Carnegie Mellon University

the LINK

THE MAGAZINE OF CMU'S SCHOOL OF COMPUTER SCIENCE

FALL 2024
ISSUE 18.2

AI AND THE Election

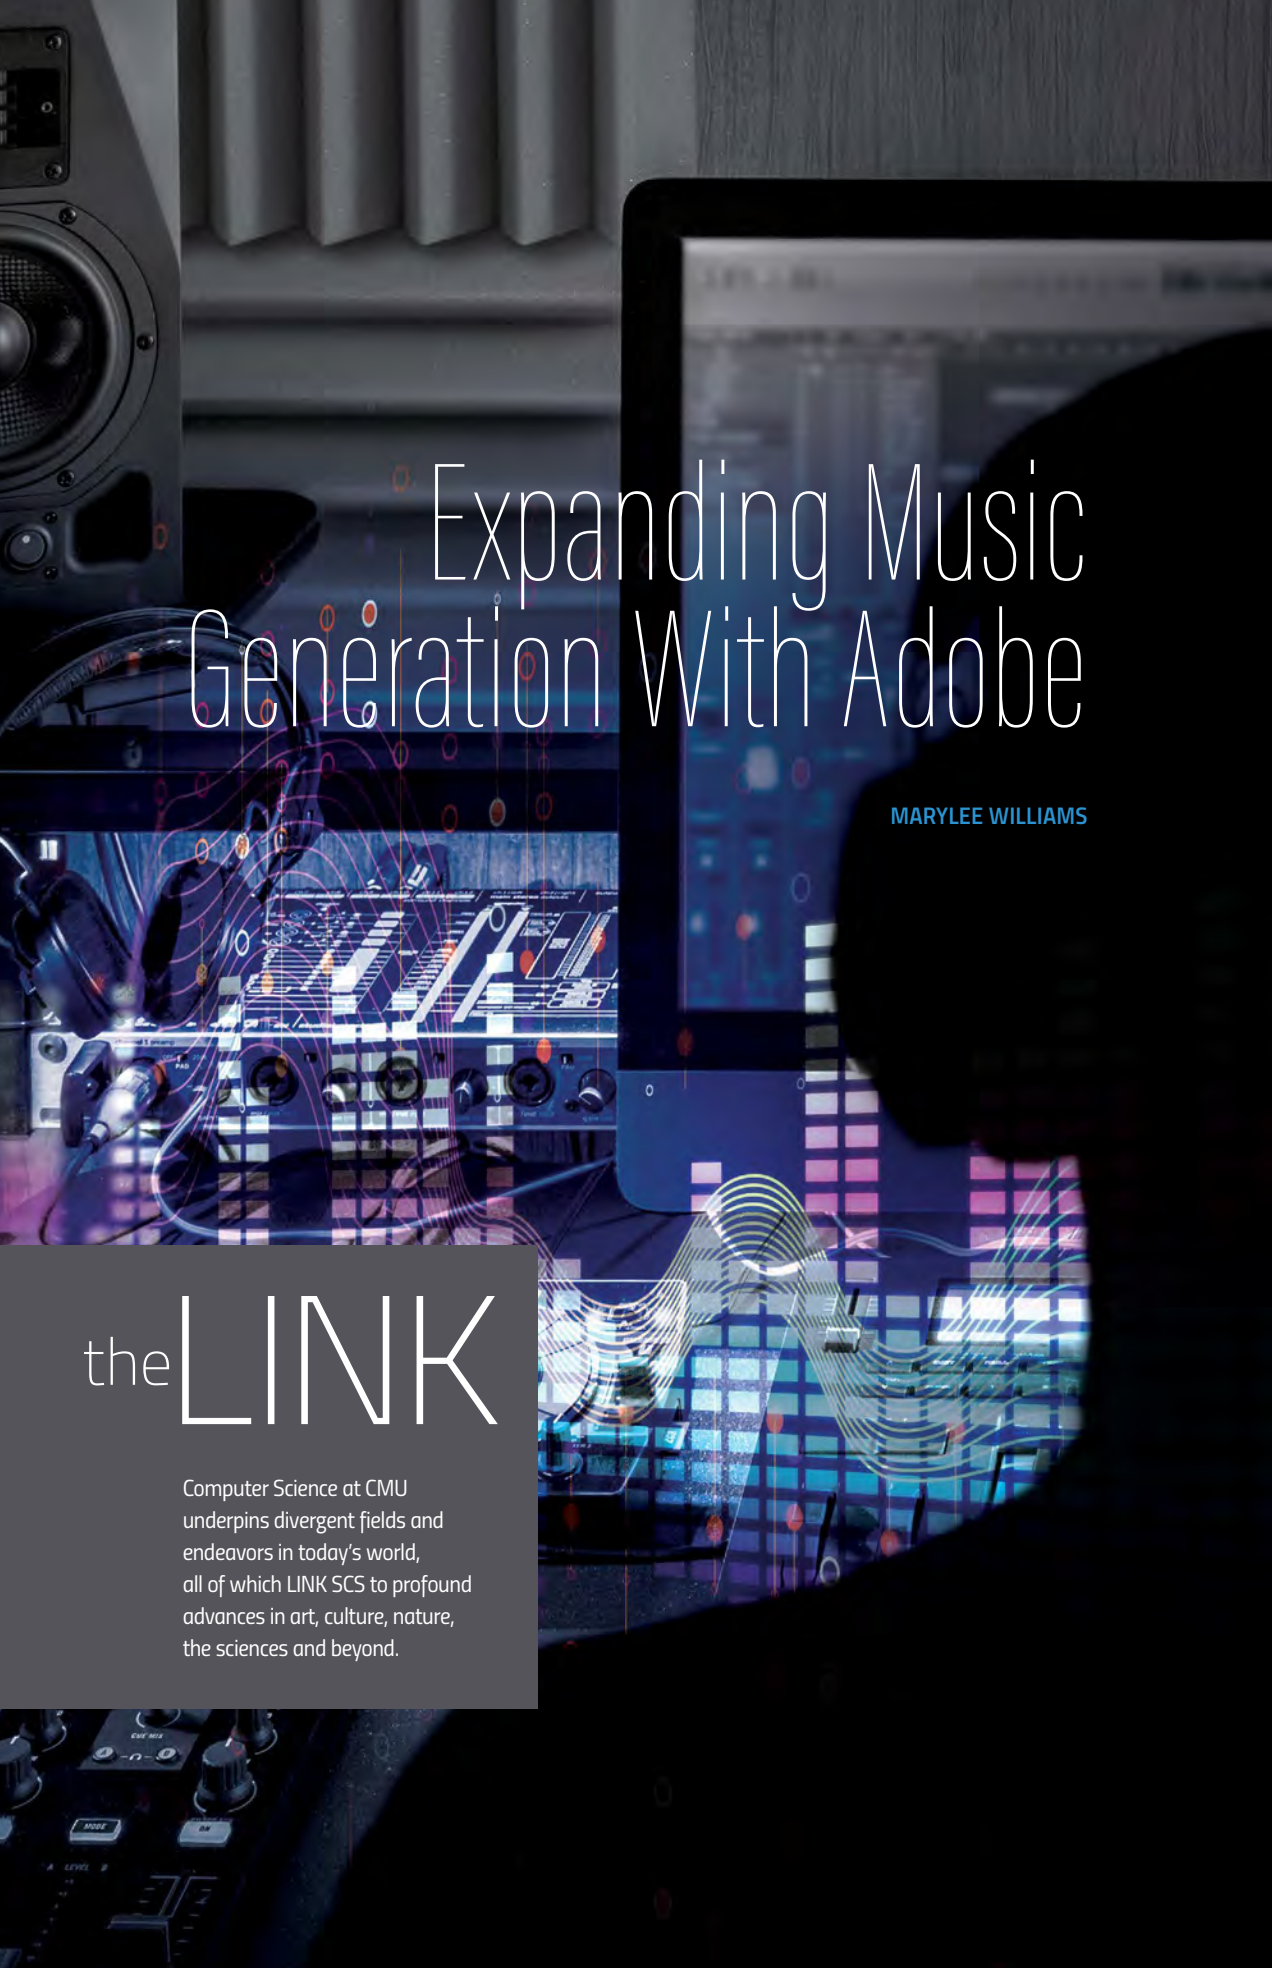

Expanding Music Generation With Adobe

MARYLEE WILLIAMS

the LINK

Computer Science at CMU underpins divergent fields and endeavors in today's world, all of which LINK SCS to profound advances in art, culture, nature, the sciences and beyond.

SCS researchers were part of the team behind one of Adobe's newest ventures, the generative AI music creation and editing tool called Project Music GenAI Control.

An avid classical piano and viola player, **Shih-Lun Wu** learned viola because all the violin seats in his school orchestra had been taken. Now a student in Carnegie Mellon University's School of Computer Science, Wu uses generative AI and machine learning to make music creation more accessible and engaging for people of all abilities.

Wu, who is pursuing a master's degree in language technologies, was part of the research team behind one of Adobe's newest ventures, the generative AI music creation and editing tool called Project Music GenAI Control. Wu co-developed an aspect of the project, Music ControlNet, with SCS faculty members **Chris Donahue** of the Computer Science Department and **Shinji Watanabe** of the Language Technologies Institute. The CMU researchers collaborated with Nicholas J. Bryan, a senior research scientist and head of the Music AI research group at Adobe Research.

Music ControlNet gives users command over aspects of generated audio like the melody, rhythm and dynamics. Users can play, compose or draw what they want these elements — or any combination of them — to look like. The user then combines that information with a text prompt, such as "happy jazz" or "sad country," that they send to a generative AI model to produce audio. Music ControlNet can also co-create with the user, improvising on a partial melody or other inputs.

Wu and colleagues adapted previous research into pixel-level controllable image generation to create Music ControlNet. The melody, rhythm or dynamic references and text inputs are transformed into an image-like representation of the music, which is then converted to audio. This work was accepted to IEEE Transactions on Audio, Speech and Language Processing, a journal for audio and music work.

"This is just the beginning," Wu said. "Ultimately, with more effective and comprehensive controls, we hope to accelerate the creative workflow for music professionals, and we can make music creation easier for the general public."

Adobe's Project Music GenAI Control also involved researchers at the University of California, San Diego. Learn more about the project on Adobe's website. ■

the LINK

Fall 2024 | Issue 18.2
The Link is the magazine of Carnegie Mellon University's School of Computer Science, published for alumni, students, faculty, staff and friends.

Publisher
Martial Hebert

Editor
Kevin O'Connell

Contributing Writers
Aaron Aupperlee, Sarah Boden, Sheri Hall, Danielle Hoffman, Kevin O'Connell, Chris Quirk, Kelly Saavedra, Byron Spice, Steph Stricklen, Brian Thorton, Matt Wein, Marylee Williams

Design
Susie Hernishin

Photography
Carnegie Mellon University Archives, Elan Mizrahi, Nate Smallwood

Illustration
Susie Hernishin

Office of the Dean
Gates Center for Computer Science
Carnegie Mellon University
4902 Forbes Avenue
Pittsburgh, PA 15213
cs.cmu.edu

magazine@cs.cmu.edu

40

12

contents

Safeguarding AI for a Promising Future

The perception of how humans interact with AI is changing, placing people at the center of every conversation.

While respecting and safeguarding against its potential threats and misuse, the future of AI holds boundless promise. From revolutionizing healthcare and personalized medicine to enhancing education accessibility and transforming industry through automation, AI has the potential to reshape our world for the better. It is the Industrial Revolution of our time.

Within SCS and across CMU, our researchers and students stand at the forefront of pioneering technologies that not only push the boundaries of innovation but also prioritize ethical considerations and societal impact. As AI is applied to nearly every field, new avenues for growth and discovery open up. The addition of AI tools to new fields not only changes how we approach that field, but also allows for the development of new AI tools and concepts, which we can then apply to other fields in new ways.

As AI continues to evolve, so must our efforts to ensure that these technologies are deployed with consideration for their broader implications. We remain actively engaged in interdisciplinary research initiatives that explore the ethical, legal and social dimensions of AI, working closely with policymakers, industry leaders and community stakeholders to establish guidelines and frameworks for responsible AI development and deployment.

At CMU, it's often difficult to discuss the specifics of AI because the topic itself is too broad here. We utilize AI across so many aspects of our university.

In this issue of The LINK, you will find articles that detail a few areas of AI application in SCS, as well as efforts to safeguard humans. The cover story on AI and Elections informs us about the potential threats and safeguards we are building to ensure the structures of our democracy are not corrupted by fraud, from internal threats and those coming from outside the U.S. I hope you'll also read the article on how AI has already changed scientific discovery and will continue to do so in ways we have yet to fully comprehend. Vincent Conitzer, professor in our Computer Science Department, has co-written a wonderful book entitled Moral AI, which inspired the article about how we ought to be thinking about the protections we set up to ensure the safety, fairness and equity of AI systems. I hope you enjoy reading about this profound work.

Moreover, CMU remains dedicated to educating the next generation of AI practitioners and leaders who not only possess technical expertise but also a deep understanding of the ethical responsibilities associated with AI. Through our rigorous curriculum and hands-on research opportunities, we empower our students to become ethical stewards of AI, equipped to address complex challenges and drive positive change in society — now and for the future.

Martial Hebert
Dean, School of Computer Science

AI AND THE Election

SHERI HALL

Across the political spectrum, candidates running for election this year are using artificial intelligence to connect with voters and criticize their opponents.

Within hours of President Joe Biden announcing his reelection bid in April 2023, the Republican National Committee released a video ad imagining what would happen if Biden won the 2024 presidential campaign. In the top left corner of the ad, there was a note in a small font that announced the ad was "Built entirely with AI imagery." At least one Democratic Congressional candidate is using a generative AI robot to call potential voters.

Without a doubt, generative AI has tremendous potential to influence voters this election year and going forward. Researchers at Carnegie Mellon University's School of Computer Science are using this opportunity to educate voters on how to identify and evaluate AI-created content, and to better understand how AI influences elections.

> **"Compared with the last election, the ability to use artificial intelligence to create fake images, phone calls and stories online is much higher and much more sophisticated," said Kathleen M. Carley, director of the Center for Computational Analysis of Social and Organizational Systems in the School of Computer Science. "There is also more potential to use generative AI for good."**

This means informed voters must be able to identify materials created with AI and interpret the objectives of the creators. Carnegie Mellon researchers working to guide the public in understanding how AI is being used this election cycle are also striving to help policymakers build laws and rules around its use.

"Generative AI is a powerful technology," said **Hoda Heidari**, the K&L Gates Career Development Assistant Professor in Ethics and Computational Technologies and co-leader of CMU's Responsible AI initiative. "But we don't fully understand yet under what conditions it works as intended, and under what conditions it doesn't. We're looking into reliable, valid ways of evaluating this technology in a range of settings — and that includes how it is used in political campaigns."

Kathleen M. Carley, Director of the Center for Computational Analysis of Social and Organizational Systems

Hoda Heidari, the K&L Gates Career Development Assistant Professor in Ethics and Computational Technologies and co-leader of CMU's Responsible AI initiative.

