contracts, and privacy policies to constrain its use. A self-sovereign model reverses that relationship: individuals retain their data and disclose only what a transaction requires, with cryptography providing assurance that the disclosed claim is valid.

This model could also relieve individuals of burdens they rarely manage effectively. Most people care about privacy but routinely accept lengthy privacy statements and cookie notices without reading them. The panel suggested that the strongest incentive for adoption may therefore come from businesses seeking to reduce liability, avoid costly disputes, and verify claims without acquiring sensitive information they must then protect.

The growth of stablecoins and AI agents makes this possibility especially timely. An AI agent could receive a limited payment budget rather than access to an owner's entire bank account, but parties would still need ways to establish whom an agent represents and whether its transactions satisfy applicable rules. Privacy must consequently coexist with legitimate requirements to prevent fraud, money laundering, and other abuse.

Zero-knowledge proofs are not new, but advances in performance and programmability have made them increasingly relevant to practical systems. The panel used a "Where's Waldo?" example to introduce the basic idea before turning to applications. A medical service, for example, might prove that it evaluated committed patient data using an approved model without revealing either the patient's information or the provider's proprietary model. Similar methods could support private audits, tax filings, supply-chain verification, sustainability claims, and product provenance.

Compliance Beyond a Yes-or-No Predicate

Felten then challenged the appealingly simple view that compliance consists of encoding a list of rules and proving that each transaction satisfies them. Actual requirements come not only from laws and regulations, but also from contracts, company policies, audit obligations, conflicts of interest, and reputational commitments. Some requirements may be confidential, while others depend on earlier transactions or patterns of activity rather than a single event.

Compliance is also rarely a collection of unambiguous yes-or-no decisions. Organizations often manage degrees of legal and operational risk. Their enforcement mechanisms must change as adversaries learn how existing controls work, creating an ongoing arms race against fraud and abuse. Organizations may need to update their rules, keep parts of their detection strategies private, and measure both the amount of evasion and the tactics used by evaders.

A cryptographic proof can demonstrate that a specified rule was followed, but it cannot determine whether that rule adequately captures an organization's legal and social obligations. The discussion therefore pointed toward deeper collaboration among computer scientists, lawyers, policymakers, and domain experts — not the replacement of one profession by another.

Identity Without Oversharing

Miers shifted the discussion from payments and compliance to identity. Today, proving that someone is old enough to enter a venue or access an online service can require revealing an identity document containing substantially more information than necessary. A privacy-preserving credential could instead prove only that the person is over the required age and that the claim is backed by a valid government credential.

Proof of personhood raises a harder version of the same challenge. Online services increasingly need to determine whether they are interacting with a human rather than a bot, but invasive monitoring would undermine the civil liberties these systems are intended to preserve. Attackers may also buy legitimate accounts, steal credentials, or recruit real people to pass checks for them. Zero-knowledge proofs can reduce disclosure, but they do not by themselves solve credential theft, network-level surveillance, or the continuing contest between attackers and defenders.

From Possibility to Practice

Audience questions explored why self-sovereign privacy may be more viable today than it was during the cypherpunk movement three decades ago. The panel pointed to faster and more accessible proof systems, the growing economic value extracted from personal data, and the arrival of mainstream businesses with practical reasons to protect it. Blockchain infrastructure may increasingly function as an unseen settlement layer rather than as a product users must consciously adopt.

Adoption nevertheless remains uneven. Succinct proofs have become common within blockchain systems, while large institutions in areas such as healthcare remain cautious about integrating them into established workflows. Panelists identified possible applications ranging from supply-chain compliance and media provenance to arms control. They also discussed specialized alternatives, including redactable signatures, that may prove more efficient than general-purpose proof systems for particular tasks.

Questions about artificial intelligence exposed another important boundary. Proving that a particular computation was executed correctly is becoming increasingly practical; proving a broad property of a model — such as that it can never exhibit a prohibited behavior — is much harder. Likewise, cryptography cannot put information that is already public “back in the bottle,” but it can reduce how much additional information future systems collect.

The panel closed by emphasizing the gap between what privacy-enhancing technologies can now support and the mental models that still shape policy. Computing researchers have an opportunity not only to improve these systems, but also to help policymakers recognize that privacy, security, and regulatory oversight need not always be competing goals.

Cybersecurity for the Public Good

Thursday, July 23, 2026 | 2:00 pm–3:30 pm CT

Moderator

- **Susan Landau,** Tufts University

Panelists

- **Zakir Durumeric,** Stanford University
- **Yoshi Kohno,** Georgetown University
- **Daniel Lopresti,** Lehigh University

Key Takeaways

- **Define cybersecurity around people, harms, and more:** Beyond preventing unauthorized access, cybersecurity must address disinformation, online abuse, privacy violations, and non-consensual AI-generated content.
- **Assess risk in context:** A technology that benefits one user can endanger another, making lived experience, power relationships, and conflicting stakeholder goals essential to security analysis.
- **Close the deployment gap:** Many serious failures occur not because the field lacks effective technical building blocks, but because smaller organizations and critical-infrastructure providers lack the resources, expertise, and usable tools needed to deploy them.
- **Teach security as a mindset and a public responsibility:** Adversarial thinking should appear early and throughout computing education — and reach students outside computer science who will increasingly build or govern technology with AI-enabled tools.

Summary

Susan Landau structured the session around a series of questions: What counts as cybersecurity? Whose risks matter? How should AI change the field? And what responsibilities do researchers, educators, industry, and government bear? The discussion widened the frame from securing systems to protecting the people and institutions that depend on them.

From Secure Systems to Secure People

Durumeric described cybersecurity as protecting people from harm, including threats outside traditional categories of intrusion or malware. Disinformation and non-consensual AI-generated sexual imagery still involve adversaries, targets, and technology-mediated attacks. Kohno framed cybersecurity both as a condition and as a worldview: approach technology with healthy skepticism, ask what can go wrong, mitigate foreseeable harms, and continue evaluating systems after deployment. Lopresti emphasized the practical end state — a system people can use without fearing attack.

That people-centered framing makes risk contextual. Kohno cautioned that there is no generic "individual" whose exposure can be assessed in isolation. Continuous location sharing may help members of a healthy family, yet create grave danger in an abusive relationship; its implications may change again when reproductive-health travel crosses state lines. Security research must account for lived experiences, unequal power, and legitimate goals that may conflict.

The "public" in public-good cybersecurity calls for a population-level view. Lopresti drew on the CCC paper *Lessons for Cybersecurity from the American Public Health System*, pointing to public health's attention to populations, threat vectors, measurable outcomes, prevention, and coordinated response. Durumeric shifted the focus from proving that a system could be secure to producing better outcomes for people, organizations, and critical infrastructure. Kohno connected this scope to technologists' moral responsibility to understand and mitigate harm.

The Last Mile Is the Main Problem

A recurring theme was the distance between strong technical ideas and dependable protection in practice. Water utilities, healthcare providers, and other essential organizations often operate with limited cybersecurity staffing and resources. Large technology and financial firms can invest heavily in protection, while smaller organizations struggle to integrate and maintain even well-established tools. Durumeric described this as a divide between the best-resourced one percent and the other 99 percent: substantial innovation exists, but it does not reliably reach most organizations. Work on deployment, usability, maintenance, and organizational fit may appear less novel than inventing a new primitive, yet it is often where security breaks down.

The panelists urged researchers to involve stakeholders early and understand not only what they request, but why. Organizations operate under different constraints, and elegant solutions can fail when designers misunderstand their setting. Engagement with agencies such as the Cybersecurity and Infrastructure Security Agency can reveal why problems that appear solved in principle remain hard in practice.

AI increases the urgency of this implementation gap without erasing cybersecurity's core principles. An audience question about Mythos, a new Anthropic LLM, sharpened the contrast between finding vulnerabilities and building secure systems. Durumeric observed that an AI system can be highly useful for discovering flaws without being able to produce dependable fixes. An attacker may consider five successes in ten attempts a strong result; defenders cannot safely deploy a patch that works only half the time. The panel largely agreed that AI lets adversaries find vulnerabilities, craft phishing messages, and scale attacks more quickly while lowering the expertise required. Defensive uses remain promising, but are less mature than attackers' ability to use AI at scale.

Education as Civic Infrastructure

The educational message was direct: security should not be "bolted on" after students have spent years learning to build systems. Students should learn from the beginning to ask what happens when inputs are malformed, systems fail, or malicious actors gain access. Because AI tools will enable people across many disciplines to produce software, this mindset must reach

beyond computer science to students who will develop, procure, regulate, and manage technology.

An audience contribution about Spelman College's cybersecurity clinic provided one model. The clinic brings together students from computing and other disciplines to help under-resourced organizations assess risks and develop security plans. This approach treats cybersecurity not simply as secure software development, but as a service activity responsive to an organization's particular needs. It also surfaced open questions about whether labor-intensive clinics can scale, how such work can be integrated across curricula, and how universities can steer more graduates toward critical-infrastructure organizations that cannot match large technology companies' salaries.

The discussion closed without drawing a hard boundary between "cybersecurity for the public good" and "computing for the public good." Instead, security emerged as an integral quality of responsible computing — one that shapes which problems researchers choose, whose experiences inform the design, and whether technical advances ultimately reach the people and institutions that need them most.

Resources

- Computing Community Consortium, Lessons for Cybersecurity from the American Public Health System
- Spelman College, Spelman SPEAR Cybersecurity Clinic

Boundless Frontiers: Co-Creating Global Research Futures

Thursday, July 23, 2026 | 4:00 pm–5:30 pm CT

Moderator

- **Katie Siek,** Indiana University

Speakers

- **Joanna McGrenere,** University of British Columbia
- **Aaron Quigley,** Australian National University
- **Ellen Zegura,** National Science Foundation and Georgia Tech

Key Takeaways

- **The policy tone has shifted:** U.S. science policy once championed the international flow of scientific information. Recent policy documents mention international work mainly in terms of cost sharing, competition, and lost talent.

- **Barriers are real and growing:** Research security rules, visa delays, shrinking funding, and international students' worries about coming to the U.S. all make cross-border research harder.
- **Collaborations start with people:** Lasting partnerships often begin with visiting graduate students, conferences, and small, intensive workshops such as Dagstuhl seminars.
- **Institutions can reduce the friction:** Seed programs, targeted funding, sabbatical support, and rules that are harmonized across countries make international work easier to start and sustain.
- **NSF is moving from ad hoc to strategic:** Joint calls with other countries' science agencies and joint workshops remain key tools, and NSF is now taking a more deliberate approach to its international partnerships.

Summary

Katie Siek opened with an audience poll asking for the first word that comes to mind about international collaboration in computing research. Answers ranged from languages and travel to geopolitics, along with benefits such as excitement and cultural exchange. She then traced how U.S. science policy has framed international work. In *Science: The Endless Frontier* (1945), Vannevar Bush wrote that "the Government should take an active role in promoting the international flow of scientific information." The National Academies' 2020 follow-up, *The Endless Frontier: The Next 75 Years in Science*, called large, international, data-sharing teams central to modern science, while noting that training, incentives, and institutions had not kept pace. By contrast, Siek pointed out that the White House Office of Science and Technology Policy's 2026 report, *Science: A New Golden Age*, uses "international" only once in its main text, in the context of cost sharing, and uses "global" mostly alongside competition and leadership. The report also worries that the U.S. trains global talent at great expense only to lose it to foreign governments. Against that backdrop, the session set out to identify barriers, mechanisms, and levers for international collaboration, collect practical institutional examples, and gather input for a best-practices memo.