Yonatan Bisk, Assistant Professor of Computer Science at the Language Technologies Institute

Types of Misuse

A major aim of research for SCS identifies the different ways generative AI can be misused, Heidari said. "This technology has been made publicly available without adequate guardrails," she said. "We have already seen numerous examples of how it can be used in malicious or otherwise questionable ways."

It's important to recognize that AI itself is not harmful, Carley explained; but people can use AI in harmful ways. Researchers categorize the misuse of generative AI into two main areas. The first is disinformation, which refers to inaccurate information purposely spread by people or robots with the intent to mislead. Secondly, there is misinformation, which is also inaccurate, but the sender doesn't realize it and they spread these falsehoods without malicious intent.

A special class of disinformation known as deepfakes involves content that looks and sounds real, but has been completely generated by a computer. Deepfakes come in the form of audio, video, images, memes and even news stories that contain disinformation but appear to be completely true. Researchers have found evidence that online actors are using AI, sometimes in the form of bots, to spread disinformation, including deepfakes, on a daily basis. For example, clear data shows that foreign entities are using generative AI on quasi-news and social media platforms to influence U.S. voters.

Carley, whose research combines cognitive science, social networks and computer science, is currently conducting a global investigation to analyze a type of deepfake called "pink slime," or hyper-localized fake news websites that contain completely made-up stories.

> "With the demise of many local newspapers, we've seen a sharp increase in these fake local news sites," said Carley.

"You will see a news website in two completely different towns that use the exact same story and words, but just change the photos and quotes as if they came from local officials. Our project is trying to identify these sites and then quantify how much misinformation they are spreading.

"Unfortunately, we have seen them have an impact in getting people to donate money to fake causes, and also in getting people to believe a candidate has more support than they actually do," she said.

This project is just one example of how CMU researchers are working to identify and quantify the effects of AI misuse — an important first step to ensuring that AI is used fairly and productively.

Will AI Affect Election Results?

Despite specific examples of misuse, it is extremely difficult to quantify AI's overall effect on elections and the democratic process because currently there is no concrete way to measure its cumulative impact over time.

"If you pick out any one instance of GenAI misuse, one could argue it is unlikely to have changed the outcome of the election," Heidari explained. "But over time and as instances accumulate, the technology can contribute to the erosion of trust in the democratic process. We may get to a point where many voters won't believe what they see or hear, simply relying on their gut feelings and emotions, instead of facts and reason, to decide how to vote."

> "Ideally, we want the democratic process to be a rigorous debate and exchange of information about candidates and their policies, rather than an uninformed majority vote."

Heidari compares the impact of generative AI on elections to how social media has led to declining mental health of adolescents over the past decade. "In both cases, we are dealing with diffuse harm," she said." With this type of harm, it's difficult to pinpoint whether any one interaction has tipped the scale. There aren't many identifiable victims, yet over time, we observe a huge impact."

Broadly, Heidari's research addresses issues of accountability and governance of AI, and specifically how to evaluate the negative impacts of AI tools on individuals and society. For instance, she studies a technique called red-teaming for GenAI, which involves stress-testing the model to assess the risk of it producing problematic content, such as misleading information about the electoral process.

"The results of the risk assessment can then be utilized to decide, 'How can I prevent harms by adding appropriate guardrails?'" Heidari said. These guardrails can be technological, such as prompt-based filters, or policy-based, such as guidelines and regulations surrounding the use of GenAI in elections.

Both are important in ensuring the responsible use of AI, said Heidari. "Both are important, and they work in complementary ways," she said.

Generative AI for Good

Although there is ample potential for the misuse of generative AI, researchers agree it's critically important to highlight the positive uses for the technology as well.

"The goal is to create robots that are actually helpful to society," said **Yonatan Bisk**, assistant professor of computer science at the Language Technologies Institute. Bisk's research focuses on how robots use language to communicate with humans.

There are many examples of generative AI content used for good, such as alerting communities when a wildfire is approaching. In elections, generative AI can help candidates reach voters who might otherwise be disenfranchised by translating campaign messages into different languages or condensing policy proposals into summaries that are easier to understand.

"The problem is, we're not seeing it used for good as much," Carley said. "I know a lot of political consultants are afraid to use generative AI because they view it as complicated, and they don't want to get it wrong."

Carly explained that this hesitancy tips the scales in the wrong direction because bad actors are using AI freely.

> **“For researchers, the question becomes, what brakes can we create that help people to use this technology in good ways?”**
> — YONATIN BISK

To get there, technology creators must first understand all of the applications — both good and bad — for the technologies they create. According to Bisk, this often happens through a process called threat-modeling, where researchers try to imagine unintended and malicious uses for the systems they build.

Understanding the potential threats of AI systems is an important first step in regulating them, said Bisk. “The interesting questions are who is inventing, who is legislating, who is regulating that space? My biggest goal is to ensure that our elected officials are well-educated on the implications of the projects they are funding, and are constructing the appropriate policies to regulate these systems.”

Recognizing and Evaluating AI content

As generative AI becomes more widespread, it’s critically important for voters to learn how to recognize and evaluate computer-generated content.

CMU’s Block Center for Technology and Society houses a research area on how to responsibly harness AI and analytics for social good. And CMU’s Center for Informed Democracy and Social-cybersecurity (IDeaS) houses a research area on how to detect, measure the impact of, and mitigate online harms that threaten democracy such as disinformation, hate and terrorism. These two centers have jointly created The Responsible Voter’s Guide to GenAI in Political Campaigning to help protect the integrity of the democratic process.

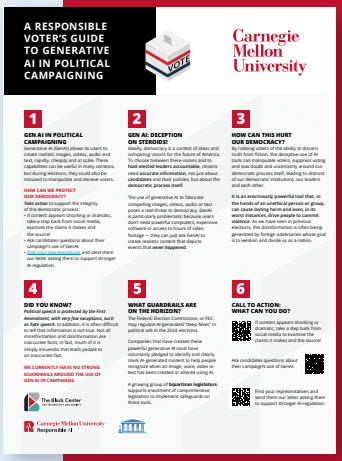
A RESPONSIBLE VOTER'S GUIDE TO GENERATIVE AI IN POLITICAL CAMPAIGNING

Carnegie Mellon University

1 GEN AI IN POLITICAL CAMPAIGNING

2 GEN AI: DECEPTION ON STEROIDS!

3 HOW CAN THIS HURT OUR DEMOCRACY?

4 DID YOU KNOW?

5 WHAT GUARDRAILS ARE ON THE HORIZON?

6 CALL TO ACTION: WHAT CAN YOU DO?

Watch the Block Center video discussing the Voter’s Guide.

> **“The biggest part we can play at the moment is increasing awareness,” Heidari said. “As AI researchers, we should make sure we don’t contribute to the hype, but instead, present a voice of reason.”**

There are currently no federal rules about using AI in political campaigns. The Federal Communications Commission has proposed requiring politicians to disclose the use of AI in TV and radio ads. At the same time, the Federal Election Commission proposed banning candidates from using AI to deliberately misrepresent opponents in political ads. So far, neither proposal has been enacted.

This leaves U.S. voters on their own in identifying and deciphering AI content this election cycle. CMU researchers have some tips for how to navigate this difficult landscape.

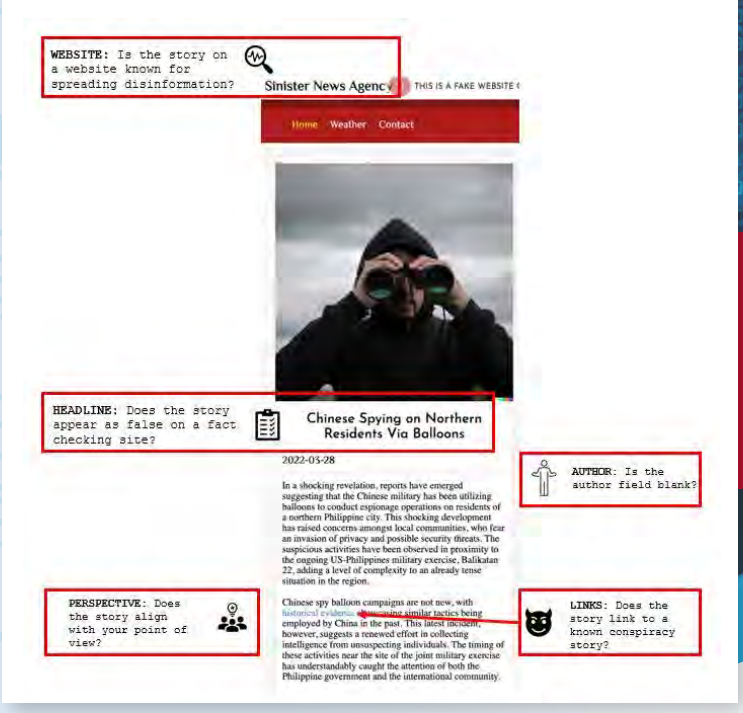

An example of a website intentionally spreading disinformation, with annotations highlighting red flags and offering critical thinking questions.

In addition to the Voter's Guide, the IDeaS center has also published guides on detecting disinformation and identifying content generated by AI language models.

While it can be difficult to identify malicious computer-created content, one strategy is to pay attention to your emotional reaction to the material, Carley explained, because people working to spread disinformation often attempt to play on people's emotions.

"If you're reading a story, and you find you're starting to get really excited or sad or angry, that emotional shift is probably because you're being played by the story," Carley said. "That's a sign you should take a break from the media."

When consuming media, it's also important to consider your own biases, noted Bisk. He points to the fact that people are more inclined to believe viewpoints that they agree with, and not believe viewpoints they disagree with. The same is true when evaluating whether content was generated by AI. "If you see a video of someone talking and you agree with what they are saying, you are more likely to think the video is real," he said. "If you disagree with what they are saying, you are more likely to think the video is a bot or a fake."

Thinking about the motive of the person or organization delivering the content is also useful, Bisk said.

"If I'm presented with a clip — whether it's audio or video — of a candidate making a nonsensical statement, I should ask, 'What is the incentive for that person to say those things? Does what you're seeing fall within the character of the candidate? Is it consistent with your model of who they are?" Candidates are only going to put out material they believe will get them elected, Bisk said. "If a particular clip seems out of character or designed to simply provoke outrage, it may not be real."

As always, it's important to consider the source of the information you are consuming, said Carley. "People tend to look at sites that are very similar, and get the same story from all of them, even though it may be inaccurate," she said. "That's why it's important to look at a variety of sources, even if you think you're going to hate them."

"Most importantly, people should be aware that misuse of generative AI is out there. "To voters, I would just say, 'Be careful,'" Carley said. "And to policymakers, try to use these new technologies for good." – KATHLEEN M. CARLEY

MORALITY AND ARTIFICIAL INTELLIGENCE

DANIELLE HOFFMAN

It might seem as if Artificial Intelligence is everywhere these days, and that its emergence is a recent phenomenon. But since the early days of computer science, when punch cards were being fed into room-sized mainframes, the concept and the importance of AI has shown up consistently in academia, industry and across our broader culture. Depending on the media you consume, you might be thrilled by the possibilities or terrified of the implications.

The implications go far beyond a monotone robot voice droning. "I'm sorry, Dave, I'm afraid I can't do that," as HAL, the famous rogue AI, says in "2001: A Space Odyssey." Modern AI encompasses everything from students using ChatGPT to write papers, to the Turkish media circulating convincing deepfake videos of political opponents, to a health care system using an AI trained on biased data, resulting in white patients being dramatically prioritized over black patients with the exact same level of illness.

The need for ethical boundaries around AI has never been more important.

Carnegie Mellon University, often credited as a "birthplace of Artificial Intelligence," is well equipped to help guide this conversation. As Anand Rao, a professor of applied data science and AI in the Heinz College of Information Systems and Public Policy said: "One of the reasons I came to CMU was the way the university is very deeply rooted in the technology, while at the same time looking clearly at the implications on society."

But how can CMU help make sure that the tech leaders of tomorrow will be aware of those implications? How can we ensure that the AI we build makes moral choices? And beyond that lies a greater question: What is morality when it comes to AI?

To help answer some of these questions, we turn to **Vincent Conitzer**, renowned artificial intelligence researcher and ethicist at CMU and co-author of "Moral AI and How We Get There" along with Jana Schaich Borg, a neuroscientist, and philosopher Walter Sinnott-Armstrong.

Vincent Conitzer, Professor of Computer Science and Director of the Foundations of Cooperative AI Lab (FOCAL).

Defining AI

American computer and cognitive scientist John McCarthy coined the term in 1955, calling it "the science and engineering of making intelligent machines." McCarthy went on to say that "intelligence is the computational part of the ability to achieve goals in the world." Meanwhile, AI pioneer Alan Turing's famous test states that if a machine can have a conversation with a human, and the human cannot distinguish if they are conversing with another human or with a machine, that machine has demonstrated human-level intelligence.

Today Turing and McCarthy would likely call our powerful computing machines and modern algorithms very intelligent indeed — yet we keep moving the goalposts of what constitutes true intelligence. Conitzer and his co-authors call this the AI effect. "Once we know how to write a computer program that can perform a task we once thought required intelligence, such as navigating mazes, we no longer think the computer program's solution to the task is really intelligence. Instead, it becomes just another algorithm."

Therefore, the authors of "Moral AI" use the definition provided by the US National Artificial Intelligence Initiative Act of 2020: "A machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations or decisions influencing real or virtual environments with sufficient reliability."

Anand Rao, Distinguished Services Professor of Applied Data Science and AI in CMU's Heinz College of Information Systems and Public Policy

Toward Making AI Fair

Anand Rao, Distinguished Services Professor of Applied Data Science and AI in CMU's Heinz College of Information Systems and Public Policy, worked in AI and business for 30 years before landing in academia, and thus, has had a front-row seat to moral AI discussions over the years, many centering on AI that will act ethically in difficult situations.

One common thought experiment involving ethical situations is the trolley problem. There are innumerable iterations, but the classic version is: there are five people tied to a train track, and a trolley is headed straight for them. You can flip a switch to divert the trolley to another track, but there is one person walking along that track. Do you knowingly kill one person to save five?

"Early on in the late '80s and '90s, no one really worried too much about ethics," Rao said. "If you brought up the trolley problem, they'd say 'that's a philosopher's problem, and not relevant to what we are doing.' The focus was much more around how do we plan efficiently to get the robot from point A to point B? I'm really happy it moves from this room to the other room, rather than worrying about whether it will harm someone."

Rao said it wasn't until 2014 that people started to get serious about ethical AI. That was the year the Oxford philosopher Nick Bostrom published his sobering polemic "Superintelligence: Paths, Dangers,

> **"A machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations or decisions influencing real or virtual environments with sufficient reliability."**
>
> **— US NATIONAL ARTIFICIAL INTELLIGENCE INITIATIVE ACT OF 2020**

Strategies," which landed on the New York Times' best seller list for science books and prompted Tesla and SpaceX founder Elon Musk to assert that artificial intelligence was potentially more dangerous than nuclear weapons.

The need to apply ethics to AI for engineers and other technical professionals that might not have received any training, studied ethics or fairness, or been given specific guidelines, is an important point.