The panelists brought three national perspectives. Aaron Quigley, dean of the ANU College of Systems and Society, described the Australian National University, founded after World War II, and its range of programs. Aaron Quigley formerly served as the Science Director and Deputy Director of CSIRO's Data61 in Sydney Australia where he expanded Australia's Quantum Next Generation Graduates program and established the Responsible AI Research Centre. Joanna McGrenere, an HCI researcher and co-head of computer science at the University of British Columbia, leads a department of about 70 tenure-track faculty, 3,000 undergraduates, and 300 graduate students. Her five-year Inria & Université Paris-Saclay International Research Chair gave her a firsthand look at how the European and Canadian systems differ. Ellen Zegura described NSF's international efforts, noting that science diplomacy is a key reason the agency invests in them and that the landscape is complicated right now.

Barriers: Security, Mobility, and Funding

An audience poll on the biggest barriers surfaced research security, international students' concerns about the risks of bringing themselves and their families to the U.S., researcher mobility, visas, and graduate student funding. Zegura said funding for international work has declined recently. Quigley argued that institutions that truly value international researchers will find ways to build programs for them. As dean, he has funded international graduate students, and he urged leaders to get creative with incentives even when budgets are tight.

Research security came up repeatedly. Zegura noted that the U.S. government maintains a list of foreign entities of concern, and McGrenere said Canada has a similar list. Quigley pointed to the AUKUS partnership and raised concerns about universities collaborating with institutions tied to China's military, known as the "Seven Sons of National Defense," calling that the kind of collaboration institutions want to avoid.

How Collaborations Start and Last

Asked how strong collaborations get started beyond funding, the panelists pointed to people and places. Zegura said her own major international collaborations grew out of graduate students visiting Georgia Tech, including two students from France. Quigley described pitching to host a ubiquitous computing conference in Dublin because researchers in mainland Europe did not know his group well; the effort connected him with a wide range of European funding programs. When Siek asked how to build the case for international collaboration, Quigley advised faculty to do their homework.

Zegura acknowledged that NSF's international strategy has historically been ad hoc, much like faculty-initiated efforts at universities, but said the agency is now thinking more strategically. Its main tools include joint calls with other countries' science foundations, where NSF has typically served as lead reviewer, including in collaborations with India, and joint workshops, such as virtual workshops with the Japan Science and Technology Agency (JST). She was careful to add that this does not mean NSF is promoting virtual collaboration as a substitute.

Incentives and Institutional Levers

On incentives, Quigley encouraged leadership to incentivize international collaboration. Indeed, he described how he provided faculty with funds to collaborate with international colleagues and addressed barriers faculty were experiencing to make international collaborations more feasible. He also described a modest program with India that he supported as dean, which drew many applicants and visitors and helped put his university on partners' radar. He noted that international students who move between countries become natural bridges between nations and institutions. Asked whether people outside academic institutions in the U.S. can apply for NSF funding, Zegura said some programs accept proposals from nonprofit organizations. On maintaining ethics and access, McGrenere described Canada's harmonized approach but called the problem largely unsolved, with the practical goal of reducing friction. Quigley observed that researchers rarely take the sticky questions to ethics committees such as IRBs.

Quigley and several audience members praised Dagstuhl seminars as isolated but welcoming settings where people talk, share ideas, and build new connections. Quigley offered to help colleagues propose Dagstuhl workshops, noted that Japan has similar venues, and highlighted the generous sabbatical support available in Japan and Korea.

Bold Ideas and Commitments

Asked what they would change with a magic wand, McGrenere called for harmonization across countries, since international collaboration keeps getting stickier. Quigley named visas: the process is so slow that only large technology companies can afford to wait. He also noted that remote work, widely allowed during COVID, is no longer supported at many institutions. For their own commitments, Quigley pledged to broaden the scope of funding opportunities, citing direct funding from Korea, and Siek pointed to a CRA program (DSP).

Conclusion

The panel described international collaboration as under pressure but far from over. Policy is shifting toward competition and security, while visas, funding cuts, and compliance rules add friction. Yet the panelists' own experiences showed that durable partnerships still grow from visiting students, conferences, workshops, and small, well-targeted programs. Their ideas, from harmonized rules and faster visas to strategic funding and support for researcher mobility, will feed into the session's best-practices memo.

Resources

- Vannevar Bush, Science: The Endless Frontier (75th anniversary edition, NSF)
- National Academies, The Endless Frontier: The Next 75 Years in Science (2020)
- White House OSTP, Science: A New Golden Age (2026)

Track 3: Novel Computing Directions

CS Enrollment Trends and the Reality of Jobs for CS Graduates

Wednesday, July 22, 2026 | 11:00 am–12:30 pm CT

Moderator

- **Vivek Sarkar,** Georgia Institute of Technology

Panelists

- **Nancy Amato**, University of Illinois Urbana-Champaign
- **Mary Hall,** University of Utah

- **Ran Libeskind-Hadas,** Claremont McKenna College
- **Ben Zorn,** Microsoft Research

Key Takeaways

- **Read the decline carefully:** Enrollment pressure is real but uneven. Leaders should distinguish a correction after extraordinary growth from changes in student interest, international mobility, and perceptions of the job market.
- **Broaden the map of computing careers:** A CS degree does not lead only to a job in the technology sector. Finance, healthcare, manufacturing, government, entertainment, and other domains increasingly need graduates who can apply computational thinking to their core problems.
- **Preserve fundamentals for an AI-mediated workplace:** Algorithms, systems, specification, testing, and problem decomposition become more — not less — important when graduates must direct and validate AI-generated work.
- **Build multiple pathways into computing:** Interdisciplinary degrees, certificates, bridge programs, and industry-linked learning can reach new populations and support lifelong reskilling without diluting a strong CS degree.
- **Turn uncertainty into an evidence agenda:** Departments need disaggregated data on placement, international students, program demand, and employers outside the technology sector — not simple predictions about whether AI will create or eliminate jobs.

Summary

Vivek Sarkar opened the Novel Computing Directions track by recalling an earlier CRA concern: how departments could respond when CS enrollment seemed to double every few years. The current moment presents the opposite anxiety. Applications are falling at some institutions, entry-level hiring has tightened, and headlines often equate a CS degree with coding for a technology company. Sarkar argued that this obscures both the breadth of computational thinking and the growing need for computing talent across the economy. He organized the discussion around enrollment, interdisciplinary education, job readiness, and other opportunities to address enrollment and job placement challenges.

An Uneven Enrollment Correction

The institutional evidence resisted a single national story. Mary Hall reported Utah's first decline after a decade of growth, but the pattern varied sharply: CS enrollment fell 12 percent, computer engineering grew 7 percent, and software development and data science held comparatively steady. Overall master's enrollment fell 25 percent, closely tied to recent headwinds in international enrollment, while an interdisciplinary master's program centered partly on AI grew 26 percent. Nancy Amato described a different picture at Illinois. Applications had softened, but the centrally managed program expected its largest incoming class even as its acceptance rate rose from roughly 6 to 7 percent to 10 percent. She saw a possible

opportunity to widen access after years of restricting entry.

Audience questions exposed the stakes for departments whose budgets depend on enrollment, especially tuition from master's students. Suggested responses included cross-campus partnerships, combined bachelor's/master's pathways, and outreach to new populations. Yet the ethical question of recruiting international students into an uncertain market remained unresolved; answering it requires placement data separated by program, degree level, and student status.

Computing Across Disciplines

Ran Libeskind-Hadas presented Claremont McKenna's integrated sciences model as a proof of concept: computing and domain faculty co-teach, embed computation in disciplinary courses and research, and build collaborations across silos. Amato described Illinois' CS+X and X+Data Science programs, while Hall highlighted Utah's AI, data science, games, robotics, human-centered computing, and certificate pathways. A late audience exchange clarified their purpose. Traditional CS graduates can already succeed in non-technology industries; blended degrees broaden who enters computing and what problems they are prepared to address, rather than replacing the CS degree.

Job Readiness Beyond the First Job

The panel resisted both complacency and predictions of a job apocalypse. Hall's placement data showed weaker outcomes but also opportunities in finance, fintech, engineering, and government. Libeskind-Hadas similarly reported that economics students with computational sophistication had an advantage in a newly difficult market.

Zorn argued that CS education's enduring value lies in teaching people to specify, decompose, test, and validate. His phrase "prompts are programs" captured the shift: anyone delegating work to AI must understand the request and judge the result. As teams coordinate multiple agents, these habits become relevant across knowledge work; for computing graduates, reliability, security, privacy, and evaluation become central professional skills.

Asked about layoffs and vanishing entry points, panelists cited pandemic-era overhiring and wider economic uncertainty while acknowledging that graduates now face fewer choices. Their prescription was to prepare students for durable careers: preserve foundations, teach responsible AI use, and make careers beyond technology firms visible. Business-school partnerships and entrepreneurial capstones can also help students explain how technology creates value.

Other Opportunities for Enrollment and Job Placement

The final discussion widened to lifelong learning. Amato highlighted Illinois' iCAN bridge program for adults without a computing degree; panelists also discussed certificates, employer-supported study, and opportunities for graduates whose skills need renewal. Access to AI tools remains unresolved. Zorn argued that improving open-weight models may reduce dependence on paid tokens, while Sarkar dryly noted that GPUs are not free either. Accreditation and shortened degrees raised the same tension: create flexible routes without

weakening foundations or disadvantaging students with fewer prior credits.

The session ended with an agenda rather than a forecast: gather longitudinal, disaggregated enrollment and placement evidence; develop sustainable responses to volatile master's demand; broaden access to AI infrastructure; and study which interdisciplinary, bridge, and accelerated pathways work for which students. Sarkar invited continued CRA activity around these questions. The closing message was cautiously optimistic: computing remains broadly valuable, but departments must explain that value more clearly and redesign pathways around the world students are entering.

Research in AI and the Future of Arts and Entertainment

Thursday, July 23, 2026 | 11:00 am–12:30 pm CT

Moderator

- **Jeff Burke**, University of California Los Angeles

Panelists

- **Brian Magerko**, Georgia Institute of Technology
- **Moshe Mahler**, Carnegie Mellon University
- **Sarah Gimbert Sterman**, University of Illinois Urbana-Champaign

Key Takeaways

- **To support creativity in art and entertainment, AI should support processes rather than outcomes:** Current models are often trained to make "better" artistic outputs with a single prompt. However, in practice, art results from a creative process where the artist iteratively manipulates and refines a medium using tools. Future research should focus more on how AI can be used to build tools that support this process, rather than how to directly generate outcomes.
- **Creativity has many subcomponents that are often overlooked:** Instead of talking about creativity as a whole, it is helpful to consider the sub-activities such as idea generation,

Summary

This session involved a panel discussion of how research in AI could impact the future of arts and entertainment. In particular, the discussion focused on creativity, with panelists considering how generative AI can both enhance and suppress human creativity across various disciplines. In general, the panelists agreed that the interaction between AI, art, and creativity is complex and that future research should focus on supporting the creative process rather than producing a final product.

The session started with a short five-minute lightning talk from each of the panelists followed by a question and answer session including both moderator-seeded and audience questions. Here, we summarize each of the three lightning talks in the order they were presented in the session, as well as each question asked to the panelists.

Lightning Talk — Brian Magerko

In this talk, Brian Magerko discussed his nuanced perspective on the concerns and implications of AI in creative spaces, grounded in his 30 years of experience on AI and creativity research combined with his experience as the director of graduate studies of digital media at Georgia Institute of Technology. He emphasized the significance of how recent generative models have shifted the conversation around and capabilities of using AI in creative spaces; for most of his career, it never felt like AI and Art would get this effortless this fast. This means that many potential concerns stemming from AI use in art and entertainment, such as art diversity, artist livelihood, intellectual property theft, and livelihood of artists and cognitive impacts, are emerging and just starting to be looked at seriously by the research community. Magerko ended by talking about his own work, LuminAI, where AI is used to power an artificial dancer that can improvise in real time with humans. This work involved working with dancers, involved small models, and resulted in an augmentation — not replacement — of human creativity. Magerko argued that supporting values such as these in research are important for the future of AI and Art.