A few years later, The Future of Life Institute, a nonprofit made up of academics and tech giants, recognized both the tremendous potential and alarming risk of AI and put together the Asilomar Conference on Beneficial AI. At this 2017 conference, they developed the 21 Asilomar AI Principles, one of the earliest and most influential sets of AI governance.

"Academically, we're going very deep in many of these areas," Rao said. "There's a lot of detail in a narrow sense. Now, let's say I'm just a company person building a model in a bank. For them, it is really daunting — they know how to build a model, now you come and tell them it needs to be 'fair.'"

In "Moral AI," Conitzer and his co-authors touch on the transition from establishing principles and philosophies to practical implications. They tell us that while these principles and commitments might sound good, ethical issues persist, and until we address some of the nuances of the process of creating AI, from inadequate communication to a lack of clear metrics, the issues aren't going away.

Teaching Moral AI at CMU

Rising amid the cluster of Beaux-Arts buildings that make up CMU's campus are the glass and steel Gates-Hillman Centers, where some of the world's top minds in computer science focus on ethics, fairness and AI.

During the academic year, Vincent Conitzer, one of those top minds, runs the Foundations of Cooperative AI Lab (FOCAL). He also spends his summers at Oxford University at the Institute for Ethics in AI. Like Rao, Conitzer uses the trolley problem as an example of the thorny complexity of puzzling out ethical quandaries.

"The point of the trolley problem is that it's not that easy. And if you think that something simple will do — I'm just going to save the maximum number of people — well, it's much more complicated than that. We're running into problems that have been studied in philosophy, except now with AI systems they're real and concrete. And it's not just ethics, but also, for example, philosophy of mind."

If it's a matter of sacrificing one person to save five, the problem seems relatively easy to solve. But when we add variations found in applying AI, things get trickier. What if you're driving a car, and you know that if you swerve to avoid the child in the road, you'll kill two adults? What's the moral choice there? And then, what if it's a computer driving that car? Can we give the AI tools to make these moral decisions, given that every scenario will have thousands of variations?

And of course, students need to want to tackle those issues. Conitzer tells a story about an Intro to AI course he taught for undergraduates, wherein he sprinkled ethical instruction and assignments among the technical work.

"Some students liked it," Conitzer said, "but in the middle of the semester we asked for some feedback, and a student said: 'Can we just cut out that ethics stuff?' It illustrates a little of the difficulty. In some ways technical work feels more satisfying to the student, but it doesn't make the ethics issues any less important."

Moral AI Across CMU

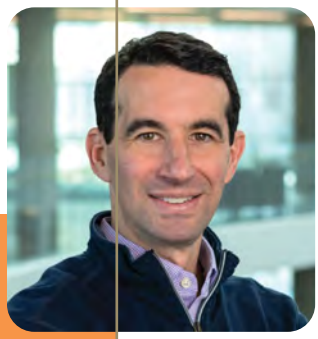

Jay Aronson, Professor of Science, Technology and Society in CMU's History Department

I want them to think about their place in the broader world, and what the implications of their intellectual efforts might mean."

— JAY ARONSON

Faculty members tackling the ethical questions of AI can be found in almost every department across campus: from Alex John London and Hoda Heidari's pioneering work in philosophy and computer science at the K&L Gates Initiative in Ethics and Computational Technologies; to the Block Center for Technology and Society, where Ramayya Krishnan, dean of the Heinz College of Information Systems and Public Policy, leads faculty in research around the future of work and responsible AI.

These questions are even addressed in CMU's History Department, where **Jay Aronson**, professor of science, technology and society thinks a lot about ethics and AI. He runs the Center for Human Rights Science and teaches an undergraduate course called "Killer Robots? The Ethics, Law and Politics of Drones and AI in War."

"I decided to teach it because it's interesting, but also, it was a little bit of a bait and switch," said Aronson. "I think a lot of CMU students are interested in autonomy and technology, and they hear the word ethics and think, 'Oh, I know what ethics is.' But when I get them in class, I spend the majority of time teaching them about international law and the potential moral and strategic hazards of using technology in war." The content is not just around technology, he adds, but around actively questioning that technology being used. The mere fact that we can develop these advanced technologies isn't reason enough to dive in; we need to consider the ramifications and decide whether we should.

"I want them to think about their place in the broader world, and what the implications of their intellectual efforts might mean," Aronson said. "Get them to think from a life-saving — rather than a life-taking — perspective. My students tell me that they crave opportunities to think through the implications of their work."

Of course, AI encompasses so much more than autonomous vehicles on the ground and drones in the air. In "Moral AI," Conitzer and his co-authors relay that there is "both bad news that gives us reason to be worried about some uses of AI, and good news that gives us reason to advocate for other uses."

That good news includes The Los Angeles Times' Quakebot algorithm, which warns its readers about California earthquakes more quickly and accurately than traditional media. The good news also includes the lives saved by a complicated kidney exchange network, the work of Tuomas Sandholm, Angel Jordan University Professor in CSD, made possible because of AI.

> **If AI could help humans make better moral judgments in even a few of these cases, it could benefit us tremendously.**
>
> **— VINCENT CONITZER**

AI might even save the planet, as researchers use AI systems to visualize future floods and wildfires, monitor forests, and improve decision-making about the climate.

Even with all this good news, as AI becomes more sophisticated, Conitzer raises some sobering possibilities: What if AI is used to control our physical environment? We're already using AI to plan tactical strikes; how much future control over cybersecurity or autonomous weapons systems will we hand over to AI? Yet, Conitzer remains cautiously optimistic, especially when it comes to not only reflecting on, but improving upon, human morality.

"Humans make important moral judgments," Conitzer points out. "Problems arise in decisions about which targets to attack in war, which criminal defendants to grant bail, whether to brake (and which way to turn) a car in an apparent emergency, and many other cases. If AI could help humans make better moral judgments in just a few of these cases, it could benefit us tremendously." ■

Moral AI
And How We Get There

Can we build and use AI ethically?

That's the timely topic of "Moral AI and How We Get There" by SCS Professor Vincent Conitzer, co-authored by Duke University professor of neuroscience Jana Schaich Borg and Duke University philosophy professor Walter Sinnott-Armstrong.

"Moral AI" resists painting apocalyptic scenarios or waxing poetic about utopias; instead, the authors take a more sensible approach, clearly laying out what we know about AI right now, the ways in which it is already being used for good and ill, and future decisions to make. The book is not just a primer on the challenges and opportunities of ethical AI, it also lays out concrete calls to action that technology, policy and business leaders of today could put into action tomorrow. The authors encourage careful reflection on the values we want AI to reflect, and suggest ways in which we can program AI to be more moral than us.

Though written by academics, the authors of "Moral AI" wanted the book to be accessible and wrote in a conversational style. Conitzer said the goal was to write a basic introduction to moral issues around AI so that anyone could pick it up and learn from it, regardless of their background in the field. ■

Scan to read more and purchase your own copy.

PEOPLE WITH AUTISM TURN TO CHATGPT FOR ADVICE ON WORKPLACE ISSUES

BYRON SPICE

CMU RESEARCH ASKS WHETHER THAT USE OF AI MAKES SENSE

A new Carnegie Mellon University study shows that many people with autism embrace ChatGPT and similar artificial intelligence tools for help and advice as they confront problems in their workplaces.

Andrew Begel, Associate Professor in S3D and the HCII

But the research team, led by the School of Computer Science's **Andrew Begel**, also found that such systems sometimes dispense questionable advice. And controversy remains within the autism community as to whether this use of chatbots is a good idea.

"What we found is there are people with autism who are already using ChatGPT to ask questions that we think ChatGPT is partly well-suited and partly poorly-suited for," said Begel, an associate professor in the Software and Societal Systems Department (S3D) and the Human-Computer Interaction Institute (HCII). "For instance, they might ask: 'How do I make friends at work?'"

Begel heads the VariAbility Lab, which seeks to develop workplaces where all people, including those with disabilities and who are neurodivergent, can successfully work together. Unemployment and underemployment are problems for as many as nine out of 10 adults with autism, and many workplaces don't have the resources to help employees with autism and their co-workers overcome social or communication problems as they arise.

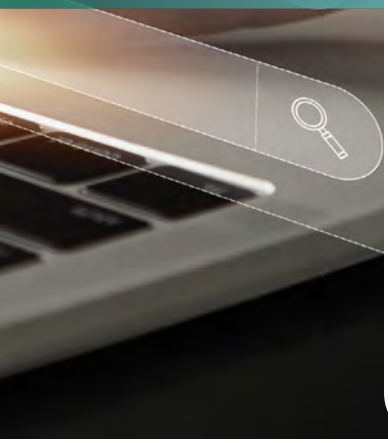

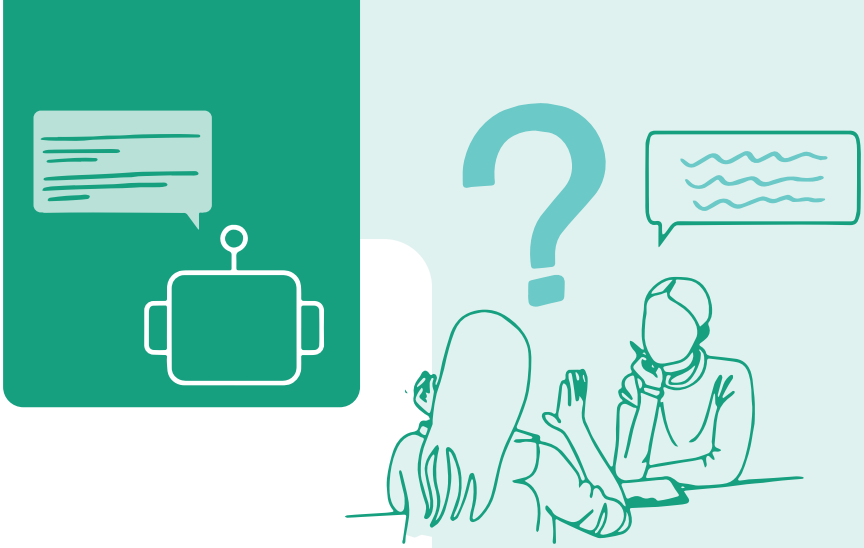

To better understand how large language models (LLMs) could be used to address this shortcoming, Begel and his team recruited 11 people with autism to test online advice from two sources — a chatbot based on OpenAI's GPT-4 and what looked to the participants like a second chatbot but was really a human.

Somewhat surprisingly, the users overwhelmingly preferred the real chatbot to the disguised counselor. It's not that the chatbot gave better advice, Begel said, but rather the way it dispensed that advice.

"The participants prioritized getting quick and easy-to-digest answers," Begel said.

The chatbot provided answers that were black and white, without a lot of subtlety and usually in the form of bullets. By contrast, the counselor often asked questions about what the user wanted to do or why they wanted to do it. Most users preferred not to engage in such back-and-forth, Begel said.

"I THINK, HONESTLY, WITH MY WORKPLACE, IT'S THE ONLY THING I TRUST BECAUSE NOT EVERY COMPANY OR BUSINESS IS INCLUSIVE."
— ANDREW BEGEL

Participants liked the concept of a chatbot. One explained: "I think, honestly, with my workplace, it's the only thing I trust because not every company or business is inclusive."

But when a professional who specializes in supporting job seekers with autism evaluated the answers, she found that some of the LLMs answers weren't helpful. For instance, when one user asked for advice on making friends, the chatbot suggested the user just walk up to people and start talking with them. The problem, of course, is that a person with autism usually doesn't feel comfortable doing that.

Patrick Harrington, Assistant Professor in the HCII

Sanika Moharana, Ph.D. student in the HCII

JiWoon Jang, Ph.D. student in the HCII

Results from the experiment were presented by first author and HCII Ph.D. student **JiWoong (Joon) Jang** at the Association for Computing Machinery's Conference on Human Factors in Computing Systems (CHI 2024) in Honolulu. In addition to Begel and Jang, co-authors include HCII Ph.D. student **Sanika Moharana** and **Patrick Carrington**, an assistant professor in the HCII.

It's possible that a chatbot trained specifically to address the problems of people with autism might be able to avoid dispensing bad advice, but not everyone in the autism community is likely to embrace it, Begel said. While some might see it as a practical tool for supporting autistic workers, others see it as yet another instance of expecting people whose brains work a bit differently than most people to accommodate everyone else.

"THERE'S THIS HUGE DEBATE OVER WHOSE PERSPECTIVES WE PRIVILEGE WHEN WE BUILD TECHNOLOGY WITHOUT TALKING TO PEOPLE."
— ANDREW BEGEL

"There's this huge debate over whose perspectives we privilege when we build technology without talking to people. Is this privileging the neurotypical perspective of 'This is how I want people with autism to behave in front of me?' Or is it privileging the person with autism's wishes that 'I want to behave the way I am,' or 'I want to get along and make sure others like me and don't hate me?'"

At heart, it's a question of whether people with autism are given a say in research that is intended to help them. It's also an issue explored in another CHI paper, on which Begel is a co-author with Naba Rizvi and other researchers at the University of California, San Diego. In that study, researchers analyzed 142 papers published between 2016 and 2022 on developing robots to help people with autism. They found that 90% of this human-robot interaction research did not include the perspectives of people with autism. One result, Begel said, was the development of a lot of assistive technology that people with autism didn't necessarily want, while some of their needs went unaddressed.

"We noticed, for instance, that most of the interactive robots designed for people with autism were nonhuman, such as dinosaurs or dogs," Begel said. "Are people with autism so deficient in their own humanity that they don't deserve humanoid robots?"

Technology can certainly contribute to a better understanding of how people with and without autism interact. For instance, Begel is collaborating with colleagues at the University of Maryland on a project using AI to analyze conversations between these two groups. The AI can help identify gaps in understanding by either or both of the speakers that could result in jokes falling flat or creating the perception that someone is being dishonest. Technology could also help speakers prevent or repair these conversational problems, Begel said, and the researchers are seeking input from a large group of people with autism to get their opinion on the kind of help they would like to see.