Lightning Talk — Moshe Mahler

An Associate Teaching Professor at Carnegie Mellon University's digital media group, Mahler also spent ten years at Disney research, and he is also a filmmaker and studio owner. In his talk, Mahler described his experience both as a researcher and a user of artistic creation tools. He noted that the AI industry really focuses on artistic tools that create efficiency or quality gains. However, he'd prefer that future tools focus on artist-driven design principles. Mahler emphasized that advances in AI present an opportunity to rethink how artists make content, but that this is different from the democratization of content creation. We as researchers and AI practitioners have done a good job in lowering the barrier for entry into various artistic entertainment tasks. However, it is a more open question what can we do to raise the ceiling for the masters of craft. Mahler argues that building AI-powered tools for expert artisans is an important direction for research and investment in the future.

Lightning Talk — Sarah Gimbert Sterman

This lightning talk examined how computer scientists can support creative growth across different settings. Sterman, an Assistant Professor at the University of Illinois Urbana-Champaign, summarized some of her recent research, arguing for the importance of designing AI systems that support the creative process rather than those that are tailored solely to producing final outputs. This includes the state of creative struggle; with AI in writing, for example, much of the current focus is on getting it right and fluent the first time. However, in practice the goal isn't fluent text; it's to express your ideas, and ideas are typically not fully thought through up front. Instead, ideas get fleshed out through the writing process, a process that can be lost with current AI tools. As a result, there is a need for systems that support this

process more directly.

Question and Answer

Question: What is the research you wish people would fund for supporting new human creativity that is currently being missed?

Magerko and Mahler agreed that current computers struggle to truly improvise or collaborate with humans. Instead they are subservient. Understanding collaboration and improvisation computationally could potentially help support human creativity. Sterman added that another interesting direction to explore would be researching how to support artists, including those without programming backgrounds, in making their own digital tools. When an artist makes something with physical media, they are also often making their own tools (e.g., violin makers make their own jigs).

Question: How important is it that we understand and mitigate the effects of AI on human individual, social, and collective creativity?

All panelists agreed that this topic has important educational and long-term implications. Magerko brought up recent research showing that when one group of writers can use AI and another cannot, those without AI demonstrate greater creativity, with effects that can persist over time. The panelists agreed that this is a direction needing additional research and support.

Question: How do we teach creativity and computing together in the current landscape? What insights from your research on understanding creative processes transfer to teaching and learning, and how can we use those insights to help students become more creative in their own disciplines?

The panelists again emphasized the importance of process over outcomes. Many of these insights predate generative AI. For instance, there is a long history showing the importance of starting by sketching by hand before moving to digital art. Panelists also noted that it is important to fund technical and artistic joint projects where excellence in both fields is valued. This can be challenging given current funding structures. Each discipline must be willing to support excellence in the other. Mahler noted that a lot of his students come from an engineering background, and feel like they are not creative people. Creativity is a muscle that you have to exercise.

Question: What would it take for an AI agent to be genuinely capable of creative improvisation?

The panelists agreed that the agent would need all of the social context, including both the larger social moment and at the individual level. True improvisation and artistic creation by AI may be impossible; art may be inherently human, as it is about communicating yourself.

Question: How might AI tools foster human creativity rather than leading to skill atrophy? And what are the obstacles to making AI more of a collaborative partner in the creative process?

The panelists noted the importance of supporting activities people cannot do rather than replacing challenging activities they can do. In addition, the discussion should move beyond the word “create” to consider how and where AI is integrated into the process. Mahler noted

that partial automation is preferable to full automation. This dates back to before AI. It's a balance of how much the tool should do. Sterman also noted that she objects to describing AI as a collaborator; collaboration is relational and AI can mimic that, but it can't really engage in that sort of thing. Instead, she prefers thinking about it as a tool.

Question: Does technology need creativity at all? Isn't efficiency the whole point of technology?

The panelists agreed that many everyday uses of technology are not primarily about efficiency. We use technology for being creative and consuming the results of other people being creative all the time. The point of technology is to be transformative — efficiency can be a driver of that transformation, but it is not the only thing. Writing, for example, made people more efficient, but it also allowed them to express deeper ideas.

Yes, And: Computing Visioning Improv

Thursday, July 23, 2026 | 2:00 pm–3:30 pm CT

Organizers and Moderators

- The executive board of the CCC (CRA's Computing Community Consortium) including **William Gropp** (University of Illinois Urbana-Champaign), **David Jensen** (University of Massachusetts Amherst), **Michela Taufer** (University of Tennessee, Knoxville), and **Kevin Butler** (University of Florida)

Summary

The *Computing Community Consortium (CCC)* is a CRA committee which leads visioning activities to bring the computing community together and help shape the long-term research directions. The goal is to help build consensus and to develop resources to help push those directions forward. So far, the CCC has helped to organize 93 workshops with more than 4,000 attendees. From these workshops, the CCC produces strategic plans and reports for the community at large.

In this session, the executive board of the CCC led one such visioning session. To do so, the organizers introduced the audience to the original game of *speculative execution.* Designed to spur brainstorming about visions for the future of computing, this activity involved the audience splitting into groups of two to five. Then, each group was handed a couple of cards with candidate research domains and potential technological approaches. Then groups were tasked to brainstorm and pitch a potential vision involving the provided cards to the room. Groups were asked to imagine a research advance that can improve the given application by 10x in capability, efficiency, or cost. At the end of the session, the audience voted on the pitches that they liked the best.

Overall, the game was well received by attendees. During the brainstorming, the room was full of chatter with each group considering their specific scenario and planning their presentation. In addition, pitches were full of jokes that spurred audience laughter and engagement. In this summary, we briefly list some of the ideas presented in the session. However, overall,

speculative execution was an enjoyable (if slightly non-traditional) conference session experience!

Presented Visions

Real-time Epidemiological Observation. Several groups described approaches to improving disease monitoring and response. For instance, one group proposed a cloud-based, distributed, epidemiology observation center for the spread of these diseases which uses an AI model to protect privacy, and to adopt solutions for the individual community, so it can detect outbreaks early and prevent diseases from spreading. Similarly, another group proposed building personal, wearable sensors on your wrist that can share your data in a dashboard format to determine where there may be disease hotspots, both for general disease monitoring as well as for informing individual travel decisions.

Robo-Pothole Filler: One group envisioned a future with automated robotic construction infrastructure which can help patch and fix issues such as potholes or sidewalk cracks efficiently and effectively.

Automatic Language Translation: One group proposed leveraging existing technology in cochlear implants combined with AI models to convert sounds in one language that you don't understand into one that you do.

Environmental Educational Mixed Reality: One group proposed combining low-powered computing with wearables to build educational interactive experiences to augment learning for elementary children when interacting with their environments.

Executable Email: One group proposed that email should just become data, which can be executed or automatically routed to the people who actually need to see the email and can act on it.

The Pivotal Role of Computing Research for the Future of Healthcare

Thursday, July 23, 2026 | 4:00 pm–5:30 pm CT

Moderator

- **Raymond Ng,** University of British Columbia

Speakers

- **Fayika Farhat Nova,** Parkview Health
- **Tarek Haddad,** Medtronic
- **Martin Michalowski,** University of Minnesota

Key Takeaways

- **AI can expand healthcare capacity and enable more personalized care:** By organizing complex clinical knowledge, screening large volumes of medical data, reducing false alerts, and automating time-consuming review, AI can help scarce specialists focus on the patients and decisions that most need their expertise.
- **Healthcare AI must be evaluated by more than technical performance:** Accuracy and speed are necessary but insufficient. Systems must also produce clinically meaningful explanations, support rather than disrupt human decision making, fit local workflows, and improve outcomes that matter to clinicians and patients.
- **Deployment is itself a computing research challenge:** Generative AI is harder to validate than traditional, narrowly defined models. Privacy, regulation, reimbursement, local validation, equitable access, co-design, and governance must be addressed as part of system development rather than after a model has been built.

Summary

Raymond Ng opened the panel by emphasizing that artificial intelligence will play an important role in healthcare, but that the contribution of computing research extends beyond AI alone. The speakers approached this opportunity through decision support, medical devices, human-computer interaction, and natural-language processing. Across these perspectives, they shared a common message: *healthcare does not simply need more powerful models. It needs computing systems designed around clinical decisions, human workflows, patient needs, and the realities of deployment.*

From General Guidelines to Patient-Centered Decisions

Martin Michalowski introduced multimorbidity as an example of a problem that requires computing research. Patients increasingly live with several chronic conditions, yet clinical guidelines are generally written for one disease at a time. Applying several guidelines to the same person can produce conflicting or impractical recommendations. Computing methods can help combine this knowledge and identify a plan tailored to the patient's priorities, including cost, convenience, and the likelihood that the patient can follow it.

Michalowski then turned to explainability. As decision-support systems become more capable, their recommendations can also become harder to understand. Existing tools may show how a model generally reasons or which features influenced a prediction, but that does not necessarily tell a physician why a recommendation makes clinical sense for this particular patient.

The risk is not merely that an explanation will be incomplete, but that it may be confidently wrong. Michalowski described examples in which an AI system mischaracterized a drug's mechanism or referred to a real clinical trial while supplying fabricated numerical results. He argued that computing researchers must develop ways to verify machine-generated claims and determine whether explanations actually help clinicians make decisions. His group presented automated measures that reportedly aligned with physician judgments in more than 80 percent

of cases.

Extending Specialty Care

Tarek Haddad framed medical-device AI against a widening gap between the supply of specialists and demand for care. 40 percent of U.S. electrophysiologists are expected to leave active practice within five years, while the worldwide population aged 65 and older is projected to grow by more than 300 million over the next decade. AI cannot replace this workforce, but it can automate repetitive work and help clinicians direct their attention more effectively.

Haddad highlighted Medtronic's AccuRhythm™ AI as one example. Its heart-arrhythmia algorithms reportedly reduced false device alerts by 90 percent while retaining 99 percent of true alerts. This illustrates how a carefully bounded AI system can reduce workload without sacrificing sensitivity.

Generative AI could go further by assisting with documentation, triage, workflow management, decision support, and chronic-disease management. It also creates a harder safety problem. Traditional medical-device AI often performs a narrowly defined task with a limited set of possible outputs. Generative systems accept and produce open-ended language, making their behavior much more difficult to test comprehensively. Haddad identified review agents, trusted medical knowledge, and domain-specific training as promising ways to make these systems more reliable, while emphasizing that new validation methods will be essential.

Optimizing Care, Not Just Tasks

Fayika Farhat Nova argued that task optimization is not the same as care optimization. Healthcare technologies may predict risk, generate documentation, recommend actions, or automate individual steps. Care, however, is delivered through a changing system of people, handoffs, policies, exceptions, and patient circumstances. Improving one task can create additional work or risk elsewhere.

Human-computer interaction therefore changes the questions researchers ask. Instead of focusing only on whether a system can be built, accelerated, or made more accurate, researchers must also ask what problem people are solving, how work actually happens, whether clinicians and patients understand and trust the system, whether it fits the workflow, and what should not be automated.

Nova distinguished technical measures, such as accuracy, speed, scalability, and reliability, from human measures, such as usability, trust, workflow fit, safety, cognitive burden, and adoption. Both are needed to understand healthcare impact. As she put it, "A technically excellent system has little value if people cannot or will not use it." She also highlighted what current systems often fail to represent: professional judgment, local practices and exceptions, patient goals, social context, and practical constraints. Making these factors visible to a computing system is a central human-centered AI challenge.

Making Clinical Text Computable

Ng focused on the large amount of healthcare knowledge contained in unstructured clinical

documents. Machine learning in healthcare still relies heavily on structured data, even though reports and notes contain information that could improve patient care and operational efficiency. Extracting that information is difficult because reports are not standardized: clinicians use different terminology, abbreviations, and forms of negation to describe similar cases.

Ng presented an analysis on how an automated extraction pipeline could replace manual review of approximately 120,000 reports per year, saving more than 1,000 person hours, and allow processing of an individual report approximately 600 times faster than manual review. The scale of the example showed how NLP can redirect highly specialized staff from repetitive review toward higher-value work.

He also argued for healthcare-specific language models. Sending patient information to an externally hosted model can create privacy and governance concerns. A locally controlled model can be deployed privately and adapted to clinical vocabulary, abbreviations, and context. The goal is not simply to apply a general chatbot to healthcare, but to build language technologies around the domain's privacy, reliability, and knowledge requirements.

Deployment, Access, and Governance

Audience questions highlighted the gap between a promising model and a sustainable healthcare system. Haddad noted that regulators have substantial experience with traditional medical-device AI, but that appropriate validation pathways for generative AI remain unsettled. Approval alone is not enough: hospitals may not adopt useful systems if their use is not reimbursed. Nova similarly warned that mental-health applications marketed in therapeutic terms can be widely available without adequate validation, making patient education and safeguards essential.

The discussion of false-alarm fatigue showed why accuracy cannot be considered in isolation. AI paired with expert review may reduce workload and improve performance, while immediate access to a somewhat less accurate tool may be preferable to waiting weeks for a specialist. Cost and geography further shape that choice. Michalowski pointed to the potential value of small models that can operate on basic phones in rural communities, where the alternative may be little or no timely access to care.

The panelists also cautioned against developing a technical solution and only afterward searching for a clinical problem. Michalowski raised concerns about systems being deployed without sufficient validation in the environments where they will operate. Ng described co-design with clinicians from the beginning, including revising technical goals as researchers learned the organization's language, workflow, and pain points. Nova's remarks explained why this matters: local practices and exceptions are part of the system requirements, even when they are absent from the initial specification.

Two audience questions remained open. How can AI improve rural healthcare without widening existing disparities? And when a computing system consistently outperforms unaided human judgment, as audience members suggested may already occur in areas such as mammography and drug-interaction checking, when should its use become part of the expected standard of care? These questions connect technical evidence to access,

accountability, professional responsibility, and policy.

The panel concluded with cautious optimism. AI can help healthcare professionals manage growing demand, reduce avoidable work, and make better use of clinical information. Realizing that promise, however, will require computing researchers to look beyond model performance and design for clinical usefulness, human trust, local conditions, and shared governance.

Track 4: Computing and the Now

Navigating the "Mess": Strategic Resilience for Computing Research Organizations

Wednesday, July 22, 2026 | 11:00 am–12:30 pm CT | Presentation slides

Organizers

- **James Allan**, Chair, CRA Board of Directors; University of Massachusetts Amherst
- **Yannis Ioannidis,** President, Association for Computing Machinery (ACM); University of Athens

Panelists

- **William Gropp**, University of Illinois, IEEE-CS
- **Stephen Smith**, Carnegie Mellon University, AAAI
- **Carol Woodward**, Lawrence Livermore National Labs, SIAM

Key Takeaways

- **"Mess" Aspect 1: AI is Outpacing our Structures:** Faculty have insufficient time or support to learn AI, curriculums have no room left over for AI content, authorship rules don't cover AI-assisted research, reviewers are overwhelmed by AI-generated submissions, and universities can't compete without industry-scale compute.
- **"Mess" Aspect 2: The Narrative Around Enrollment and the Workforce:** CS enrollments are down, and still falling; this may be in part because the public thinks CS is "just coding" and academic units lack real-time data to push back. In addition, visa rules and anti-tenure bills are squeezing programs. Simultaneously, faculty are buried under rising teaching loads.
- **"Mess" Aspect 3: Funding and Policy Volatility:** Given the unpredictability of the federal funding landscape, institutions can't rely on NSF funding staying stable. Overall, federal research dollars are decreasing and award accept ratios are collapsing. In some states, units are also fighting harmful state legislation alone, and funding gaps are now derailing tenure cases.

- **A Coordinated Response is Needed:** There is an opportunity for professional societies to work together to help researchers and institutions navigate all three aspects of the AI-induced shifting landscape. Individual organizations are already undertaking various initiatives, and cross-organizational initiatives such as the CRA's *Community Led Ethics for Accountability in Research* (CLEAR) are already bringing organizations together. There is a need for more such collaborations.

Summary

This panel included representatives from several computing professional societies and organizations. The goal of this panel was for individual societies to share what they have been doing to attack a shifting research landscape full of funding and policy volatility where AI is outpacing organizational structures and computing enrollments are shrinking. By starting this dialogue, societies may both be able to better support their members and also potentially collaborate on future initiatives. Ioannidis and Allan started the panel by summarizing the goal and discussing the current "mess" (as described in the key takeaways above). Then, each organization leader talked for five minutes, describing relevant societal activities and responses to each of the three themes listed above. Talks are listed in the order they occurred during the session.

William Gropp, IEEE-CS: IEEE-CS is frequently encountering AI's impact throughout the publishing process. This includes instances of fraud and cheating, but also higher-level conversations about whether and when using AI during the publication process is ethical and when it is not. In addition, Gropp noted that writing a paper has become easier in some respects, contributing to a surge in conference submissions, making reviewing challenging. To work toward solutions, IEEE-CS has been collaborating with the CRA to have formal discussions, including producing "Beyond Code," a report about AI and the future of Code. IEEE-CS is also trying to encourage societal engagement and membership through conference-based short-term mentorship initiatives. Gropp concluded by underscoring that while AI and the shifting landscape poses a lot of challenges, there are also a lot of opportunities. In particular, there is a real opportunity for the different professional societies to work together to try different policy approaches and advocate for the computing research community.

Steve Smith, AAAI: AAAI also has initiatives targeting many of the themes discussed in the opening summary. Smith underscored that AAAI is very much an advocate for talking to Congress and helping communicate the importance of basic research; right now, the focus is on industry where companies are making significant advances, but this is only possible due to basic research that began in the 1980s. AAAI is committed to supporting research that forms the foundation of the advances of the future. Regarding AI and the publication process, AAAI is facing a scaling problem related to the number of papers submitted to their annual conference. In three years (2024–2026), AAAI has gone from 8,000 to 30,000 submissions (25,000 post desk rejects). This necessitates 75,000 reviews. To address this problem, AAAI experimented with using generative AI during the review process by adding a fourth review generated by ChatGPT. They had two rounds of reviews. The first involved two humans and ChatGPT. If any of the three voted yes, the paper went on to a second round with a third human reviewer. They also used AI to summarize reviews for area chairs, which Smith noted was probably the biggest

positive impact of their approach. Finally, they used AI to recognize papers that were trying to beat the system, finding multiple instances of the same paper being submitted with different sets of authors. Overall, the experiments were a success and a few other conferences are using the same review pipeline this year.

Carol Woodward, SIAM: SIAM is a professional society for applied mathematics that represents computational researchers. Woodward underscored that AI represents a major shift in research practices — more than any time before, this is a time to lean on each other, and the professional societies can play an important role. For instance, SIAM recently adopted a publication policy centered around the use of AI. This new policy focuses on when a paper gets submitted; if you are an author, you are asserting the integrity of the work. So even if AI is used, the author still owns the work and mistakes rest with them. While not as extreme as reported by AAAI, SIAM is also seeing more submissions and is experimenting with how to integrate AI while protecting the intellectual property of the authors. Woodward also notes that SIAM has a committee on public policy where they send a group of 25 people to capital hill each year to explain how important applied math research is to the development of AI, and underscoring the importance of basic research. SIAM also participates in CLEAR and other initiatives for coordinating between professional societies. Professional societies are really where we are seeing the voice of research coming together as decisions are being made.

James Allan, CRA: The CRA has several initiatives aimed at helping navigate the mess. This includes initiatives to gather data, get people together, and generate reports to share findings and discussions with the broader computing community. CRA does not run conferences; it convenes summits. Thus, the CRA is not faced with the same crises as described by AAAI or SIAM, but its member organizations are. CRA is trying to help by organizing National AI Research Resources (NAIRR) events to help faculty learn about the resources available to them. There are also many reports on relevant topics coming out in the near future. The CRA is also an advocacy organization, advocating for computing research with the federal government to try to keep funding levels high. For instance, the CRA has a government affairs policy group that has a blog so that people who are interested can keep abreast of what is going on.

Yannis Ioannidis, ACM: The ACM agrees that there is a need for stakeholders to get together and discuss potential actions. To help with this, they will be hosting an AI leadership summit in fall 2026. Ioannidis notes that large technology companies are not hiring as much as they once did, but other industries are hiring computer scientists more. Thus, ACM is seeking conversations with CEOs of smaller, non-computing companies to understand what they want from computing students, discussions that may influence computing curricula. To empower research without needing prohibitively expensive computing infrastructure, the ACM is advocating for AI companies to make frontier models actually open source. The ACM also has a policy council who is trying to advocate to policymakers. As with other organizations, the ACM is grappling with publication policies. In the past 18 months, it has changed its policy on AI use three times, each time making it more accepting. Right now, if you use it for the research itself, you have to declare it. If you use it for the text, you do not. There are paradoxes we are dealing with regarding AI in research reviewing. Ioannidis notes that there are a lot of ethical challenges with AI in research, and that discussions are ongoing.

Questions and Answers:

Q: Given the discussion about paper reviews, is there a good way to publish guidelines on how we navigate the ever changing rules regarding AI and Research?

A: Smith acknowledges that it is a hard question, and asks if CRA is a place where such guidelines could be developed. Allan indicates that CRA is definitely interested in getting people together to talk about these sorts of things, but what people think is the right thing to do is evolving. There are a number of challenges, and further discussions are needed before we can determine what to do. Woodward agrees that she thinks it is a very fluid area, so clear guidelines are unlikely to emerge across disciplines for a while. She notes we are going to need reviewing processes that look at research significance which requires field knowledge. She is also concerned about provenance; AI can do the work, but it can't reference all of the things that it has drawn on to do this new work, and that is a significant problem. Gropp notes that this is the extension of the challenge of a system that rewards a least publishable unit. This system was a problem, but we were living with it. Now, it is overwhelming us. Ioannidis notes that there are already research results purely produced by AI. We have to change the way we evaluate research and researchers.

Q: To what extent are reputation-based approaches being considered to improve the integrity of CS research processes (e.g., some public consequence for authors that violate ethical standards around AI use). And what are legal issues associated with sharing information about unethical behavior among professional societies and employers?

A: Ioannidis notes that legally, sharing things between professional societies is different from sharing between employers. He wants to work toward a rule where everything you do for all organizations will be tracked in the same domain space. Then it can be used to look for suspicious behaviors and decide consequences. For instance, there are people who have been sanctioned at the ACM and cannot go to conferences for 30 years. However, notifying employers may not be possible. Smith shares that in their call to papers in 2026, AAAI included that if they detected certain unethical behaviors by authors, there would be some sort of punishment. This year, they just rejected the paper, but there is a need to evolve to more long-term penalties like the ACM. Figuring that out at AAAI is in process.

Q: Other fields do most of their reviewing work in journals. They are not subject to time pressures. Should we look at other field publishing models?

A: Smith indicates that AAAI has tried to encourage publishing in journals by having journal-first conference sessions. However, it is hard to break the habit in AI. Gropp states that IEEE-CS has also looked into prioritizing journal publications. However, the rate of progress in the field can make the journal publication process and timeline a challenge. That doesn't mean we shouldn't do it, but we do need to look at this as a community and rethink it. Ioannidis has been advocating that we should go back to journals. However, he has experienced resistance; for the change to occur, it has to be a community effort. Globally, funders, deans, and professors all have to agree. Woodward notes that applied math is mostly in journals, but even so, they are still seeing an uptake in submissions. Regardless of whether a discipline uses conferences or journals, she noted the culture around writing up the "least publishable unit" needs to change, and there needs to be more focus on the significance of a contribution.