"We've built a video calling tool to which we've attached this AI," said Begel, who has also developed an Autism Advisory Board to ensure that people with autism have a say in which projects his lab should pursue. "One possible intervention might be a button on this tool that says, 'Sorry, I didn't hear you. Can you please repeat your question?' when I don't feel like saying that out loud. Or maybe there's a button that says, 'I don't understand.' Or even a tool that could summarize the meeting agenda so you can help orient your teammates when you say, 'I'd like to go back to the first topic we spoke about.'" ■

PHOTOGRAPHY BY **ELAN MIZRAHI**

2024 COMMENCEMENT

Congratulations to the School of Computer Science Class of 2024 on your graduation! As you celebrate this remarkable achievement, we are thrilled to reflect on the dedication and innovation you've demonstrated throughout your journey. Your hard work has prepared you to tackle the challenges of tomorrow, and we can't wait to see all that you will accomplish in the years ahead. Remember, you are now part of a vibrant community of thinkers and creators — go out and make your mark on the world! ■

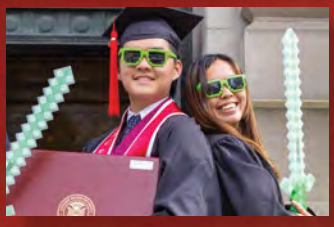

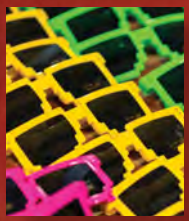


CMU
ABHINAYA
ΗΩΧ
CLASS OF 2024

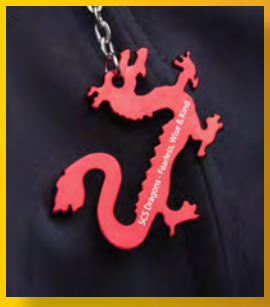
KS Dragons - Fearless, Wise & Kind

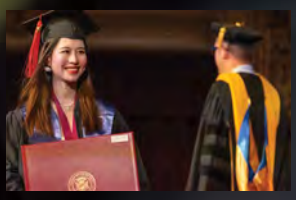


CARNEGIE MELLON UNIVERSITY

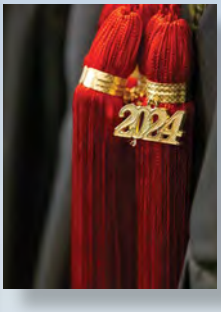
2024


Carnegie Mellon University
Congratulations
to the
Class of 2024!

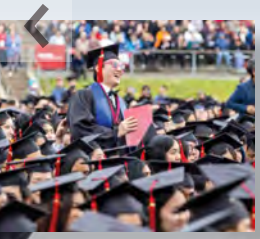

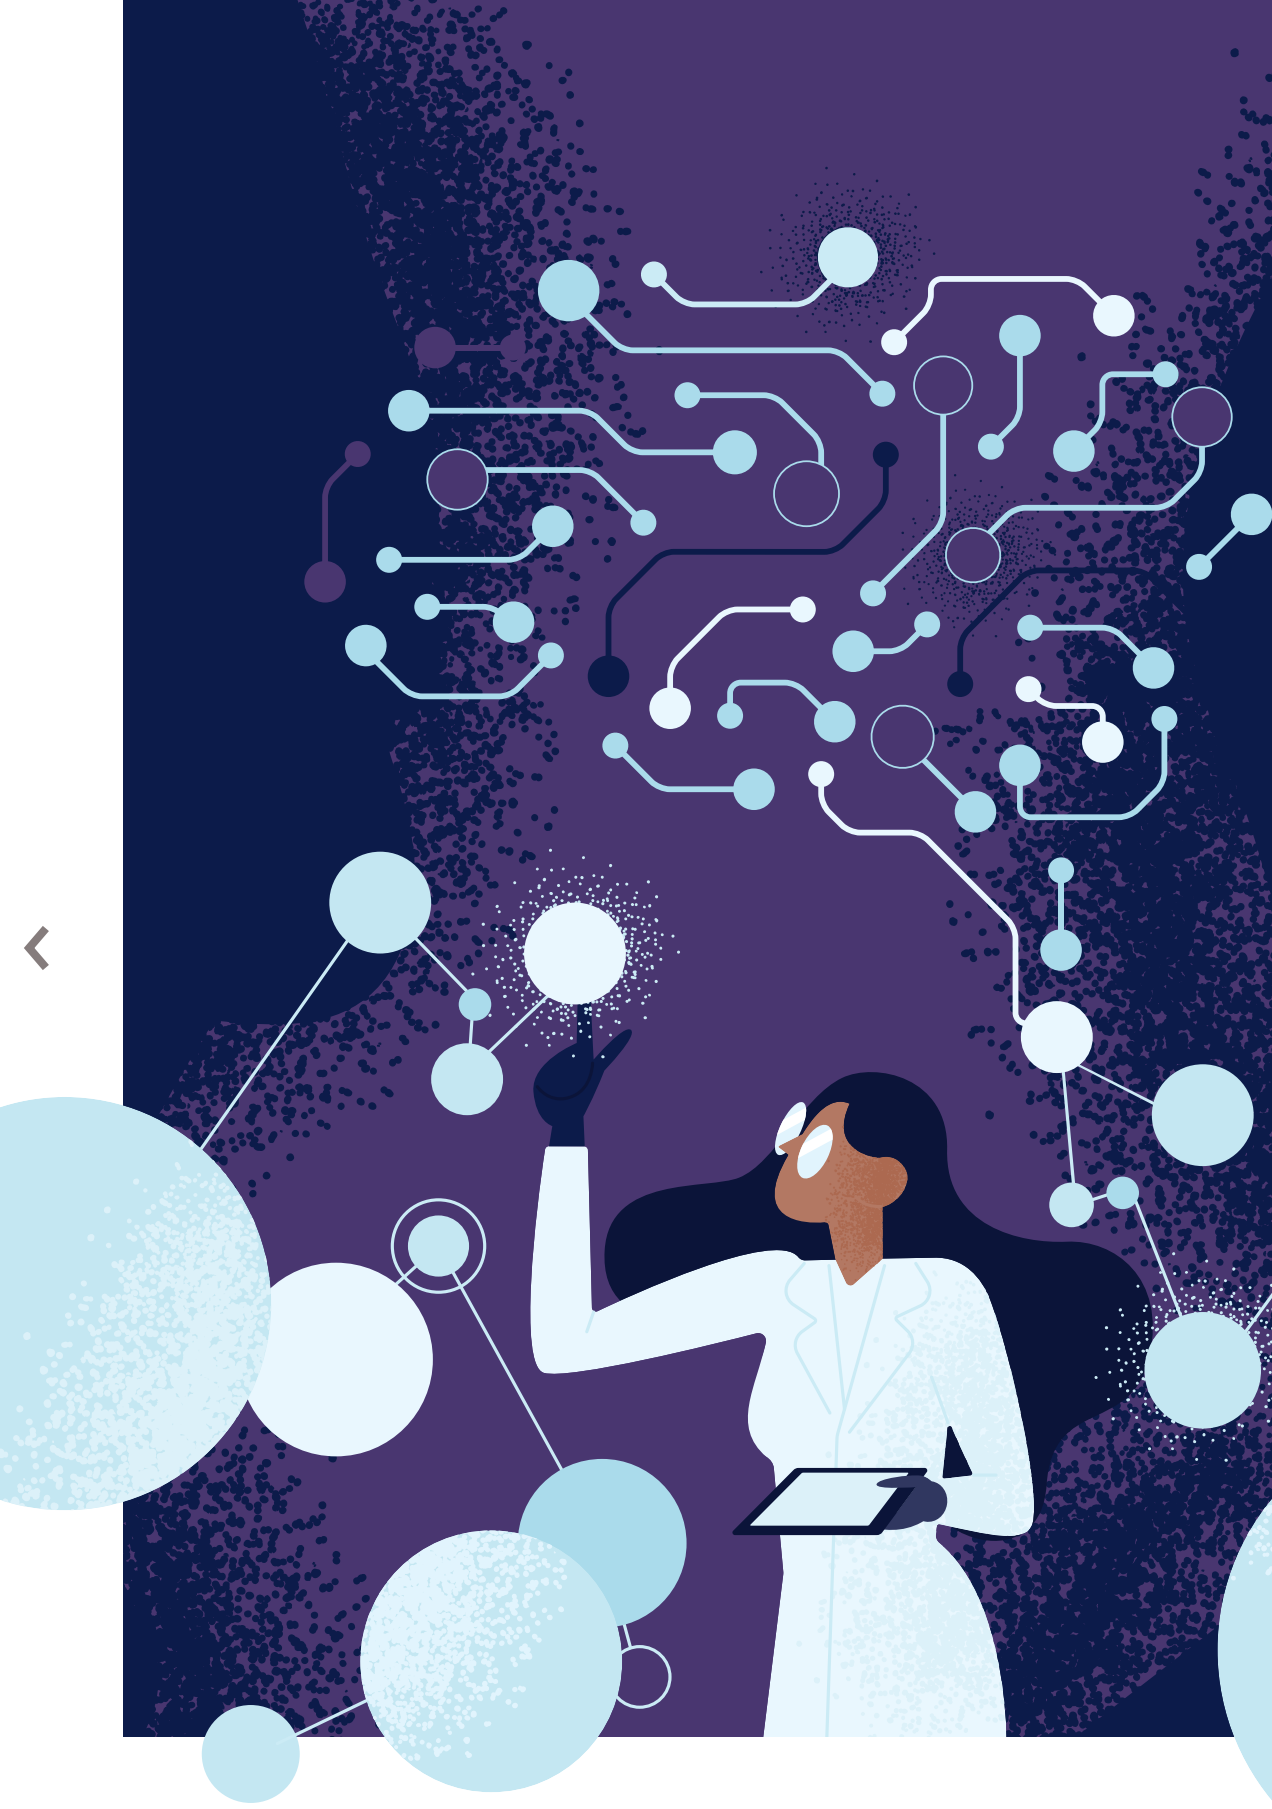

CHRIS QUIRK

PARTNERING AI AND SCIENCE

AS THE SCIENTIFIC METHOD EVOLVED, IT HAS SERVED HUMANITY WELL: OBSERVE, DEVELOP A HYPOTHESIS, RUN AN EXPERIMENT, ANALYZE THE RESULTS, AND THEN REPEAT. FROM THE ROOT OF THAT BASIC FORMULA, COUNTLESS DISCOVERIES HAVE BLOSSOMED BECAUSE OF THE POWER INHERENT IN THE SCIENTIFIC METHOD'S RIGOR.

In recent years, computer scientists have amplified that power, such that artificial intelligence now can increase the speed of the very deliberative scientific method by orders of magnitude. Scientists in the School of Computer Science and across CMU are taking full advantage. Attending new techniques and methods, these researchers are creating questions about how AI is changing the way we think about science, and how AI may offer new approaches to scientific inquiry.

As an AI researcher and leading figure in the development of computer vision, **Martial Hebert**, dean of the School of Computer Science, has watched with interest as scientists in myriad disciplines have put AI to work in their labs and inquiries. "AI is able to do a massive amount of data processing that no human could do in any reasonable amount of time, thereby providing conclusions and information that help scientists target what their next area of focus might be," said Hebert. "It is absolutely transforming the way we look at scientific discovery."

Martial Hebert, Dean of SCS

Barnabás Póczos, Assistant Professor in MLD

"For hundreds of years people did science the same way. Now we have this AI companion which can help us make scientific discoveries faster, quicker, cheaper and better," said **Barnabás Póczos**, an assistant professor in the Machine Learning Department (MLD). "At some point when people write scientific papers, they may credit AI as a collaborator, because they may not have been able to make these scientific discoveries alone."

> THIS NEEDS TO BE A COLLABORATIVE ENDEAVOR TO FIND AI METHODS THAT WORK IN THE CONTEXT OF THE SCIENTIFIC METHOD.

Rachel Mandelbaum, CMU Professor of Astrophysics and Cosmology

But how is AI assisting? First and foremost, researchers are taking advantage of AI's raw computational muscle. **Rachel Mandelbaum**, CMU professor of astrophysics and cosmology, is working on one of the most exciting projects in her field, the Legacy Survey of Space and Time (LSST). Powered by a 3,200-megapixel camera now being installed at the new Vera C. Rubin Observatory in Cerro Pachón, Chile, slated to open next year, the LSST will map the southern skies in a way astronomers could only dream previously. The survey will look at every point in the sky about a thousand times over 10 years. "It's going to be like a 10-year color movie of the entire visible sky," said Mandelbaum.

The LSST will find about 10 million objects in the sky that have changed — every single night. "So every night there will be 10 million alerts for that. It is going to produce an enormous data set," Mandelbaum explained. Sorting all that data will demand AI-powered data processing. "Astronomers will have to figure out very quickly whether a change is something interesting that we should be following up with a different telescope, or whether it is a plane that flew overhead, and not so interesting."

AI is indispensable to the success of the LSST and extracting salient information out of the almost boundless trove of data. "There are a lot of exciting opportunities for using AI for scientific discovery in astronomy and other fields, but we can't use AI out of the box," Mandelbaum said. "This needs to be a collaborative endeavor to find AI methods that work in the context of the scientific method."

As an ardent player of Go, the ancient strategy board game with an almost incomprehensible number of possible game positions, **Jeff Schneider**, a research professor in the Robotics Institute (RI), recalls the watershed moment when a computer defeated international master Lee Sedol in 2016. "I'm a pretty poor Go player, but I truly thought that I would never see a computer that could beat me," said Schneider. "For me it was mind blowing, not only that it could beat me, but beat the world champion."

Today, Schneider's work parlays AI in a bid to make nuclear fusion a reality. Fusion is commonly dubbed the holy grail of energy production, because of its ability to produce nearly limitless power at low cost. Fusion differs from nuclear fission, which splits atoms and is the process used in current nuclear energy technology. Rather, fusion melds atoms together, generating massive amounts of energy from tiny amounts of fuel. No greenhouse gasses are released, only helium, and the waste product tritium that fusion consumes and expels is far less radioactive and dangerous than the waste fission produces.

Jeff Schneider, Assistant Professor in RI

Hosein Mohimani, Associate Professor in the Lane Computational Biology Department

Abhinav Adduri , Ph.D. Student in Computational Biology

True nuclear fusion, long a dream of the scientific community, is the energy system of the sun and the stars, welding atoms together in their cores at tremendous temperatures and pressures. Recreating that environment on Earth is a supreme challenge to say the least, and one that Jeff Schneider has been working on diligently. "This is big science," he said. If harnessed, fusion could supply the energy needs of the entire planet, without endangering the climate.

Schneider is seeking a fusion solution using a tokamak, a device that suspends a torus — a doughnut-shaped plasma of hydrogen — within it, using powerful magnetic fields. "We're trying to get it to these extreme temperatures and pressures, when there isn't anything that could physically confine the plasma," he explained. "The magnetic field holds the plasma in place." The challenge is that the dynamics of achieving that are complex and unstable, and there are billions of variables involved.

> **IT'S REALLY CHANGING THE WAY SCIENCE IS DONE.**
> — JEFF SCHNEIDER

Given the high cost of tokamaks, researchers have to plan carefully and use their precious time allotments with extreme efficiency. Schneider and his colleagues now use a modified large language model, which has been loaded with decades worth of data and results from prior experiments run on the tokamak, to help fashion and sharpen their own tests. "It's helping us build models and get insights using machine learning, and it's really changing the way science is done," Schneider said. "It's a much more effective process and things are improving a lot."

In the field of health sciences, as the number of dangerous infections that are resistant to many antibiotics rises — like Methicillin-resistant Staphylococcus aureus (MRSA) — the urgency increases to identify new drugs to combat them. **Hosein Mohimani**, associate professor in the Ray and Stephanie Lane Computational Biology Department, and **Abhinav Adduri**, a Ph.D. student in Computational Biology, have found success with the help of AI, targeting a particularly nasty bug known as Candida auris.

C. auris, a yeast resistant to many antifungals, preys particularly on the less healthy, and thus, often wreaks its havoc in healthcare facilities. In essence, antimicrobial resistance occurs when pathogens develop cross-resistance to the means that drugs employ to destroy them. "Maybe this drug attacks the cell wall. Then, if a bacteria can protect its cell wall, that renders a whole class of antimicrobials inactive," Adduri said. "We then need new antimicrobials with novel mechanisms of action." The high volume methods of AI virtual experimentation that the researchers built increase their odds of finding gaps in the armor of pathogens.

Mohimani and Adduri created AI tools to massively speed up drug research, and as a result, more quickly find antibiotics with a high probability of effectiveness. They rely on AI to suggest possible designs for drugs, in the process helping them focus their search in the most fruitful directions, eliminating vast swaths of likely fallow areas. "We generate data on organisms that give us hints about the capacity of certain microbes to produce interesting antibiotic molecules," explained Mohimani. "Then we use AI methods to analyze that data and figure out what is the most promising."

> **WE USE AI METHODS TO ANALYZE THAT DATA AND FIGURE OUT WHAT IS THE MOST PROMISING.**
>
> — HOSEIN MOHIMANI

Mohimani sees nothing unusual in employing AI as a kind of scientific partner. "Ever since the discovery of penicillin, a lot of antibiotics that we are currently using in our medicine have been discovered by random luck," he said. "We are here to discover novel antibiotics in a systematic way, by using high throughput methods to generate a lot of data, and then using AI algorithms to analyze it." The human in the loop, so to speak, is still critical to fabricating the drugs. Mohimani and his colleagues enhance the probability of success for the compounds they identify by using existing natural materials as a base, rather than letting the algorithm construct its own antibiotic from scratch. "It makes them easier to synthesize," Adduri explained. "AI algorithms don't really have a great notion of the pragmatic aspects of chemical synthesis ability."

Andreas Pfenning, a genomic researcher, created a process for finding disease cures that is a kind of feedback loop, closing a virtuous circle of human-AI collaboration. Pfenning, an associate professor in the Ray and Stephanie Lane Computational Biology Department, works in comparative genomics finding new avenues to identify important information among the dense thicket of animal genomes, particularly neural functions, which could lead to therapies for Alzheimer's disease or addiction.

Andreas Pfenning, Associate Professor in the Ray and Stephanie Lane Computational Biology Department

Faced with the daunting amount of information in genomes, researchers looking for key segments have, quite logically, sought shortcuts. "Since the 1960s, if people are trying to identify whether or not a certain part of the genome is important or functional, they have used models of how the individual nucleotides of the genome evolve," said Pfenning. Researchers typically focused their attention on genetic information passed on from generation to generation, as that which is preserved in a genome is likely important. Further researchers built mathematical models that presumed that if you are seeking a part of a genome related to a disease, you would find it in a similar location in the genome across different species.

But what if those principles didn't always hold? Rather than taking a mathematical approach, Pfenning applied machine learning, which unearthed connections in unlikely sections of genomes. In doing so, Pfenning identified previously unknown parts of the genome related to vocal learning and the capacity to mimic sounds. "This is the foundation of human speech," Pfenning said.

> WHAT I ARGUE FOR IS A TIGHTER INTEGRATION OF AI AND SCIENCE, RATHER THAN AI COMING IN AND TAKING OVER THE SCIENTIFIC DISCIPLINE OF BIOLOGY.
>
> — ANDREAS PFENNING

The traditional models were, in fact, missing important information. "A common theme in a lot of our work is trying to reframe traditional biological problems. The data that we collect can provide input to our algorithms. And then the algorithms can make predictions that we can go into the laboratory to test," he said. "We still need to conduct scientific experimentation, and what I argue for is a tighter integration of AI and science, rather than AI coming in and taking over the scientific discipline of biology."

Despite the accumulation of benefits, there remain questions and caveats around AI and research. A white paper published last year by The Royal Society, Britain's leading scientific academy, found in its review of the use of AI in science and engineering that while the increasing adoption of AI applications in STEM fields shows great potential, policy and ethical questions, like the risks of potential discrimination inherent in data systems, have to be part of the broader conversation.

Others raise concerns about the "black box" nature of AI — the algorithms cannot at this time tell us how they arrived at a particular answer, or why they took a specific action. Hebert's outlook remains pragmatic. "If you use AI to drive cars, you want to understand exactly what it's doing so you can prove that it's doing the right thing," he said "In biology, if you have the AI predicting that particular key configuration would behave in a certain way, you would do experiments to verify that."

Tom Mitchell, Founders University Professor in MLD.

There is also the problem of what **Tom Mitchell**, Founders University Professor in the Machine Learning Department, calls "monothink," where the utility of successful AI tools can become its own enemy. "AlphaFold is the go-to tool for protein folding. And we all know it's not perfect, but it's way better than anything else that's out there," said Mitchell. "But if every scientist in the field is using the same AI tool, it's kind of like if every researcher in a Ph.D. group were getting advice from the same advisor. If that advice is sometimes wrong, then everybody's misled."

In addition, due to the fragility of AI, and its propensity to produce hallucinations — false or fabricated results — Mitchell says we must recalibrate the way we think about computers. "For the first five decades of computer science, up until 2000 or so, it really was true that if a commercial system released some software, they kind of guaranteed that it worked. There were proofs of correctness, and we got into a mode of thinking of computers as something that would be perfect," he said. "Now we're starting to use computers in a very different way. ChatGPT doesn't always work in the sense that it gives the perfect answer, just like people. We're starting to use computers in a way much more similar to how we use our human colleagues, rather than as oracles that are infallible." ■

> WE'RE STARTING TO USE COMPUTERS IN A WAY MUCH MORE SIMILAR TO HOW WE USE OUR HUMAN COLLEAGUES, RATHER THAN AS ORACLES THAT ARE INFALLIBLE.
>
> — TOM MITCHELL

SARAH BODEN

Thinking About AI & Work

HOW AI CAN MAKE WORK MORE EFFICIENT AND CREATIVE WHEN FUTURE USERS HELP WITH ITS DESIGN

Artificial Intelligence won't make humans obsolete in the workforce,

says **Nik Martelaro**, an assistant professor in the Human-Computer Interaction Institute (HCII).

Though powerful, algorithms are rigid, and at some point, technology always fails. As an example, he points to the iPhone, which was first introduced in 2007.

"This is a fantastic product. But still, it screws up all the time," said Martelaro, whose research focuses on interaction design and interaction with autonomous systems.

Instead, Martelaro sees AI's role in the future of work as a collaborator with humans. We bring dexterity and cognitive flexibility, while AI delivers speed and precision.

HCII researchers explore how humans and computers communicate and interact, including some who are researching how AI can help people do their jobs. A few examples include Françeska Xhakaj's work creating tools that allow teachers to better identify which students need help, and Adam Perer's exploration of how machine learning might aid physicians in clinical decisions. Also, Haiyi Zhu researches how AI can address systemic bias to improve the well-being of gig economy workers, such as Uber drivers and Instacart shoppers.

But for AI to make work more efficient, the people already doing these jobs must be integral to designing the technology, said **Sarah Fox**, assistant professor in the HCII and Director of the Tech Solidarity Lab.

Sarah Fox, Assistant Professor in the HCII and Director of the Tech Solidarity Lab

> "In situations where AI is poorly designed, it could degrade or displace work. And that's very concerning."

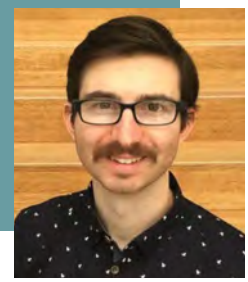

Nik Martelaro, Assistant Professor in the HCII.

Jodi Forlizzi, Herbert A. Simon Professor in Computer Science and HCI

For example, **Jodi Forlizzi**, the Herbert A. Simon Professor in Computer Science and HCI, leads a project that illustrates the importance of centering workers and their experience in the design of AI, in her partnership with UNITE HERE!, the largest hospitality union in the U.S. Other collaborators for the project include Fox, as well as researchers at University of Illinois Urbana-Champaign, Michigan State University and New Mexico State University.

Hotel housekeepers reported frustration with a new algorithmic management software that determined the order in which rooms were to be cleaned. According to the employees, they'd sometimes be scheduled to clean several checkout rooms in a row. These rooms must be fully turned over for new guests, a more labor-intensive assignment than a room that has the same guest for multiple nights.

Fox explains that when housekeepers decide the order of their workloads, they say they are more efficient and tired less quickly. A housekeeper might accomplish this by spacing the cleaning of checkout rooms with other rooms with the same guests staying multiple nights, which requires a lighter touch. Additionally, workers might opt to make all the beds, and then clean the sinks, and then vacuum all the carpets for an entire block of rooms.

The housekeepers reported the platform seemed to be designed by someone who assumed that changing the sheets or cleaning the shower at their own home was the same as working in a hotel, which is not the case. The latter requires a high degree of precision and cleanliness.

"What [the housekeepers] impressed upon us was they believed the designers and technologists who developed this original software didn't appreciate or recognize their work as expert work," said Fox.

A more ideal scenario with AI design will include workers from the jump. That's why Fox collaborates with the AFL-CIO Technology Institute, the Transport Workers Union (TWU) and the Amalgamated Transit Union (ATU) to investigate how AI can be used on city buses.

The Technology Institute was established by the AFL-CIO, the national labor federation representing some 12.5 million workers from 60 unions, including TWU and ATU.

Collectively, both unions represent more than 300,000 U.S. workers, and their members include bus operators, train operators, conductors, trackworkers, car cleaners, mechanics and other frontline transit workers.

Martelaro, along with CMU research associate **Hunter Akridge** and Ph.D. student **Alice Tang** collaborate on the project.

Adding autonomous vehicle technology to city buses may seem simple since there's ample research into using AI in four-door sedans. But that is not the case as public transport buses serve a different purpose. At 40 feet long, these accessible vehicles stop every few blocks, even during rush hour traffic.

To include the operators in the design of the AI, roughly 50 TWU and ATU drivers across North America recorded voice diaries of their workdays. In these diaries, the transit workers shared how AI could help them drive safer and increase awareness of their surroundings.

"I always felt like driving is like martial arts," said one driver in Portland, Oregon. "You never know when a car is gonna pull out in front of you. You never know when a little kid is going to pull out in front of you."

Drivers also gave suggestions, such as equipping buses with blind spot detection and improvements to lane departure warnings, which monitor when a bus gets too close to the edge of the lane. Fox noted that the current warning systems used in public transit do not factor in the size of buses, making the frequent alarms distracting.

But drivers need more than assistance with technical skills. These civil servants must also care for the welfare of their riders, who include children, those with disabilities and people without homes.

A recent HCII paper examining what would be lost if public transportation was fully automated explored the importance of passenger management. For example, one transit driver recalled an incident in which a frightened girl was being bothered by a visibly intoxicated male passenger.

"It just didn't sit right with me," said the driver, who intervened by instructing the girl to move and sit near the front door.

When the bus arrived at the girl's stop, the driver only opened the front doors but kept the rear exit locked. The intoxicated man was seated toward the back; therefore, he couldn't follow the girl off the bus. Then, the driver drove away and didn't let the man exit for another eight blocks despite his yelling threats.

If there hadn't been a human driver, that girl may not have been able to leave the bus safely. However, some visionaries foresee a future that advocates for fully automated public transportation. It's already happening to a limited extent in Scotland, though buses are still required to have a human safety driver aboard in case of emergencies.

Part of the AFL-CIO's motivation for working with HCII is that they don't want AI to compromise the safety and dignity of transit work. In 2022, Amanda Ballantyne, director of the AFL-CIO Technology Institute, discussed her organization's collaboration with CMU. She said "the current failed model of corporate-dominated research" too often fuels inequality and leaves workers behind.

That's a legitimate concern as history is replete with examples of new labor-saving technologies leading to substantial reductions in employment for workers, said **Brian Kovak**, Professor of Economics and Public Policy in Carnegie Mellon University's Heinz College of Information Systems and Public Policy. However, Kovak notes that it's too early to know what tasks AI might be best at, and therefore, it's hard to predict who is at risk of being displaced.

Brian Kovak, Professor of Economics and Public Policy in the Heinz College of Information Systems and Public Policy.

Aniket Kittur, Professor in the HCII

David Chuan-en Lin, Ph.D. student in the HCII

Hyeonsu Kang, Ph.D. student in the HCII

> **Rather than replacing jobs entirely, AI will likely change what types of skills are required for different occupations," said Kovak.**

That's also how Fox sees things potentially playing out with public transportation and points to some parallels with aviation. While AI has alleviated some stress and strain on pilots, these jobs now require more technical training and expertise.

Beyond having algorithms help with the tedious or labor-intensive aspects of a job, or even designing technology to make work safer, AI can also encourage people to be more creative.

One project out of Nik Martelaro's lab is a collaboration with the Toyota Research Institute, in which an AI program refines a designer's sketches. As the person draws new lines of a car, the AI swiftly produces a rendering based on those inputs. This creates a volley of ideas between the person and AI, where the designer receives suggestions from the algorithm. While the AI does not autonomously design new vehicles, it does help to generate concepts that a person might not have thought of on their own.

The collaborators for this project include HCII professor **Aniket Kittur**, CMU Ph.D. students **David Chuan-en Lin** and **Hyeonsu Kang**, as well as Yan-Ying Chen and Matt Hong at Toyota.

Martelaro's work also explores the use of collaborative AI in the field of mechanical engineering. The software to generate the geometry for new designs already exists, but the technology doesn't know what the user wants to create. So HCII assistant professor **Kenneth Holstein** and Ph.D. student **Frederic Gmeiner**, along with Martelaro, work toward creating AI that asks engineers questions to help them think more critically before drafting a new project: What is this machine's purpose? How much weight should it support? What materials are being used?

Writing a checklist that augments with each new task requires understanding and awareness of one's own thought process.

"A lot of people don't actually do that metacognitive step, they just jump in," said Martelaro.

The concept of a software platform asking users questions about their objectives has been around for some time. In 1996, Microsoft Office introduced "Clippy," the animated paperclip intended to help people with tasks like formatting documents or writing letters. However, Clippy's legacy is one of eye rolls and mockery, as the design was too prescriptive and ultimately unhelpful.

The HCII project's goal is more open-ended — Gmeiner, Holstein and Martelaro want to help designers to think more flexibly.

Kenneth Holstein, Assistant Professor in the HCII

Frederic Gmeiner, Ph.D. student in the HCII

> **I'm not here to automate your job," says Martelaro. "I'm here to make you better at your job." — Nik Martelaro** ■

Wikibench Allows Wikipedians To Improve AI Evaluation Datasets

Every two seconds, someone makes an edit to a Wikipedia page. And while many of those edits help curate and develop the internet's encyclopedia, some of them aren't made in good faith.