Q: CS is not alone in dealing with AI-generated content in research papers and reviews. Are computing societies working with societies in other sciences to devise new structures and policies? What cross society activities have you started with each other?

A: Allan notes that the audience is experiencing one such activity right now! He also reiterates the CRA's CLEAR program. Smith notes that AAAI has started an initiative with INFORMED, a journal publishing organization. Tracy Camp mentioned CRA education-related work involving multiple societies. Allan notes that every February, there is a society leadership meeting.

Resources

- CLEAR taskforce
- Beyond Code: Engineering Trustworthy Software Systems with AI at Scale

Industry Expectations: What Industry Wants from Computing Students

Thursday, July 23, 2026 | 11:00 am–12:30 pm CT

Moderators

- **Jon Bell**, Northeastern University
- **Rebecca Hwa,** George Washington University

Panelists

- **Fayika Farhat Nova**, Parkview Health
- **Titus Winters**, Adobe
- **Kayla Woitkowski**, Nutanix
- **Gabriela Cruz Thompson,** Intel

Key Takeaways

- **AI fluency is becoming a baseline, but judgment is the differentiator:** Employers increasingly expect graduates to use AI tools, but they distinguish strong candidates through critical thinking, communication, validation, systems understanding, and the ability to connect technical work to business and stakeholder needs.
- **Fundamentals still matter, but their emphasis and assessment must change:** The panel disagreed over how much of the traditional CS curriculum should remain, but broadly supported shifting attention from syntax and routine implementation toward code understanding, testing, debugging, design, and evaluation of a student's problem-solving process. Technical interviews increasingly ask candidates to perform

code review and maintenance tasks using AI tools, moving away from LeetCode-style programming challenges.

- **Internships remain a critical talent and trust pipeline:** Companies continue to use internships to develop future full-time employees and bring AI-native knowledge into their organizations, but uncertainty about entry-level roles and declining trust in AI-assisted applications are making hiring more conservative and experiential.
- **Academia-industry alignment requires sustained, multi-party relationships:** Industry needs are dynamic and often driven by immediate projects rather than fixed job descriptions. Meaningful alignment requires ongoing engagement among faculty, career services, recruiters, business leaders, and the managers who hire and supervise graduates.

Summary

During this session Rebecca Hwa and Jon Bell of the CRA-Industry Practitioner-to-Professor Working Group organized a panel on how artificial intelligence is changing what industry expects from computing graduates. The discussion drew on perspectives from Titus Winters of Adobe, Gabriela Cruz Thompson of Intel, Fayika Farhat Nova of Parkview Health, and Kayla Woitkowski of Nutanix. The panel was framed around three questions: what skills graduates now need and whether traditional fundamentals retain their value; how companies recruit and evaluate interns when AI systems can perform many intern-like tasks; and how academia and industry can communicate more systematically about changing workforce needs. Across these topics, the panelists emphasized that AI has moved from an experimental technology into everyday workflows, while curricula, hiring practices, and assessment methods are adapting at different rates.

AI Fluency Is Becoming a Baseline, but Judgment Is the Differentiator

The panelists broadly agreed that familiarity with AI tools is becoming an expected part of a computing graduate's skillset, but tool use alone will not distinguish a candidate. Titus Winters argued that syntax and basic code production have little standalone value now that AI can perform much of that work. He instead emphasized code understanding, design, testing, validation, teamwork, communication, and ethics. Kayla Woitkowski similarly reported that intern hiring managers continue to prioritize critical thinking and the ability to operate under ambiguity, particularly because the clearly defined assignments and grading criteria of school do not resemble the shifting requirements of many businesses. Fayika Farhat Nova added that many students already have adequate AI skills expected at their levels; the differentiator is whether they question the output, understand the larger system, and can connect their technical choices to intended outcomes and stakeholders. Gabriela Cruz Thompson reinforced the need for problem solving, testing, debugging, listening, and understanding the business process that a technical system is intended to support.

The panelists also described AI familiarity as an opportunity for new graduates to contribute immediately. Thompson and Woitkowski argued that AI-native employees can become internal ambassadors who help more experienced workers understand and adopt new tools. Winters

cautioned, however, that AI-native experience is not equivalent to judgment. He has seen particularly strong results from experienced professionals who combine deep technical instincts with sustained use of AI, suggesting that students must also develop taste and the ability to recognize whether a result is appropriate. Employers look for these qualities by asking candidates to explain difficult problems they have solved, how they worked with others, where they used AI, and how they validated the result. Nova described using realistic questions about incomplete healthcare datasets to test both foundational knowledge and reasoning, while Woitkowski emphasized watching whether a candidate demonstrates curiosity, coachability, and learning during the interview rather than judging only whether they reach the correct answer.

Fundamentals Still Matter, but Their Emphasis and Assessment Must Change

Winters offered the strongest challenge to the existing CS curriculum, arguing that much of it is retained through tradition rather than current need. He questioned the value of having students repeatedly implement standard data structures or recreate systems that professional developers would normally reuse, and provocatively suggested that computing education should "tear it up, burn it down." The skills that matter are the ones that curricula were always meant to teach: communication, ethics, higher-order reasoning, design, testing, correctness. Code understanding is much more important than learning syntax. From his perspective, no undergraduate can master the full breadth of computing in four years, and graduates will spend substantial time learning during their first years of employment. Curricula should therefore emphasize durable concepts, continuous learning, and problem solving grounded in meaningful applications. Panelists also reported division, even within the same company, with some agreeing that the curriculum needs a major reset and others strongly defending the continued importance of basic computing knowledge. Without sufficient foundations, students may be unable to understand, test, or repair the systems that AI helps them produce.

Despite this disagreement, the discussion suggested a shared direction for redefining fundamentals. The panelists placed greater value on understanding code and data, testing and validation, debugging, system design, and deciding where human attention is required than on memorizing syntax. Woitkowski argued that education should assess the process by which students solve a problem rather than only the final artifact. She suggested that supervised computer laboratories could allow instructors to observe how students use AI, potentially with undergraduate teaching assistants helping provide the required support. Her own experience building a basic application with an AI-assisted development platform illustrated the distinction: the tool quickly produced an initial system, but deeper technical knowledge was necessary once real problems appeared. Winters similarly proposed renewed attention to pair programming because it provides a way to transfer judgment and design instincts. Together, these comments point toward curricula and assessments that permit AI use while requiring students to explain, validate, revise, and defend the work it helps produce.

Internships Remain a Critical Talent and Trust Pipeline

Woitkowski framed internship hiring as a component of long-term workforce planning rather than a source of temporary labor. Organizations examine where employees are likely to leave, what entry-level roles will exist in the following years, and how interns can be prepared to enter

those positions. Many companies aim to convert approximately 70 percent of their interns into full-time employees and use expected entry-level openings to determine the size of each intern cohort. AI-native interns may also help companies experiment with new tools and understand how AI can be integrated into existing work. At the same time, uncertainty about whether the traditional pyramid-shaped workforce will survive is making organizations more conservative. Woitkowski noted that, generally, companies do not want to recruit large intern classes without credible full-time opportunities afterward. Nova argued that mid-sized organizations may provide especially valuable opportunities because they need current AI expertise but often lack the resources and specialized staff of the largest technology companies.

AI has also increased concerns about whether application materials accurately represent candidates. Woitkowski described résumés generated to mirror job descriptions so closely that candidates can advance through automated screening without possessing the represented skills. Nova shared an even more extreme example in which the person appearing at the end of a hiring process was apparently not the candidate the organization believed it had been evaluating. These problems increase the value of internships, in-person career fairs, references, and sustained interactions that allow employers to observe a candidate over time. The panelists did not object to using AI to improve a résumé, but they emphasized that every claim must be accurate and defensible. Thompson noted that hiring managers increasingly inspect GitHub profiles or other evidence of completed work, while Nova advised students to include concrete details about datasets, performance improvements, failures, and how problems were overcome. Winters summarized the underlying expectation: anything placed on a résumé or GitHub profile is an invitation for the interviewer to ask deeper questions.

Academia-Industry Alignment Requires Sustained, Multi-Party Relationships

The panel did not identify a single source through which companies determine what skills they need. Nova explained that hiring requirements often emerge from the needs of a specific project and the expertise already present on the team. Woitkowski connected those immediate needs to broader workforce planning, including expected departures, organizational restructuring, and uncertainty about future entry-level roles. Thompson argued that universities should communicate directly with first-line managers who are hiring and supervising employees because they often have the clearest understanding of current work. These needs can change too quickly to be captured through occasional job-posting reviews or isolated conversations. Nova contrasted a one-time guest lecture, which may provide a useful individual experience, with sustained industry-supported capstone projects that expose students to the complexity, ambiguity, and changing requirements of professional work.

Woitkowski described the needed model as a form of triangulation among faculty, career services, human resources, and the business or technical representatives responsible for the work. She pointed to North Carolina State University's corporate ePartners structure as a mature example in which companies can support a department through several mechanisms and participate in an advisory board that informs curricular discussions. New collaborations often begin with individual engagements such as capstone projects, and can gradually evolve into deeper relationships. Thompson added that government programs are increasingly pushing industry beyond letters of support and toward real research and educational projects that include students. The discussion of startups and smaller organizations demonstrated why

human resources and policy expertise must also be part of these relationships: Nova described lower-cost and sometimes unpaid opportunities as a way to build experience and trust, while Woitkowski raised legal and fairness concerns and advocated for paid internships. Nova further noted that employment restrictions can limit paid opportunities for international graduate students, highlighting the value of accessible experiential learning opportunities. The CRA-Industry P2P initiative was presented as one mechanism for creating more regular cross-talk through its biennial survey, working groups on shared problems, and reports to academic departments and industry members.

Sustaining Innovation: Navigating Federal Funding in Computing Research

Thursday, July 23, 2026 | 2:00 pm–3:30 pm CT | *Presentation slides*

Organizer

- **Dilma Da Silva**, Texas A&M University

Panelists

- **Irina Dolinskaya,** NSF
- **Hal Finkel,** DOE
- **Patrick Lincoln,** DARPA

Key Takeaways

- **Federal priorities are expanding around mission-driven computing:** NSF, DOE, and DARPA are investing in AI-enabled science, trustworthy computing, cybersecurity and privacy, resilient systems, quantum, and national-scale research infrastructure, with each agency pursuing these areas through its distinct mission.
- **AI and quantum are priorities, not prerequisites for federal funding:** Foundational computing, applied mathematics, systems, security, computing education, and interdisciplinary research remain part of agency portfolios, especially when researchers clearly connect their work to scientific, societal, or national needs.
- **Federal funding stability cannot be guaranteed:** Panelists acknowledged recent award reductions and emphasized that no agency has a “crystal ball,” while also pointing to continuity in many core research programs and agency-specific efforts to reduce uncertainty for funded researchers and students.
- **Large opportunities require genuine collaboration and active engagement:** Major awards increasingly demand multi-institutional synergy, a compelling justification for scale, and early interaction with agencies, reviewers, industry partners, and research consortia; the community was also urged to pursue these opportunities with greater urgency.