MARYLEE WILLIAMS

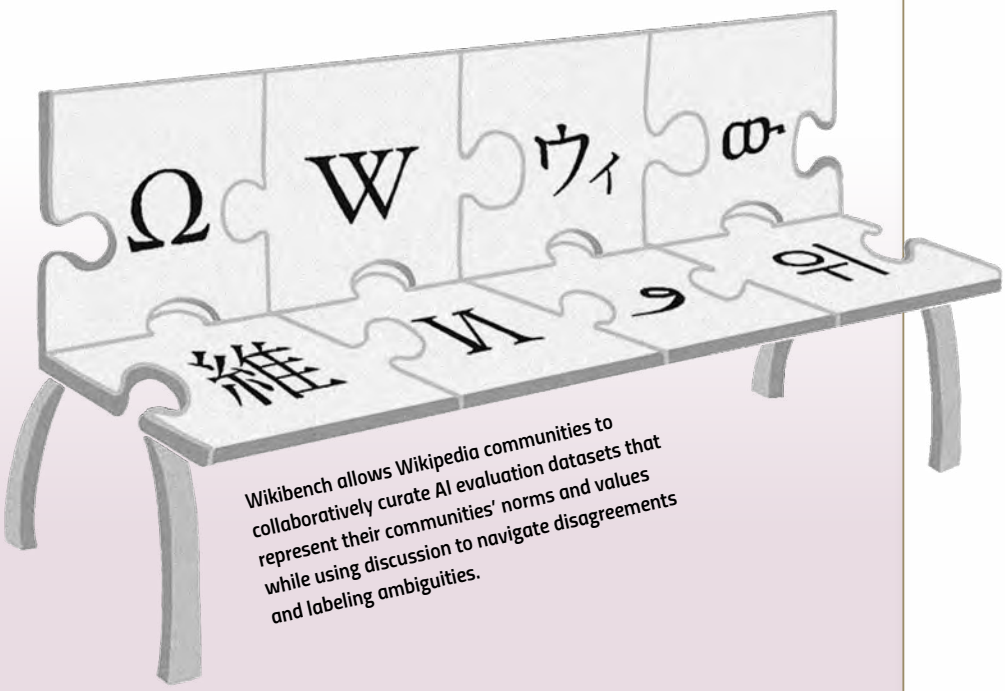

Wikibench allows Wikipedia communities to collaboratively curate AI evaluation datasets that represent their communities' norms and values while using discussion to navigate disagreements and labeling ambiguities.

To monitor recent changes — accepting the beneficial and reverting the malicious — Wikipedia patrollers have turned to artificial intelligence to predict whether an edit was vandalism or done in bad faith. While such tools can help the volunteers who edit and monitor the pages (known as Wikipedians) quickly identify and address issues, a method for Wikipedia communities to evaluate how well these tools fit their specific needs has yet to be created. The tools also don't account for varying content-moderation norms and values across different languages used on Wikipedia.

School of Computer Science researchers have developed a solution to this problem in the form of Wikibench, which allows Wikipedia communities to collaboratively curate AI evaluation datasets that represent their communities' norms and values while using discussion to navigate disagreements and labeling ambiguities. The research team from the Human-Computer Interaction Institute included Ph.D. student **Tzu-Sheng Kuo**; Assistant Professor **Kenneth Holstein**; **Haiyi Zhu**, the Daniel P. Siewiorek Associate Professor; Assistant Professor **Sherry Tongshuang Wu**; and **Meng-Hsin Wu**, a Master of Human-Computer Interaction alumna. It also includes Aaron Halfaker, a researcher at Microsoft; Jiwoo Kim of Columbia University; and Zirui Cheng of Tsinghua University.

Researchers chose Wikipedia because a demonstrated need existed among patrollers to both evaluate and train the AI systems deployed in their individual communities.

"Wikibench demonstrates that it's possible to have an approach that allows the people impacted by an AI system to be the people who evaluate the system," Kuo said.

The name Wikibench comes from the concept of benchmarking, the systematic comparison of an AI model's performance against a ground truth standard. Unlike typical AI benchmarking, Wikibench acknowledges that what accounts for good performance on socially constructed tasks like content moderation can be highly community specific.

"The goal is to create space for community members to find consensus, where possible, but also to identify areas where different community members may have genuine disagreements," said Holstein, whose Co-Augmentation, Learning and AI Lab (CoALA) studies how AI systems are designed and used by workers.

Researchers deployed Wikibench on English Wikipedia. To lower the barrier to participation and seamlessly integrate it into existing workflows, the team developed a simple plug-in that allowed Wikipedians to add or discuss labels while patrolling edits. Consider, for example, an AI model that flags an edit as vandalism. While a patroller using Wikibench might label this edit as damaging because it doesn't have a source, they may also believe the edit was done in good faith because the person wasn't trying to vandalize the page. Wikibench can also be used to curate other types of datasets that involve different labels.

Along with labeling the edits, Wikibench also creates a campaign page that allows the public to see the dataset and build consensus through discussions about disagreements. This space helps participants clarify ambiguities in labeling, potentially leading to changes in their labels. If they don't reach a solution, Wikibench preserves the information about their disagreement.

"Disagreements are OK, and some of the disagreements represent genuine differences in perspective," Holstein said. "But in other cases, when community members have opportunities to discuss seeming disagreements, they realize that they don't actually disagree." ■

Wikibench was presented at the Association for Computing Machinery Conference on Human Factors in Computing Systems (CHI 2024) in Honolulu. Learn more about the project and its future on the Wikibench research page.

FROM THE [does] Compute PODCAST

AI AND CLIMATE

Q+A *with* Priya Donti

Interview by Steph Stricklen

Priya Donti (CS 2022) is an assistant professor and the Silverman (1968) Family Career Professor at MIT EECS and LIDS. Donti is the co-founder and chair of Climate Change AI, a global nonprofit initiative to catalyze impartial work at the intersection of climate change and machine learning.

How did you get your start in this field?

I have known for quite some time that I wanted to work on climate. My high school biology teacher turned our class into a climate and sustainability class. I got motivated to work on climate as an issue of human well-being and equity that we would need to solve quickly. But I did not know exactly how I was going to work on it.

I went to college thinking I would be a material scientist and try to discover how to create better solar panels. But I really fell in love with my computer science classes and was a little bit confused because at that time I had no idea how computer science and climate fit together. But toward the end of undergrad, I stumbled upon some work from researchers at the University of Southampton, which really made a case that AI and machine learning would be critical for integrating renewables into power grids and managing all the variability in power production that comes from the variations in weather. And so, I got really excited about this idea, that computer science was going to be a critical part of decarbonizing power grids and just haven't looked back since then.

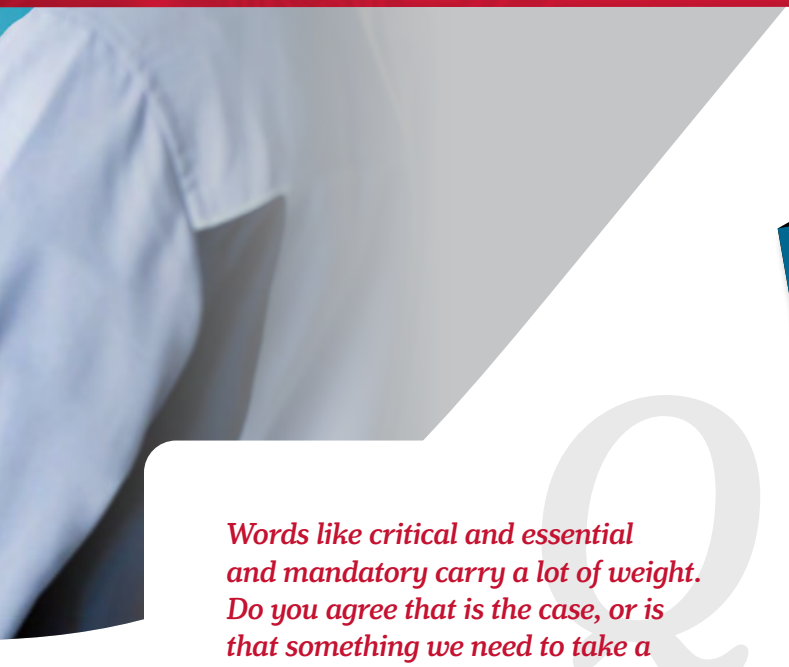

Words like critical and essential and mandatory carry a lot of weight. Do you agree that is the case, or is that something we need to take a higher altitude look at where we're going next?

I would say when it comes to the use of AI for climate action, there are places where it is critical, there are places where it's helpful, and there are places where it is inapplicable or inappropriate. And so, we really do have to distinguish, making sure that we are leveraging these tools in places where they are the right fit.

When it comes to power grids specifically though, I will say "critical" in that there are a couple of things that are happening as we move to next-generation power grids with large amounts of renewables. We have to manage the power grid at greater speed to deal with the real-time variation in how much power is being supplied into the system because solar power is varying based on the weather, wind power is varying based on the weather, and our existing optimization and control techniques where we write down the equations of the power grid and explicitly solve over them are not able to solve fast enough for that to happen.

We also need to manage the grid at greater scale because we have more devices coming online, not just distributed solar and wind power, but also electric vehicles, batteries, and similarly doing this in a top-down, physics optimization way isn't working.

On the power grid, a large amount of data is being collected through sensors. That provides an opportunity to manage the grid in a much more dynamic and data-driven way. And so, when we think about analyzing patterns automatically in large amounts of data, well, that's exactly what machine learning is.

What in your field are you most excited about right now?

I am really excited about the directions for machine learning to contribute to optimization and control of power grids. So, machine learning on power grids has been used for a very long time. The process of, for example, forecasting electricity demand to figure out how to balance the grid used to be done using rule-based systems. It's been common that you would use machine learning to do this by just learning from past data how to predict future electricity demand.

But where there has been a harder time making inroads are those places where you're actively optimizing the system. Because when it comes to a forecast, you can have multiple methods. You can look at all of them. You can figure out which one you like. You can figure out which one you trust. When you're optimizing a system, you are deciding about what to do on the system and it changes the state of the system going forward. You might break the system in the process. So, there's a lot of care that needs to be taken in that decision. I'm really excited to see that there actually is interest in thinking about this in power grids and that there has been a bigger mobilization of not just researchers, but also power grid operators and software providers who are coming together to try to address this.

What are the ways that AI is critical in solving the climate change problem?

AI is being used in a variety of ways, from helping us to take large streams of data like satellite imagery and turn it into actionable information that can be used for climate change-related decision making. For example, a coalition of entities called Climate Trace, which is using satellite imagery and on-the-ground sensor data to provide an independent third party estimate of where the greenhouse gas emissions in the world are coming from at the facility level as an input to the UN climate change negotiation. Information transparency is the goal.

One thing we haven't touched on is the role of AI in accelerating certain time intensive simulations like climate models, energy models or city planning models that model how wind is flowing through a city to figure out the overall efficiency of the city. These physical models are expensive to run. There are many uses of machine learning to better speed those things up to downscale the output so that they can give you more relevant information at the city level rather than just the country or region level. There are lots of ways, from providing information to automatically optimizing systems to facilitating science and engineering that AI is playing a role.

There are some places where I do think AI is critical. For example, a decarbonized power grid without using large-scale data-driven analysis. And similarly, colleagues have argued things like large-scale biodiversity monitoring. Those are two examples where taking large amounts of sensors and turning that into biodiversity information or optimizing and controlling a system at the scale of the power grid in a way that facilitates renewables. That's where — from a technical perspective — I'd say AI and machine learning are critical.

In places where AI is speeding things up, that could have otherwise happened at a larger time scale, that is also another form of criticality. Some of these things like getting that better battery by accelerating experimentation or getting a better climate prediction at a particular city, for example, maybe there are other ways to have done that that could have taken more time, but maybe AI speeds up the ability to do that in a way that deals with the criticality of moving quickly.

It sounds like the line between where AI is helpful and critical is thinner than the line between where AI is critical and absolutely inappropriate.

One way in which AI can be not wholly inappropriate, but maybe not the right first step can be in situations where there exists another step that you should take before leveraging AI. For example, optimizing the heating and cooling systems in buildings. You can put all sorts of fancy automation on your heating and cooling system. But if you haven't insulated the building in the first place, you've missed a critical, less high-tech way to do something very good. But then also, your automated system might be miscalibrated because then when you put the insulation in, it has to relearn how to optimize the building in that context. So, there are places where AI is helpful, but it's not necessarily the first step.

There are other places where I would say that AI is a Band-aid solution and should be — and it doesn't mean it's a bad solution. But maybe the real solution lies elsewhere. For example, there's a lot of really cool work that's looking at how to take corporate financial disclosures and mine climate-relevant information from those to help to better figure out what corporate action on this should look like, and to shape regulation. And that's a great application because it basically says this is the information we have, and we have to leverage it in some way. But if the longer-term idea is we want to use corporate disclosures to be able to analyze information at scale, maybe asking corporations to give their information in a structured format rather than text is the longer-term solution. So, figuring out where AI is being used to compensate for the fact that we have a system that is imperfect, and doing that, but then realizing what it means to make that system better rather than just nightly, without thinking about it, digging into the AI side as the actual perfect solution.

And then a place I would say AI is truly inappropriate is where it is being used to pretend that a decision is objective when really a decision is value-laden. When it comes to climate related policymaking, there are lots of hard decisions we're going to have to make in terms of how we spend our money, with who are the winners and losers of how we spend that money. And putting a bunch of data into an AI system and expecting it to spit out what the optimal policy is, is something that people have pitched to me before but is a way to basically shunt a lot of the decisions to actually the biases that exist in the underlying data rather than facing them head on. AI can maybe provide an assist to that. I talked about information transparency, for example, but you can't expect an AI system end-to-end to kind of make a big policy decision for you. ■

For more on AI and Climate Change, listen to the full interview with Priya Donti on SCS's Does Compute podcast.

INTERVIEW BY
KEVIN O'CONNELL

Mark Stehlik Announces Retirement, *Sort of.*

Mark Stehlik, University Teaching Professor, Assistant Dean for Outreach, Director of the CSD Undergraduate Program, co-founder of CS Academy and distinguished figure in the School of Computer Science and in the broad field of computer science, has announced his plans to retire at the end of the 2024-25 academic year. Sort of retire, that is. As a beloved faculty member, advisor, researcher and leader in computer science education, Stehlik's retirement gives those of us at SCS who have been influenced by him the chance to reflect on a few of the highpoints and the stories of his profound career.

Are you announcing your retirement officially, or what are we calling this?

STEHLIK: So, I'm doing this "phased retirement" thing. This is my last full-time year because I needed to set a date to stop advising. I didn't take on any new cohorts after the entry year 2021. And there's two purposes for me being half-time next year. One is to mentor those remaining students to get them to the finish line, either next December (2025) or by next May (2026). The other is to be here for the person who's going to replace me.

Your webpage states your preferred title is "Education Engineer." Tell me about the significance of that moniker, especially with as many advisees as you've had.

STEHLIK: The idea in teaching is to figure out where a student is and to try to make them better. But in advising, it takes on a whole new schema, because it's "let me figure out why your parental relationship is broken. What can we do about that? How you can be insulated from that and still thrive here? Let's figure out what courses you can take but not kill you by overloading you." I'm an experimentalist, which is sort of being an engineer. A student comes to me and tells me they're going to take six courses this semester. And I'm like, okay, let me see how well you did in the past couple of semesters and what kind of courses you took. If they did all right, I'll give them a sixth course and then if it turns out it doesn't work, well, then we're going to drop them back. So, there's a sense of tinkering with things. I don't know all the answers for everyone, but I want to find out what your answers are. And so the way to do that involves some bit of engineering, trying to understand the problem and then constructing an edifice that actually solves that problem for that student.