Summary

During this session, Dilma Da Silva of Texas A&M University facilitated a panel of federal research leaders — Irina Dolinskaya of NSF, Hal Finkel of DOE, and Patrick Lincoln of DARPA — on how the computing community can navigate a changing funding landscape. The panelists first outlined agency-specific priorities, including trustworthy and resilient computing, AI-enabled scientific discovery, quantum, and national-scale research infrastructure, while emphasizing that foundational computing research continues across their core portfolios. The discussion then turned to the uncertainty created by recent award reductions and terminations, the limits of what agencies can promise about future funding stability, and the different mechanisms they use to manage risk. The session concluded by examining how researchers can position their work, justify the need for large investments, build genuine multi-institutional and industry partnerships, engage more directly with agencies and the federal review process, and move quickly on emerging opportunities.

Mission-Driven Priorities Across Federal Agencies

The panelists' opening presentations emphasized that federal computing investments are organized around agency missions rather than a single government-wide research agenda. Patrick Lincoln described the Information Processing Techniques Office at DARPA as focused on accelerating AI innovation, advancing privacy and security, and understanding resilient "megasystems" composed of as many as a million interacting systems. In AI, this includes developing systems that are safer and more reliable, exploring approaches beyond current transformer- and GPU-centered methods, and connecting university researchers with national-security problems through initiatives such as AI Forge. In security and resilience, Lincoln pushed beyond responding to failures after they occur toward systems whose security properties can be demonstrated and whose emergent behavior can be predicted, controlled, or made robust to environmental damage and malicious attacks.

Hal Finkel centered DOE's presentation on the Genesis Mission, an integrated effort to combine scientific infrastructure, data, and models around grand challenges in energy, discovery science, and national security. He described a longer-term vision of a federated computing environment for science that would make scientific AI models and infrastructure available to a broader community.

Irina Dolinskaya located NSF within the Genesis partnership while outlining NSF's own portfolio, including Future CoRe and its new Computing Education Research program, research on datasets for AI-enabled scientific discovery, and quantum initiatives such as Project Triad. Together, the presentations showed that AI is a shared priority across the agencies, but the motivating problems and funding mechanisms differ depending on whether the goal is basic research, scientific discovery, national security, education, or large-scale infrastructure.

Foundational Research Remains Fundable

When asked directly whether agencies were still supporting foundational work unrelated to AI or quantum, Dolinskaya and Finkel both answered yes. Dolinskaya stressed that NSF's core programs remain active and that AI and quantum received attention because they are newer

and highly visible initiatives, not because they have replaced the rest of the computing portfolio. Finkel similarly pointed to DOE's continuing support for foundational applied mathematics and computer science, including areas such as inverse methods. An additional NSF representative, Ellen Zegura, explained that work connected to the Genesis Mission can also be broadly foundational: proposals may contribute to several challenges, should often be framed as use-inspired rather than narrowly use-specific, and do not necessarily need to focus on AI if they advance the relevant scientific problem space.

The discussion also expanded the opportunity landscape well beyond the most prominent labels. Lincoln identified cybersecurity, privacy, assurance, and interdisciplinary work with areas such as biology, chemistry, and materials science as continuing priorities. Finkel resisted reducing DOE's needs to a ranked list of three topics, instead describing a connected scientific-computing landscape that includes robotics for data collection, edge computing, computing in transit, advanced networking, operation in unusual physical environments, and energy-efficient data centers. Dolinskaya encouraged researchers to look beyond CISE to TIP and other NSF directorates when appropriate. Her message to department chairs and deans was particularly important for junior faculty: they should not be left with the impression that they must abandon their research and work only on AI or quantum to remain fundable. The practical requirement is not to relabel every project, but to clearly explain how it advances an agency's scientific, societal, infrastructure, educational, or security mission.

Funding Stability and Agency Risk

The question-and-answer period confronted concerns about the termination and reduction of federal awards in 2025. Dilma Da Silva asked why researchers should trust that accepted awards would remain stable when recent events had demonstrated that funding could change after an award was made. Finkel and Dolinskaya both acknowledged that they did not have a "crystal ball." Finkel noted that relatively few existing DOE awards in core research areas had been affected by policy changes and that the basic structure of the portfolio remained largely intact, although priorities will continue to shift as science advances. Da Silva provided an important counterpoint from the academic community, observing that many CISE awards had experienced reductions and ensuring that the discussion did not minimize the uncertainty researchers had already faced.

Rather than offering a blanket guarantee, the panelists described agency-specific practices that can reduce or manage some of this risk. Dolinskaya noted NSF's movement from continuing grants toward standard grants that provide the award funding at once rather than depending on later increments. Lincoln explained that DARPA actively redirects resources based on evidence, sometimes increasing funding for promising work even as it ends other activities, and that the agency tries to preserve support for affected graduate students for at least a year. These responses distinguished the relative continuity of many core research portfolios from certainty about any individual award. The agencies could explain their current structures and practices, but they could not guarantee future political, policy, or budget conditions.

Scale, Collaboration, and Active Engagement

Large federal awards remain available, but the panelists emphasized that proposing work at a larger scale requires more than assembling a larger team or budget. Finkel noted that Genesis and other DOE opportunities include awards at several funding levels, with the largest generally supporting multi-institutional collaborations. Dolinskaya pointed to programs such as AI Institutes and Expeditions and stressed that investigators must clearly explain why the work requires a major investment, why it cannot be accomplished through separate smaller awards, and what becomes possible only when the participants work together. Da Silva reinforced this point through her experience with reverse site visits, where it quickly became apparent which proposals represented genuinely integrated teams and which were collections of loosely related projects. For large awards, multidisciplinary breadth is valuable only when the proposal demonstrates real intellectual and operational synergy.

The panel also encouraged researchers to engage with the funding enterprise well before submitting a major proposal. Lincoln highlighted mechanisms through which investigators new to DARPA can introduce their capabilities, including ERIS, open-office activities, informal conversations at conferences, and other opportunities to connect with program personnel. Da Silva urged faculty to accept invitations to serve on review panels or as ad hoc reviewers, both to support the research community and to better understand how proposals are evaluated. Other comments emphasized growing federal interest in partnerships where industry contributes expertise, resources, and meaningful "skin in the game." Finkel closed with a message of both scale and urgency for department chairs and deans: the community must develop technologies that scale across supercomputers, networks, sensors, and scientific facilities while also building programs and a workforce at national scale. These efforts require long-term investment, but maintaining national competitiveness requires beginning them quickly rather than waiting for the funding landscape to become predictable.

We've Been Here Before: Innovations in Computing

Thursday, July 23, 2026 | 4:00 pm–5:30 pm CT

Organizer

- **Raquel Hill,** Spelman College

Panelists

- **Tom Cortina**, Carnegie Mellon University
- **Amanda Holland-Minkley**, Washington and Jefferson College
- **Annabel Rothschild**, Bard College
- **Tiffani Williams**, University of Illinois at Urbana-Champaign

Key Takeaways

- In CS curricula, it is important to strike a balance between specialization and preparation in a generalizable computing core.
- As a field, we should examine how much as educators we should be beholden to industrial pressure and change. Computing is somewhat unique in how much and quickly it changes to align with shifting industrial needs.

Summary

This session was organized to discuss a sentiment published by Valerie Barr where she expressed "Feeling Cranky about AI and CS Education." In this essay, Barr discussed how computing educators are responding to advances in generative AI, and in particular, changing industrial pressures. She argues that this transformative moment is not all that different from moments in computing's past where educators have also overhauled the curriculum to respond to changing practice and innovation, a shift that has not always led to lasting educational improvements, instead leading to courses codifying what was ultimately a passing fad. In addition, she argues that fully overworking a curriculum to center generative AI has the potential to stunt students' critical thinking skills while simultaneously underpreparing them in the key skills needed to become computing experts.

The session started with a summary of this sentiment by the organizer, Raquel Hill. This summary focused on how computing education has gone through several transformations since its inception; for instance, at one point everyone had to learn Java and Object Oriented programming, while at another point Web and E-commerce courses became ubiquitous. Next, each panelist gave a quick lightning talk, summarizing their experiences and perspective. The rest of the session was organized as a series of questions posed by both Hill and the audiences followed by answers from the panel. Here, we summarize each lightning talk and key points from the question and answer session.

Tom Cortina discussed his experience as a teaching professor at Carnegie Mellon University as well as the ACM computing education co-chair. He noted that students are now more likely to use AI than attend office hours when they need extra help. He also noted that more students are using AI to complete homework quickly, often producing incorrect answers. As a result, many classes at CMU no longer require homework and instead are shifting toward assessments such as in-class exams or code interviews. In his course, for example, students still do programming assignments, but then participate in an interview about it after completion and have to maintain a diary about how they used AI.

Amanda Holland-Minkley is a Professor at Washington & Jefferson College, a small liberal arts college in an interdisciplinary computing program. She is also grappling with how AI use by students is changing computing education. In particular, she is trying to find ways to break students out of "AI-first" rather than "AI once you know what you want to do." She is also trying to encourage an academic culture in which students choose to learn with their classmates and talk to their professors. In a liberal arts setting, there is less of a "how do you do it at scale" problem, so she is also able to play with assessment changes including oral exams or required 15-minute check-ins with a tutor.

Annabel Rothschild is an Assistant Professor at Bard College and a colleague of Valerie Barr. By training, she is an AI ethicist, and she teaches a class on the history of AI. This includes three "AI waves," each of which was marked by a technical demonstration, an investment bubble, and an AI winter. Thus, the current wave of generative AI investment raises questions about whether another bubble is forming. She notes that at Bard, many students study the arts and humanities, and many of them are really anti-AI. This poses interesting and challenging conditions regarding how and where to teach about AI.

Tiffani Williams is a teaching professor at the University of Illinois Urbana-Champaign. At UIUC, she started an online one-year certificate program to teach computing fundamentals to those with a college background in something other than computing who are interested in moving into the field, but are not yet ready for a CS Masters program. Students also have the opportunity to pursue research experience with a faculty member, and to date, seven certificate students have matriculated into UIUC's PhD program. Williams notes that even without AI, as a discipline, we have been a bit exclusive with who we have educated; AI has the potential to widen these gaps.

Question and Answer

Q: As a discipline, we've been here before. Are there similarities or differences between this moment of change and previous ones? And are these differences critical?

The discussion emphasized that computing has experienced similar moments of rapid technological and curricular change before, mentioning MOOCs, bootcamps, cloud computing, and other such transformations. Williams and Holland-Minkley both noted that many of these changes were initially framed as transformative or as ways to broaden participation in computing, but that they did not ultimately replace the need for structured computing curricula or automatically make computing more equitable. Cortina noted that generative AI may lead to a more significant shift because it acts less like a traditional computational tool and more like a human collaborator, which may influence how students perceive and experience the discipline. Panelists also noted that gen AI may represent a "tipping point" where it is impossible to teach both traditional computing curricula and AI-related additions during the same program; thus it may force computing educators to reconsider what learning goals and topics are at the core of computing.

Q: Are there pressures to create new specializations, degree programs, etc? What are the sources? Are these changes warranted?

Multiple panelists noted that industry, alumni, and student families are all creating pressure for more AI-focused programs and specializations. Schools are responding to these pressures in different ways, with some maintaining a general degree with sub-tracks while others are starting new hyper-specialized programs. However, panelists cautioned against over-specializing around current trends, emphasizing the importance of educators and curriculum designers prioritizing a strong computing core that prepares students for technologies and careers that will continue to change over time. Panelists agreed that educators should prioritize the education that best serves their students, even if it is not fully aligned with external pressures.

Q: Why does the CS community respond to innovations in this way?

Cortina noted that abstraction is a core part of the CS ethos, with computer scientists often explicitly searching for a way to abstract out other parts of the discipline. Thus, when a new technology arrives that makes it possible to do something faster, it can be tempting to feel we don't need to teach the lower-level way, even if that approach may be helpful for teaching critical thinking and problem solving. Williams also notes that CS often abstracts away the people and historical context behind innovations; this can make advances feel entirely new and make it easier to overlook their human impacts. Rothschild notes that CS is an applied discipline at its core, while also being a fairly young discipline constantly facing technological shifts. As described by Holland-Minkley, this can make building CS curricula feel like "trying to build a house on a fault line." In addition, several panelists noted that compared to other disciplines, historically, CS funding and advancement is uniquely tied to industry and defense. This history and present reality of the funding landscape may also contribute to why the CS community responds to innovations in this way. Q: What is the impact on the discipline, students, and broader community if we incorrectly integrate AI into computing education?

Multiple panelists noted that equity is a concern, both across and within institutions. For instance, many schools may not have the funds to support AI-coding infrastructure subscriptions for their students; this can create inequities when some students can afford more powerful models than others. Rothschild and Holland-Minkley both note that the stakes around AI integration are larger for smaller institutions such as liberal arts colleges.

Resources

- *Feeling Cranky about AI and CS Education* by Valerie Barr, https://cacm.acm.org/blogcacm/feeling-cranky-about-ai-and-cs-education/

Lightning Talks from Emerging Leaders

Lightning Talks from Emerging Leaders

Wednesday, July 22, 2026 | 5:40 pm–6:20 pm CT

In the 2026 CRA Summit, there were two sessions reserved for lightning talks from members of the 2026 Emerging Leaders Cohort. Lasting for two to three minutes, each of these presentations involved one member of the cohort having the opportunity to describe a novel institutional model, ask a provocative strategic question, or introduce a framework or process for the consideration of the audience. In this section, the talks are listed in the order in which they were delivered during the CRA summit. The first 15 were given on the first full day of the summit, while the rest were delivered in a session on the second day.

Organizer

- **Rachel Pottinger,** Professor of Computer Science, University of British Columbia

Speakers

- **Ada Gavrilovska**, Georgia Institute of Technology
- **Anna Lysyanskaya**, Brown University
- **Corin Pitcher**, DePaul University
- **Diba Mirza**, University of California at Santa Barbara
- **Dong Wang**, University of Illinois at Urbana-Champaign
- **Elena Sherman**, Boise State University
- **Graciela Perera**, Roosevelt University
- **Hari Sundar**, Tufts University
- **Heather Lipford**, University of North Carolina at Charlotte
- **Itsik Pe'er**, Columbia University
- **Jennifer Mankoff**, University of Washington
- **Jia Rao**, University of Texas at Arlington
- **Jianchen Shan**, Hofstra University
- **Jignesh Patel**, Carnegie Mellon University

Speaker: Ada Gavrilovska, Georgia Institute of Technology
Title: AI-powered research vs. Human-powered peer review

Key takeaways: Gavrilovska highlighted a widening gap between AI-powered research and human-powered peer review. She explained that systems papers once required 18–24 months of intensive work to build prototypes, develop benchmarks, and port baselines, whereas generative AI can now shorten the journey from idea to submission to only a few months. She illustrated the resulting pressure with striking submission numbers: ATC submissions reportedly rose from 400 to 900+, OSDI submissions doubled, and the first cycle of ASPLOS 2027 received 50 percent more submissions than even an exponential-growth model predicted. She cautioned that, although reviewer pools are growing, they are not keeping pace with submissions. She urged academic leaders to rethink peer review and prepare for faculty-candidate portfolios that may look markedly different in the coming years.

Speaker: Anna Lysyanskaya, Brown University

Title: Privacy-Preserving Identity

Key Takeaways: Lysyanskaya shared what she learned from her experience serving as the chair of Brown's faculty executive committee over the 2025–26 academic year. In particular, she called on members of the CRA community to engage with faculty governance; "if we want our vision to come to fruition in our respective institutions, we need to be in a position to tell our own stories." Engaging with faculty governance can lead to increased public awareness and understanding of computing research. Lysyanskaya's research is also about providing individuals control over their own stories, specifically by enabling users to share verifiable personal data only when they choose, using privacy-preserving cryptographic protocols. Lysyanskaya argues that such protocols are ready for widespread adoption, and in some places, including the European Union, already mandated by law. As a result, she encourages CS researchers to work on making such protocols more secure and more usable in practice.

Speaker: Corin Pitcher, DePaul University

Title: Guardrails and On-Ramps: Governance and Enablement for Campus Agent Platforms

Key Takeaways: Professor Pitcher's talk was a call-for-collaboration for building a shared, governed AI substrate that could serve campus communities with varying levels of readiness and enthusiasm. He argued that such an environment must combine strong guardrails — including sandboxed execution, scoped access, auditing, and cost governance — with accessible on-ramps such as pooled resources, training, documentation, and user support. He emphasized that this approach could protect sensitive institutional data while enabling equitable AI adoption across research, teaching, and administration. His proposed blueprint would support commercial and open-source tools, cloud-based and self-hosted models, and long-running AI agents, allowing institutions to expand access without sacrificing security or oversight.

Speaker: Diba Mirza, University of California at Santa Barbara

Title: Launching UCSB's new B.S. in AI Degree

Key Takeaways: Mirza shared what she learned from her experience helping lead the development and launch of UCSB's new bachelor's degree in AI. It was UCSB's first new bachelor's degree in a decade, and its launch required sustained work by faculty members to secure administrative approval at each stage. Mirza noted, however, that throughout the process, she found it absurd that there is a gap between how far AI research has moved and how little consensus there is on how to teach AI. Thus, there is no consensus on what an undergraduate AI degree should entail. She ended with a call to the community to actively consider and discuss such questions.

Speaker: Dong Wang, University of Illinois at Urbana-Champaign

Title: Building AI that Understands Society

Key Takeaways: Prof. Wang described a research vision for building AI systems that

understand society as deeply as they understand information. He noted that AI has made substantial progress in understanding the cyberworld through capabilities such as code generation and knowledge reasoning, as well as the physical world through robotics and embodied AI. However, understanding societies — shaped by human values, relationships, and collective behavior — remains largely unexplored. Drawing on a decade of research in human-centered social sensing, he explained how multimodal language models could help analyze human behavior. He envisioned AI that collaborates with people, supports collective well-being, and responds thoughtfully to social contexts as it becomes embedded in education, healthcare, science, and public decision-making.

Speaker: Elena Sherman, Boise State University

Title: Realistic Program Benchmarks for Verification Tools

Key Takeaways: Professor Sherman's research focuses on program verification, particularly the development of benchmarks for evaluating verification tools. She emphasized the importance of creating realistic benchmarks that support generalizable assessments of a tool or approach. Modern formal-verification tools often cannot analyze large, real-world code repositories, while existing benchmarks tend to be overly simple or contrived. The challenge, therefore, is to construct benchmarks that are complex enough to be computationally meaningful yet tractable for state-of-the-art tools. Her research group identifies challenging verification problems in real-world projects and transforms them into benchmarks. These benchmarks differ from conventional ones because they retain the complexity of authentic software. High-quality benchmarks are essential for understanding the limitations of existing verification tools and determining how those tools can be improved.

Speaker: Graciela Perera, Roosevelt University

Title: Solving Problems That Matter

Key Takeaways: Prof. Perera presented applied computing as a means of solving real problems and creating tangible benefits for students and their communities. She connected her vision with Roosevelt University's social justice mission, emphasizing the preparation of socially conscious graduates equipped to address contemporary challenges. She encouraged educators to move students beyond textbooks and theory into purposeful, hands-on experiences with artificial intelligence, large language models, quantum computing, cybersecurity, and future networks. She argued that students should learn not only how emerging technologies work, but also how to apply them responsibly to advance social change. She also highlighted her work with the Computing Alliance of Hispanic-Serving Institutions to broaden Hispanic participation in computing and invited attendees to collaborate on consequential work. She closed on a personal note, asking the audience to keep the people of Venezuela in their thoughts following an earthquake.

Speaker: Hari Sundar, Tufts University

Title: AI Writes the Science. Who Checks Its Work?

Key Takeaways: Sundar began by asking how many audience members had used AI to

generate code. Most raised their hands. When he asked how many had used AI to formally verify that code, almost all hands went down. Sundar emphasized that there is a need for formal verification; testing alone isn't tractable or sufficient in safety-critical systems such as airplane software or drug design. He notes that it is vastly cheaper to check a proof than it is to find one; even if the model generating the proof is unsafe, we can check the provided certificate quickly to ensure safety in a specific instance. He ends by connecting this to science, arguing that the question of the decade isn't if AI is going to write our science, but rather it's who's going to check its work.

Speaker: Heather Lipford, University of North Carolina at Charlotte

Title: Human-Centered Computing and Education

Key Takeaways: Professor Lipford traced UNC Charlotte's growth in human-centered computing from her arrival as its first faculty hire in human-computer interaction to the creation of a dedicated Department of Human-Centered Computing and Education. She explained that her search for collaborators among the university's strong cybersecurity community led her to usable security and privacy, which became her research home. Over two decades, the HCI area expanded from a single course to multiple courses, a concentration, a graduate certificate, and a larger faculty community. She celebrated the new department, chaired by Manuel Pérez-Quiñones, as well as the parallel creation of a Department of Cybersecurity. At the same time, she acknowledged the administrative challenges ahead: as graduate program director for master's programs now spanning both departments, she will continue to have "one foot" in each while helping the newly reorganized units collaborate and grow.

Speaker: Itsik Pe'er, Columbia University

Title: Starting an Interdisciplinary Center

Key Takeaways: Cross appointed between computational biology and computer science, Professor Pe'er shared his experience starting an interdisciplinary center. He has been working to launch NSCORE: the National Synthesis Center for Organismal Resilience. This center will conduct research at the intersection of biology and computation with the goal of trying to understand how microbes respond to the environment. Doing so involves analyzing massive amounts of data that spans modalities, and thus the work will lead to advances in both the biological and computing disciplines. Pe'er ends by stressing the importance of interdisciplinary research and indicating that such work takes more than a village, instead taking a great city of support from both individuals and communities.

Speaker: Jennifer Mankoff, University of Washington

Title: Disability-Led Innovation

Key Takeaways: Professor Mankoff opened by audio-describing herself, including the chicken perched on her shoulder in her title-slide photograph. Drawing on her own experience with disability, she noted that roughly 20 percent of students have disabilities and that many more people rely on accessibility tools such as captions. She introduced the Center for Research and Education on Accessible Technology and Experiences (CREATE), whose mission is to

make technology accessible and make the world accessible through technology. She argued that this work must combine technical innovation with activism, advocacy, and a commitment to ensuring that disabled people help shape the solutions. As an example, she described a collaboration to make computer-aided design accessible through a low-resolution dynamic tactile display, enabling blind users to construct and examine precise models. She concluded by urging attendees to include disabled people in their laboratories and classrooms, emphasizing that their expertise and perspectives are indispensable — and that accessible technologies often benefit everyone.

Speaker: Jia Rao, University of Texas at Arlington

Title: Scaling with Disaggregated Shared Memory (DSM)

Key Takeaways: Professor Rao researches operating systems, distributed computing, and architecture. He is interested in scaling distributed computing with disaggregated shared memory (DSM). Traditional DSM approaches rely on software runtime to construct a shared address space and maintain data coherence, but both of these tasks are challenging, which means that the promise of disaggregated shared memory to reduce programming burden remains largely unrealized in practice. However, recent advancements in hardware-supported shared addresses across machines has reduced the first challenge for DSM. He argues that it is important for researchers to focus on targeting the second remaining challenge, and he ends by asking what applications and operating systems will look like for this emerging hardware environment.