After choosing Pace University over MIT and others (more on that later), you chose to pursue a Ph.D. in Computer Science at CMU. But before you finished you took a teaching position. How did that come about?

STEHLIK: It was the end of my first year, middle of my second year, around that time. I was not excited about the dynamic programming research I was doing. My advisor left to go to Bell Labs, which was the Google of its day. I was advisor-less, mostly rudderless, and not really happy for the first time in my life at school going back to kindergarten. Something was totally wrong and needed to change and I wasn't sure what it was.

At that time, a faculty member had to leave unexpectedly, and they needed someone to teach Introduction to Programming. I had taught calculus for people pursuing their MBAs while working on Wall Street. And it was like teaching calculus to my dad, which was a wonderful way to begin teaching.

John McDermott was the acting department head at the time. I put up my hand and said, "I think I can do this." And I think his response was, "Yeah, well, we don't really have anybody else, so sure." The story has a happy ending because at the end of the semester after faculty course evaluations, he calls me into his office and says, "Not only do you seem to like this, but you're also actually good at it." I had gotten one of the highest faculty course evaluations that a CS faculty member ever had received, mostly because I gave a -expletive-.

Tell me about when computer science as a major was transitioning to SCS becoming its own school? How did you get to be the program director of SCS' undergraduate program?

STEHLIK: We had the Robotics Institute in 1979, which was a standalone department. CSD had been in existence since 1965 — one of the first Ph.D.-granting programs in the world. So, we had these two departments, and we wanted to become a school. Part of the price of "school hood" was you had to own an undergraduate program. Prior to that time (1988) if you did a CS degree here, you did it as an applied math major.

So, the first 211 course, developed by Mary Shaw and a whole bunch of other people, was Fundamentals of Computer Science. It was followed by 212, which is Introduction to Functional Programming. Mary Shaw and company put together the first undergraduate curriculum that was then going to be joint with the Mathematics Department. We didn't think we could do this all right away, and because at the time we didn't want the administrative role of dealing with incoming first-years. We didn't know how to do any of that stuff and the Math Department did.

Nico Haberman, who was the first dean of SCS, sent out a thing saying they needed a director for the program. They interviewed a bunch of people, and they didn't like anybody. So, Nico sends me mail and says, "Some of us think you might be good for the position, but you didn't apply. How come?" And I said, "Well, because teaching full-time is what I do." And he said, "So, what could we do to make you think about this?" And I thought about it for a while. I was a teacher. That's what I want to be. That's what I am. Finally, I said, "If you let me continue to teach, I'll consider doing this other administrative thing that I think I'm not going to like, but we'll see what happens." So again, in a moment that is etched permanently in my brain, Nico looks at me and says, "So you're telling me that if we just let you work harder, you'll be okay with this?" And I'm like, no. Maybe the way I thought about this was if you let me be what I think I am, I'm willing to do this other thing. But anyway, we parted company, and I was the first program director for undergraduate education.

The last major milestone we should talk about here is the formation of CS Academy with David Kosbie.

STEHLIK: David was an ex-Ph.D. student like me. He ended up working at Apple and Microsoft, and then moved back to Pittsburgh and was doing a high school teaching gig. He and I met for lunch at what was then Lulu's Noodles, right on Craig Street. He was clearly interested in teaching. I knew him back when we had played intramural football together as Ph.D. students. He was the quarterback. I was an offensive lineman.

As always, there's a story behind it.

Anyway, he wanted to get good computer science material out to the world. So, I hired him as a teaching faculty and we co-taught courses. We wanted to get 112 out into the world. And for 15 years we've been trying to do that. But at some point, we realized that the world wasn't ready for 112. And we realize that when David, in the Fall of 2017, went to a local convening of AP and other CS teachers at Hampton High School in the North Hills. About 40 teachers

were talking about all sorts of computer science curriculum stuff that was available, Scratch and other random things. One of the teachers put up his hand and said, "We're not trained computer scientists, you are. We don't know what's good."

I found this out when I taught an AP course and was handed a syllabus that was just awful. The teacher didn't know it was awful. I knew it was awful because I was the chief reader for AP, and I taught that material here. So, the teacher pretty much said, "You're at Carnegie Mellon. That's a pretty good place. Why don't you make something? It'll probably be good and then we can use it." And so literally David walks into my office and we sketch that in red.

(Editor's note: At this point Stehlik points to a scrawl of red (pictured below) on the whiteboard of his office outlining the genesis of CS Academy.)

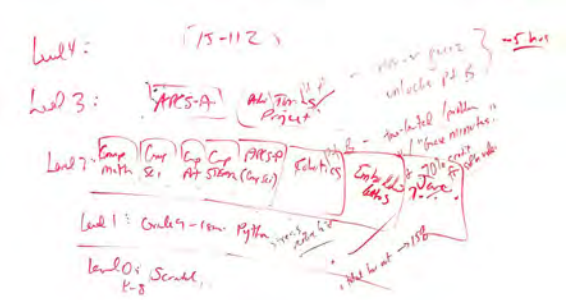

That is from November 2017. And it has levels, right? Levels 0, 1, 2, 3 and 4. Because what we realized is that 112 was our goal, but we couldn't start there. That was level 4. We had to flip it. In order to get to level 4, we had to begin at level zero. That's Scratch. That's the place where the students are at. Level 0 would be for kindergarten through 8th grade. We're not going to touch Python until we get to level 1 for ninth graders.

Additionally, we made it auto-graded, because the teachers aren't always experts. We also made it graphics-based, because that's compelling for kids. It's ninth grade algebra-ready, not even algebra completed. So, we thought we'd keep the math low in order to keep the engagement high. We started with that level 1 course. We called it CS1. We started with 14 schools, 13 of which were from that Hampton High School experience, all of whom are still with us six years later, plus one school in Rwanda at the Agahozo Shalom Refugee Village. It was created by the parents of one of my advisees for children of the Rwandan genocide. They had these kids and they wanted to do this. David and I both thought it would be a really cool thing. And our tagline in some sense for years later was, if it'll run in Rwanda, it'll run anywhere. Just give us an internet connection and some piece of $#%! computers and you can do this. And so they did.

We sent some of our own students there. David walked into our student team who was helping create CS Academy and said, "We've been asked to maybe deploy in Rwanda, but we'll only do it if we can support it. How many of you would be willing to go and spend a month, if necessary, this summer in Rwanda to help them?"

Every single kid in that room raised their hand. ■

While CS Academy has now been consumed by more than 400,000 users, all for free, by nearly 13,000 teachers across 20,000+ classrooms, we will return to this story, and all of Mark Stehlik's stories throughout the upcoming year. In a career that has lasted at CMU for this long, Stehlik has a wealth of stories. Some may even be fit to print. Some may provide answers to long-standing questions, such as: Why was Stehlik's time in Qatar cut short? Does Stehlik wear his kilt at home? Is it true that he taught a course on the finer points of the use of profanity in an academic setting?

For the answer to these questions and more, The LINK will provide an in-depth article in the next issue. Also pay close attention to events, talks and celebrations being planned throughout the year, as we approach Mark's actual retirement date. Whenever that date might actually be.

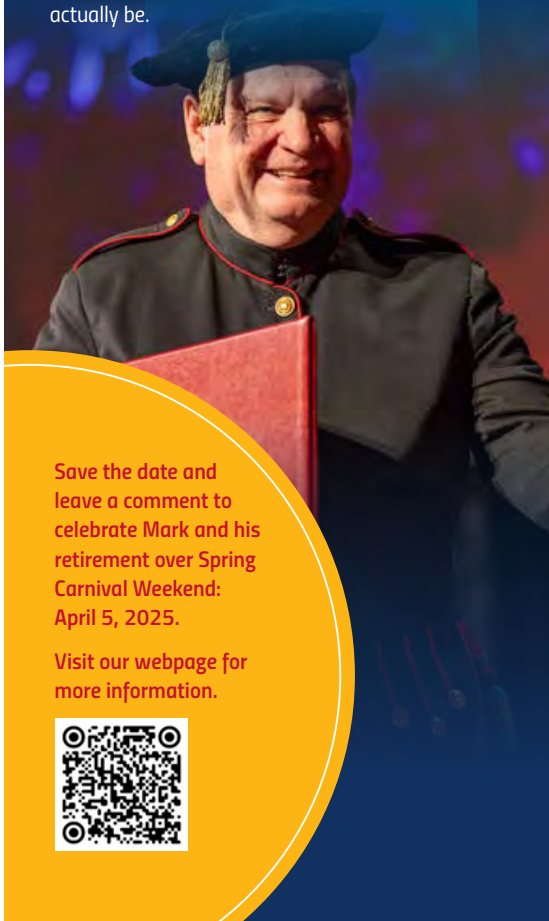

Save the date and leave a comment to celebrate Mark and his retirement over Spring Carnival Weekend: April 5, 2025.

Visit our webpage for more information.

Stehlik Leaves **His**

Mark

CMU launched its undergraduate computer science degree program in 1989, graduating its first 75 undergrads in 1992. Mark Stehlik advised them all.

Stehlik has advised close to 4,000 students in Pittsburgh and Doha, Qatar, helping them navigate the occasionally rough waters between high school and college graduation. But Mark's dedication doesn't stop there. He has:

- Trained close to a thousand AP Computer Science high school teachers;
- Led C++ summer workshops as part of an NSF-sponsored program that helped equalize the gender balance in computer science at CMU;
- Co-founded CMU CS Academy, which provides free online computer science curricula for high school and middle school classrooms that has been consumed by almost 400,000 students.

We will honor Mark at Spring Carnival in 2025. Make a donation at tinyurl.com/MarkStehlikFund

Mark's SCS legacy is indelible, but his retirement is imminent. **To celebrate his career and his impact SCS has created the Mark Stehlik Endowment for CS Academy.** Donations to the fund will support the initiative to ensure that CS Academy exists in perpetuity.

> **Like so many people who passed through SCS during Mark's long tenure, I owe him more than I could ever hope to pay back in one lifetime. I think it's only fitting to honor his impact, to pay forward his endless support for others, by building a fund that will provide ongoing support for his vision of a CS education program accessible to anyone, anywhere.**
>
> **— Jonathan Betz (SCS 1999)**

FOR MORE INFORMATION:

Jenny Belardi, *Chief Advancement Officer*
412-268-8810 ■ jbelardi@andrew.cmu.edu

Carnegie Mellon University
School of Computer Science

From left to right: U.S. Secretary of Commerce Gina Raimondo, CMU President Farnam Jahanian, Arm CEO Rene Haas, Keio University Vice President for Research Masayuki Amagai, Microsoft Executive Vice President and Chief Marketing Officer Takeshi Numoto and Japanese Minister of Education, Culture, Sports, Science and Technology Moriyama Masahito. Photo: U.S. Department of Commerce.

KELLY SAAVEDRA

CMU JOINS $110M U.S.-JAPAN PARTNERSHIP TO ACCELERATE AI INNOVATION

Carnegie Mellon University and Keio University announced they will join forces with industry partners to boost AI-focused research and workforce development in the U.S. and Japan. The partnership is one of two new university partnerships between the two countries in AI announced in Washington, D.C., at an event hosted by U.S. Secretary of Commerce Gina Raimondo.

The collaboration joins two universities with outstanding AI programs and forward-looking leaders with leading technology companies committed to providing funding and resources aimed at solving real-world problems.

CMU President Farnam Jahanian was in Washington, D.C., for the signing ceremony held in the Department of Commerce's Research Library.

"This new partnership is global in its scope and single-minded in its purpose to advance AI research and impact," Jahanian said. "It's an opportunity that is tailor-made for Carnegie Mellon University, and possible thanks to the leadership of Secretary Raimondo and Ambassador Emanuel. I am immensely proud of CMU faculty and students whose talents and brilliance have inspired this partnership, and, together, we are grateful for the chance to work closely with our partners at Keio University, Arm, Microsoft and SoftBank Group, as well as with companies across Japan, to usher in a new era of global innovation in AI."

CMU has been increasingly called upon for its AI expertise by public and private sector partners since the birth of the field. As a world hub for talent in realizing the transformative potential of AI and ensuring its safe and secure use for the benefit of humankind, CMU's AI efforts are intentionally interdisciplinary.

"Pennsylvania is a global leader in AI development and innovation — thanks in large part to leading institutions like Carnegie Mellon University that have long embraced the promise of new technologies and harnessed their power to transform the way people live and work," said Pennsylvania Governor Josh Shapiro. "Artificial Intelligence is already impacting every sector of our economy — and government leaders need to lean into its innovation to adapt to the rapidly changing technology market ethically and responsibly. That's the approach we've taken in Pennsylvania and why the commonwealth has partnered with Carnegie Mellon to develop best

practices for governing generative AI. This new partnership between Carnegie Mellon University and Keio University will build on Pennsylvania's and the United States' global leadership in AI technology — and will empower our workforce, lean into innovation and capitalize on economic opportunity."

The work of the partnership will center on specific research themes, including multimodal and multi-lingual learning, AI for robots, autonomous AI symbiosis with humans, life sciences, and AI for scientific discovery. The partnership is expected to create interdisciplinary collaborations across academic departments and include joint research projects, workshops and knowledge-sharing seminars between CMU and Keio researchers.

"Keio University has taken a leading role in Japan in AI research as well as in robotics and image processing," said Kohei Itoh, president of Keio University. "We are very excited to be moving on to the next stage in our research by forming a US-Japan collaboration with Carnegie Mellon University, a globally renowned research institute in the field of AI."

The partnership between Carnegie Mellon and Keio universities will be supported by $110 million in combined private sector investment from several American companies and a consortium of nine Japanese companies. CMU will receive funding and collaboration from Arm, Microsoft and SoftBank Group.

"Establishing AI research collaborations and industry partnerships on topics that align so closely with the research that's underway at CMU are essential for making progress in advancing the best possible impacts of AI on people's lives," said Martial Hebert, dean of CMU's School of Computer Science. "I'm excited for our faculty, researchers and students to start working with colleagues at Keio University who share our vision for AI development in some important areas of computer science and AI research."

The agreement between CMU and Keio positions CMU faculty, researchers and students to work with their counterparts at Keio to leverage shared strengths in AI and computer science to promote innovation, strengthen private sector performance, and advance their respective academic and research missions. ■

BRIAN THORNTON

CARNEGIE MELLON BREAKS GROUND FOR RICHARD KING MELLON HALL OF SCIENCES

Innovative Facility Unites Foundational Science, Computer Science and Art

A rendering of the new Richard King Mellon Hall of Sciences.

For more on the groundbreaking and details about the new building, read the full story on the CMU News website.

CMU administrators, faculty, staff and students; elected officials and community partners; and project supporters gathered on Friday, April 12, 2024 for the Richard King Mellon Hall of Sciences groundbreaking.