Speaker: Jianchen Shan, Hofstra University

Title: Advancing the Cloud OS to Bridge the VM-host Gap

Key Takeaways: Prof. Shan described a persistent cloud-computing problem: a guest operating system often sees a misleading picture of the physical machine. Virtual CPUs may be withdrawn without warning, cache and memory are shared with invisible neighbors, and memory that appears contiguous inside a virtual machine may be fragmented on the host. Traditionally, researchers addressed this mismatch through hypervisor-specific patches and coordination between the guest and host. Prof. Shan argued that this approach no longer scales across numerous, often proprietary hypervisors and is further complicated by confidential-computing technologies that deliberately remove trust between guests and hosts. His team instead uses lightweight probes inside the virtual machine to infer the true state of shared resources at runtime. This information helps the guest's memory allocator and scheduler make better decisions, improving performance and resource utilization without relying on the hypervisor. He highlighted that Hofstra undergraduates conducted this work with institutional program support.

Speaker: Jignesh Patel, Carnegie Mellon University

Title: Welcome to Lake Mystic

Key Takeaways: Professor Patel, currently the interim chair of CS at CMU, discussed how advances in AI are opening up unexpected research directions and opportunities. In particular,

he talked about regular expressions: AI agents love them and use them extensively, but not always efficiently. He described a new way to approach regular expression evaluation through a declarative database lens. Rather than treating a regular expression as a monolithic specification, we can decompose it into a semantically equivalent conjunction of declarative predicates. This creates opportunities to apply familiar database optimization techniques, including predicate reordering and cost-based execution planning. By exploiting the problem's declarative structure, these techniques can substantially reduce evaluation cost and deliver significant performance gains.

Lightning Talks from Emerging Leaders

Thursday, July 23, 2026 | 5:40 pm–6:20 pm CT

This is continued from the lightning talks from the first day

Organizer

- **Rachel Pottinger,** Professor of Computer Science, University of British Columbia

Speakers

- **Juhua Hu**, University of Washington Tacoma
- **Leonid Sigal**, University of British Columbia
- **Rasib Khan**, Northern Kentucky University
- **Matt Weinberg**, Princeton University
- **Mohammed Seyam**, Virginia Tech
- **Pamela Wisniewski**, International Computer Science Institute
- **Raffi Khatchadourian**, City University of New York Hunter College and Graduate Center
- **Reetu Das**, University of Michigan; ACM SIGMICRO
- **Sandeep Kulkarni**, Michigan State University
- **Shanon Reckinger**, University of Illinois Chicago
- **Stacy Branham**, University of California, Irvine
- **Susmit Shannigrahi**, Tennessee Technological University
- **Tamara Silbergleit Lehman**, University of Colorado Boulder
- **"YZ" Yezhou Yang**, Arizona State University

Speaker: Juhua Hu, University of Washington Tacoma

Title: “What’s the future of AI?”

Key Takeaways: Hu opened by reflecting on how often she is asked, “What is the future of AI?” In response, she described how AI education and research are taking shape at UW Tacoma. This includes a proposed M.S. in artificial intelligence organized around graduate credentials in AI Foundations and AI Applications, as well as plans for undergraduate AI education that connects AI with students’ career interests. Hu linked these efforts to her teaching in big data analytics and data science and to her research in multimodal representation learning, including context-aware methods supported by multimodal foundation models.

Speaker: Leonid Sigal, University of British Columbia

Title: Visual Perception and Reasoning

Key Takeaways: Sigal discussed the work of the approximately 140 faculty at the Center for AI Decision-making and Action (CAIDA), where he serves as deputy director, in addition to his roles as a CIFAR AI Chair at the Vector Institute and Canada Research Chair in Computer Vision and Machine Learning at The University of British Columbia. He then described his own research on the shift toward unified, general-purpose models, particularly large language models with visual capabilities. His work examines how these models encode visual information and develop emergent abilities such as grounding and grouping, while seeking to improve their consistency, reliability, safety, and visual reasoning. He also studies more efficient approaches to training and inference, including active learning and token reduction.

Speaker: Rasib Khan, Northern Kentucky University

Title: Closing the AI Confidence Gap: AI literacy isn’t another course… it’s a program design problem

Key Takeaways: Khan argued that AI literacy is not simply another course to add to the curriculum, but a program-level design problem. Responding to industry expectations that graduates use AI effectively, Northern Kentucky University is exploring incremental changes that distinguish basic familiarity with AI from the deeper problem-solving capabilities developed through computing education. Khan described a progression in which students first learn to use AI tools, then solve problems and build with AI, while developing the judgment required to use these technologies responsibly. An initial pilot identified four courses in which to apply this progression, providing a practical starting point for integrating AI literacy across the program without requiring an immediate curriculum-wide overhaul.

Speaker: Matt Weinberg, Princeton University

Title: Economics and Computation

Key Takeaways: Weinberg traced his path from algorithms and theoretical computer science into economics and computation, where the behavior of users becomes part of the computational problem. His work in algorithmic mechanism design examines how algorithms must change when the people interacting with them pursue their own goals, bringing economic

reasoning about incentives into the design itself. The same connection between abstract theory and recognizable human behavior also makes the field a natural gateway for undergraduates: students are drawn in by the applications and then move toward deeper research in algorithms and theory. Across his own research trajectory and his mentoring, Weinberg presented compelling applications not as a departure from theory, but as a way to discover and motivate new theoretical questions.

Speaker: Mohammed Seyam, Virginia Tech

Title: Software is Global. Is Computing Education? (Global collaboration as a core computing competency)

Key Takeaways: Seyam described Virginia Tech's Software Engineering in a Global Context program, which treats international collaboration as an essential software engineering competency. Students begin the semester working remotely with peers at partner universities, meet abroad for an intensive in-person development sprint, and then return to distributed collaboration to complete their projects. A recent offering brought together 23 students from four universities in the United States, Egypt, and Switzerland to develop software for real-world clients. Beyond technical development, the program gives students experience navigating cultural differences, communication challenges, ambiguity, trust, and collaboration in authentic global teams. These experiences are especially important as AI becomes more capable: AI can support technical learning, but students still need authentic experiential opportunities to develop communication, collaboration, judgment, and adaptability.

Speaker: Pamela Wisniewski, International Computer Science Institute

Title: Safety by Design: Building the Village That Protects and Empowers Youth Online

Key Takeaways: Wisniewski drew on her human-computer interaction research in adolescent online safety to challenge the computing community to become the "village" young people need to navigate online risks without being pushed offline. Her safety-by-design approach incorporates protections from the beginning of the technology design process while preserving the interactions and opportunities youth find online. Her 2019 NSF CAREER project, Safety by Design: Protecting Adolescents from Online Risks, advanced this teen-centered approach by combining adolescent developmental psychology, interaction design, and computer science and involving teens directly in co-designing interventions.

She emphasized that adolescent online safety is not solely an HCI problem; it also requires contributions from AI, security, systems and networking, software engineering, and data science. The goal is not merely to protect young people from harm, but to provide the agency and support they need to thrive as technologies evolve. The problem is too large for any one group to solve and too urgent for the computing community to ignore.

Speaker: Raffi Khatchadourian, City University of New York Hunter College and Graduate Center

Title: Making Deep Learning Code Faster Without Breaking It

Key Takeaways: Khatchadourian described how his background in programming languages and software engineering informs his work on improving the performance of deep learning systems. Rather than viewing these systems only as models, his research treats them as complex software with substantial technical debt and opportunities for systematic analysis. He develops techniques that accelerate model training and, in some cases, inference while preserving the code's intended behavior. Across this work, the resulting optimizations have produced an average training speedup of approximately two times.

Speaker: Reetu Das, University of Michigan; ACM SIGMICRO

Title: Breaking the Memory Wall: Hardware-Software Co-Design for the Data-Movement Era

Key Takeaways: Das reflected on how hardware acceleration has grown from an emerging research direction at the start of her career into a central focus of modern computer architecture. Her research addresses the "memory wall": in AI and high-performance computing, moving large models and vast, often sparse datasets between processors and memory increasingly dominates computing time, energy, and cost. Das develops in- and near-memory architectures that reverse the traditional model by bringing computation closer to where the data resides, reducing data movement and improving the efficiency of memory-bound workloads.

Speaker: Sandeep Kulkarni, Michigan State University

Title: Synchrony, Asynchrony, and More

Key Takeaways: Kulkarni connected his research on synchronization, consistency faults, distributed data structures, and synchronization-free concurrency to his teaching of foundational computing concepts. He emphasized helping students recognize common structures across seemingly unrelated problems, a skill that can improve how they reason with AI. At Michigan State University, he helps oversee the MS in Data Science, MS in FinTech and AI, and AI minor. The university is also developing BS and MS programs in AI. These efforts are supported by the Leinweber Center for Engineering and Digital Innovation and growing university-wide commitment to AI literacy. Kulkarni highlighted the importance of leadership support, lessons learned from incorporating AI into existing curricula, and adapting programs to changing enrollment patterns.

Speaker: Shanon Reckinger, University of Illinois Chicago

Title: The Science of Computing Competence

Key Takeaways: Reckinger presented her research on defining computing competence and building assessment systems that capture it. Noting that technologies are often released with greater attention to capability than impact, she questioned what today's technology leaders reveal about the competencies computing education prioritizes. She distinguished improving

existing assessments from rethinking assessment itself. Traditional approaches infer learning from exam scores; her proposed direction gathers evidence of learning that instructors interpret to support claims of competence. She highlighted oral assessments, mastery-based approaches, open-ended tasks, and teamwork as ways to move beyond highly specified, high-stakes written exams toward more authentic demonstrations of learning.

Speaker: Stacy Branham, University of California, Irvine

Title: Digital Accessibility Laws: An Opportunity

Key Takeaways: Branham highlighted an upcoming federal accessibility deadline and the substantial effort facing educators in making courses accessible to students with disabilities. She framed the deadline as both a reason to act and an opportunity to build support for accessibility improvements in computing education. In response, she introduced a CRA-E survey of classroom technologies. By pooling information across institutions, the initiative aims to reduce the burden of keeping up with changing technologies and help educators identify tools that work well for students with disabilities. At the time of the talk, the survey had received more than 450 responses from over 175 colleges and universities, covering more than 2,000 tools.

Speaker: Susmit Shannigrahi, Tennessee Technological University

Title: Teaching Networks to Manage Themselves

Key Takeaways: Shannigrahi outlined a vision for AI-native networks that move beyond static rules and isolated optimization. He noted that today's networks rely on rules tuned for expected conditions, which can miss novel threats, while existing approaches often address individual objectives rather than considering the network as a whole. His vision centers on foundation models that learn network-state embeddings from telemetry, topology, and traffic, enabling networks to generalize beyond hand-written rules and jointly balance performance, security, and cost. The goal is to move data more efficiently, respond faster to changing conditions, and adapt automatically as application demands shift.

Speaker: Tamara Silbergleit Lehman, University of Colorado Boulder

Title: Future of Computing Research is with Undergraduate Students

Key Takeaways: Lehman argued that AI is pushing universities to redefine the value of higher education, with undergraduate research offering students opportunities to create new knowledge rather than only consume it. She emphasized that the research pipeline begins before graduate school and highlighted CU Boulder's Discovery Learning Apprenticeship (DLA) and Summer Program for Undergraduate Research (SPUR) as college-backed opportunities for undergraduates. She also described the Europe-Colorado Program, which provides research experiences for students from European-based universities. Building on this model, Lehman is leading an initiative to expand these opportunities to Latin America, beginning with partnerships in Buenos Aires, Argentina.

Speaker: “YZ” Yezhou Yang, Arizona State University

Title: Assist-per-Game

Key Takeaways: With a nod to basketball’s “assists per game,” Yang introduced his Active Perception Group (APG) and its work in computer vision and “AI Evaluation as a Science.” He highlighted efforts to evaluate frontier models and assess autonomous vehicles through sensing-based approaches. Yang also pointed to opportunities for engagement with domain experts, while noting that perceptions of AI outside the research community are often inaccurate. He encouraged CRA to train AI researchers to communicate more effectively with the public, helping audiences better understand AI’s capabilities and limitations.