Carnegie Mellon University broke ground recently for construction of its Richard King Mellon Hall of Sciences, a 338,900 square-foot building for education, research, art and creativity unlike any other on its campuses.

Anchoring the corner of Forbes Avenue and Craig Street in Pittsburgh's Oakland neighborhood, the building will enliven a prominent thoroughfare and serve as the new western gateway to campus. The unique facility will house departments from the School of Computer Science and the Mellon College of Science, as well as the Institute for Contemporary Art Pittsburgh, creating an atmosphere that facilitates collaboration across disciplines to advance data-driven research and creative inquiry. Construction is expected to be completed in 2027.

"Modern science is at a crossroads, and the future of U.S. innovation is inextricably linked to our nation's capacity to drive scientific and technological advances," CMU President Farnam Jahanian said. "The Richard King Mellon Hall of Sciences will position Carnegie Mellon University to excel at this crossroads and usher in new levels of collaboration while creating new paradigms for discovery and accelerating our institutional mission of leveraging knowledge for society's gain. We are deeply committed to advancing this vision and grateful to our supporters for their transformative investments in our future."

Construction of the Hall of Sciences has been made possible by a $75 million lead grant from the Richard King Mellon Foundation, part of a historic $150 million partnership that was the largest single grant in the foundation's history.

"Building on CMU's historic work in the life sciences and the physical sciences, the new Hall of Sciences will be a place where scientists here and around the world collaborate to advance humanity. Here, they will make discoveries and uncover insights and advancements that will improve the human condition," said Sam Reiman, trustee and director of the Richard King Mellon Foundation. "And — since this is CMU — we know the academic work practiced in this building also will generate new life-science companies. And they, too, will improve and save lives."

The Hall of Sciences is the cornerstone of CMU's future of science initiative, a decade-long effort to revolutionize and accelerate research and education, leading to innovations and breakthroughs that will benefit humankind. The Biological Sciences and Chemistry departments, along with the interdisciplinary Neuroscience Institute, will have new homes in the building, along with parts of the Ray and Stephanie Lane Computational Biology Department, Language Technologies Institute and Machine Learning Department from SCS. By co-locating faculty and students from both MCS and SCS, the project will facilitate new research partnerships that will combine AI, data analytics and foundational sciences to answer critical questions and solve previously unsolvable challenges.

"This will be a hub, an intersection, of all the greatness of Carnegie Mellon. The sciences, engineering, computer science and fine arts, intersecting in one place," said CMU Trustee Ray Lane. "I expect the next 30 years to be about the sciences, computation, engineering and fine arts coming together into this hall to make meaningful changes in the human condition. It's going to be a great future." ■

AARON AUPPERLEE

CMU TechBridge: **CODING BOOTCAMP**

The School of Computer Science has launched the CMU TechBridge Coding Bootcamp to provide access to computer science education and career opportunities for high school (or equivalent) graduates.

The bootcamp will extend the reach of CMU's computer science education offerings beyond its traditional undergraduate and graduate programs. CMU has partnered with TalentSprint to implement the program to help aspiring professionals build the skills required for launching and advancing their tech careers.

"CMU is dedicated to making a meaningful impact on society by extending its expertise to those who need it most," said Ram Konduru, director of executive and professional education at SCS. "CMU TechBridge is tailored to meet the increasing demand for high-trust, high-impact coding bootcamps for entry-level tech professionals. We are partnering with leading ed tech firm TalentSprint to offer this program, which will impart skills necessary for participants to secure and succeed in tech jobs."

There are more than 377,000 job openings for software and tech roles annually, and these careers attract premium salaries, according to data from the U.S. Bureau of Labor Statistics. Further, the tech industry is moving toward skill-based hiring in addition to hiring traditional computer science degree holders.

This coding bootcamp will give participants industry-relevant skills for entry-level tech jobs without requiring a traditional computer science degree, removing barriers for people who might not otherwise have the time or resources. The program covers introductory computer science concepts with an emphasis on problem solving. Participants will learn to write code, build front-end interfaces and deploy programs using tools currently used in the industry.

This full-time, hands-on learning program will be delivered over four months. A hybrid program, CMU TechBridge Coding Bootcamp will include online classes led by instructors and mentors and two optional visits to CMU's Pittsburgh campus. The program will be delivered by SCS faculty, TalentSprint and industry practitioners. Upon successful completion, participants will receive a certificate from SCS.

"We are delighted to partner with Carnegie Mellon's School of Computer Science, a global leader in computer science education, to launch this coding bootcamp," said Santanu Paul, co-founder and CEO of TalentSprint. "This program will build on the success of TalentSprint's coding bootcamps around the world. For participants, it will open the doors to highly differentiated skills and premium careers in tech. For employers, it will provide access to a larger pool of software developer talent trained by CMU faculty." ■

To learn more, visit the program page.

Looking to Bridge Your Tech Talent Gap?

CMU's TechBridge Coding Bootcamp offers your company CMU-trained, job-ready coders, trained by SCS faculty. Our graduates come equipped with the skills to elevate your team.

Skills-First Hiring

As the pace of the tech industry speeds up, companies need in-demand skills. And they need them fast. CMU TechBridge: Coding Bootcamp graduates complete the 16-week industry-relevant curriculum, designed, developed and delivered by SCS faculty in partnership with TalentSprint. This hybrid bootcamp prepares them with the job-ready coding skills you need:

- Programming Fundamental with Python
- Java Script Fundamentals and Advanced Techniques
- HTML and CSS Web Development
- Frontend Development with React
- Server-Side Development with Node.js
- Back-end Development: Databases, Authentication and APIs
- Full Stack Integration
- Cloud Development and DevOps
- Software Engineering Practices

The first cohort started September 23, 2024

Hire from TechBridge today!
Where expertise and innovation converge.

LEARN MORE AT: bootcamps.cs.cmu.edu/coding-bootcamp

OR CONTACT: Ram Konduru, Director, Executive & Professional Education | rk@cmu.edu

Carnegie Mellon University
School of Computer Science
Executive & Professional Education

talent sprint

Nicolas Christin

CHRISTIN NAMED HEAD OF SOFTWARE AND SOCIETAL SYSTEMS DEPARTMENT

Plans To Foster Sustainable, Community-Oriented Growth as Software Reliance Grows

MARYLEE WILLIAMS

Software is ubiquitous in everyday life, but that hasn't always been the case. Just 25 years ago, cellphones came with the game Snake installed, Apple released its clamshell design iBook and LiveJournal hit the internet. That same year, 1999, Carnegie Mellon University's School of Computer Science founded the Institute for Software Research (ISR).

So much has changed since then. Apple sells smartphones that can download any game — including Snake — and blogging is now mostly vlogging. At CMU, ISR is now the Software and Societal Systems Department (S3D), with 57 core faculty members and two Ph.D. programs.

Nicolas Christin is thrilled to be at the center of it all as S3D's new department head.

"This department is at the interface of computer science and real-world impact. As a department head, you have to know everything that's going on," said Christin, who started his new role on September 1. "All the work I hear about is super cool and exciting. I'm like a kid in a candy store."

Christin joked that his security research might not have been housed in S3D when the department was initially founded two decades ago because his work looks at how software impacts society. Now, S3D has academics working on security, concurrent computation, software engineering, game theory, policy and more. For the last 15 years, Christin's work has focused on measuring and analyzing phenomena on the internet, such as the distribution of drugs online or cryptocurrencies, and providing this data to policymakers who can enact change.

"As computer scientists, we know how to capture data and we know how to analyze that data," Christin said. "Ultimately, from this measurement and analysis, you can answer questions like, 'Every year, how many millions or billions of dollars go through the internet to buy and sell narcotics?' And that's interesting to whom? Well, to pretty much everybody, and in particular to policymakers. The U.S. Congress and the European Parliament, for instance, want to understand if they can intervene and what the best intervention would be."

Christin joined CMU in 2005 as a systems scientist. He became an assistant research professor in 2013 and, most recently, a professor in S3D with a joint appointment in the Department of Engineering and Public Policy and a courtesy appointment in the Electrical and Computer

> **I KNOW THAT HIS ENERGY AND RESEARCH LEADERSHIP WILL BE INSTRUMENTAL IN LEADING S3D.**
> — MARTIAL HEBERT

Engineering Department. He's a core faculty member of CyLab, CMU's security and privacy research institute, and the director of S3D's Societal Computing Ph.D. program. Before coming to Pittsburgh, he was a post-doctoral fellow at the University of California, Berkeley, and earned his Ph.D. in computer science from the University of Virginia.

Christin succeeds James Herbsleb, who was appointed in 2019 and shepherded the department through its rebranding from ISR to S3D, managed changes during COVID-19 shutdowns and oversaw an expansion in programs.

SCS Dean Martial Hebert said he was excited for Christin to step into this role.

"I want to thank Nicolas for taking on this responsibility," Hebert said. "I know that his energy and research leadership will be instrumental in leading S3D."

Christian Kästner, an associate professor and director of the Software Engineering Ph.D. program, led the search committee for the position. He said the committee was looking for someone who understood the department's breadth, would lead it fairly and could engage with all the stakeholders.

"We wanted someone who understands the two Ph.D. programs and is committed to the department," Kästner said.

As the director of the Societal Computing Ph.D. program, Christin says he's seen firsthand how the breadth of the research and work in software has changed within S3D. As the department head, he will focus on sustainably managing the department's growth and ensuring that faculty, staff and students are involved in S3D's direction.

"I think this is one of the defining characteristics of our faculties, that people are curious — not just intellectually curious about their own domain, but intellectually curious about stuff that they don't know much about," Christin said. "This helps foster and nurture dialogue between the various department constituencies. One of the big challenges will be to make sure that this intellectual objective translates into more integration between the department's various lines of research."

Over the next few years, Christin wants to spearhead initiatives that bring advanced courses to undergraduates. He said students should have access to classes that reflect the diversity of disciplines interfacing with computer science, such as privacy, ethics and network science.

"I hope that we can enrich the computer science curriculum by providing some offerings that migrate a bit from the Ph.D. level to an undergraduate level," he said. "This is often what happens in science: you get some field of study that is brand new and super advanced, and you need a master's or a Ph.D. to understand it. As time goes by, the discipline matures, and you understand better what you're doing. At that point, it's possible to translate topics in the field and disseminate them to more junior students. That's what I'm hoping for in S3D." ■

Explore Custom and Open Enrollment Training Programs with SCS Executive & Professional Education

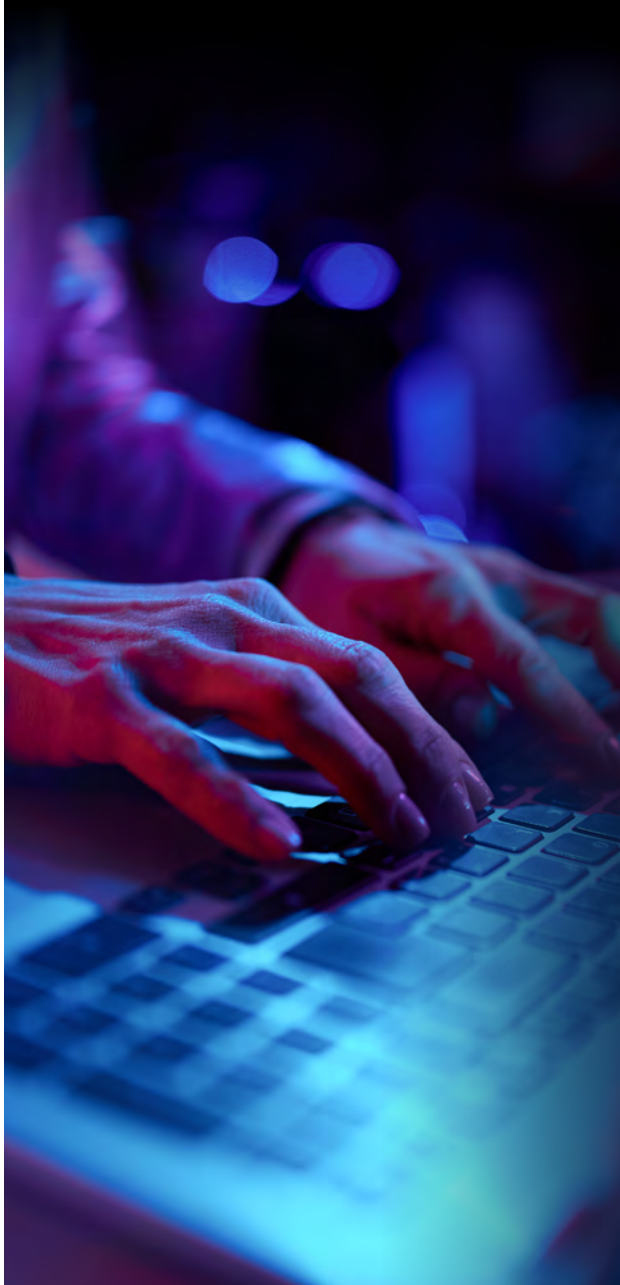

Our faculty have pushed the boundaries of computer science since the inception of the field, in topics ranging from artificial intelligence and robotics, to machine learning and language technologies – and everything in between. Who better to guide you on your enterprise and personal upskilling journey than the experts shaping the future of computer science?

Sample Custom Programs

- Robot Autonomy
- LLMs & Prompt Engineering
- Generative AI
- Safe Autonomy
- Computational Biology
- Memory Safety
- Automated Bug Finding

Open Enrollment Programs

- Artificial Intelligence
- Algorithms & Data Structures
- Computer Vision
- Deep Learning
- DevOps
- Machine Learning
- Natural Language Processing
- Python
- Responsible AI
- Software Architecture

CMU alumni and employees receive a 30% discount on our Open Enrollment programs.

Carnegie Mellon University
School of Computer Science
Executive & Professional Education

Visit exec.cs.cmu.edu for more details.

Zico Kolter

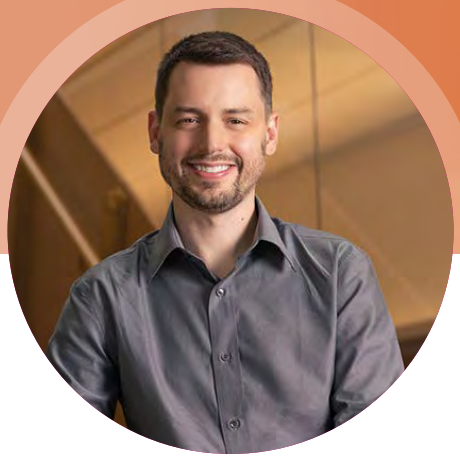

Generative and transformative tools will power the future of computing, and machine learning technologies underpin all the learning, evaluation and improvement of these systems. Research moves quickly from the lab to the real world, where it could transform fields ranging from biology to business.

Zico Kolter is ready to lead that transformation as the new director of Carnegie Mellon University's Machine Learning Department (MLD).

"The pace is unlike anything else I've seen," said Kolter, who has worked at the forefront of machine learning and its real-world applications for almost 20 years. "In typical scientific research, the timeframe between basic research, an idea and its deployment in the world is 10 or 20 years — maybe longer. Machine learning researchers are making advances that we'll see in production a few months after they're published.

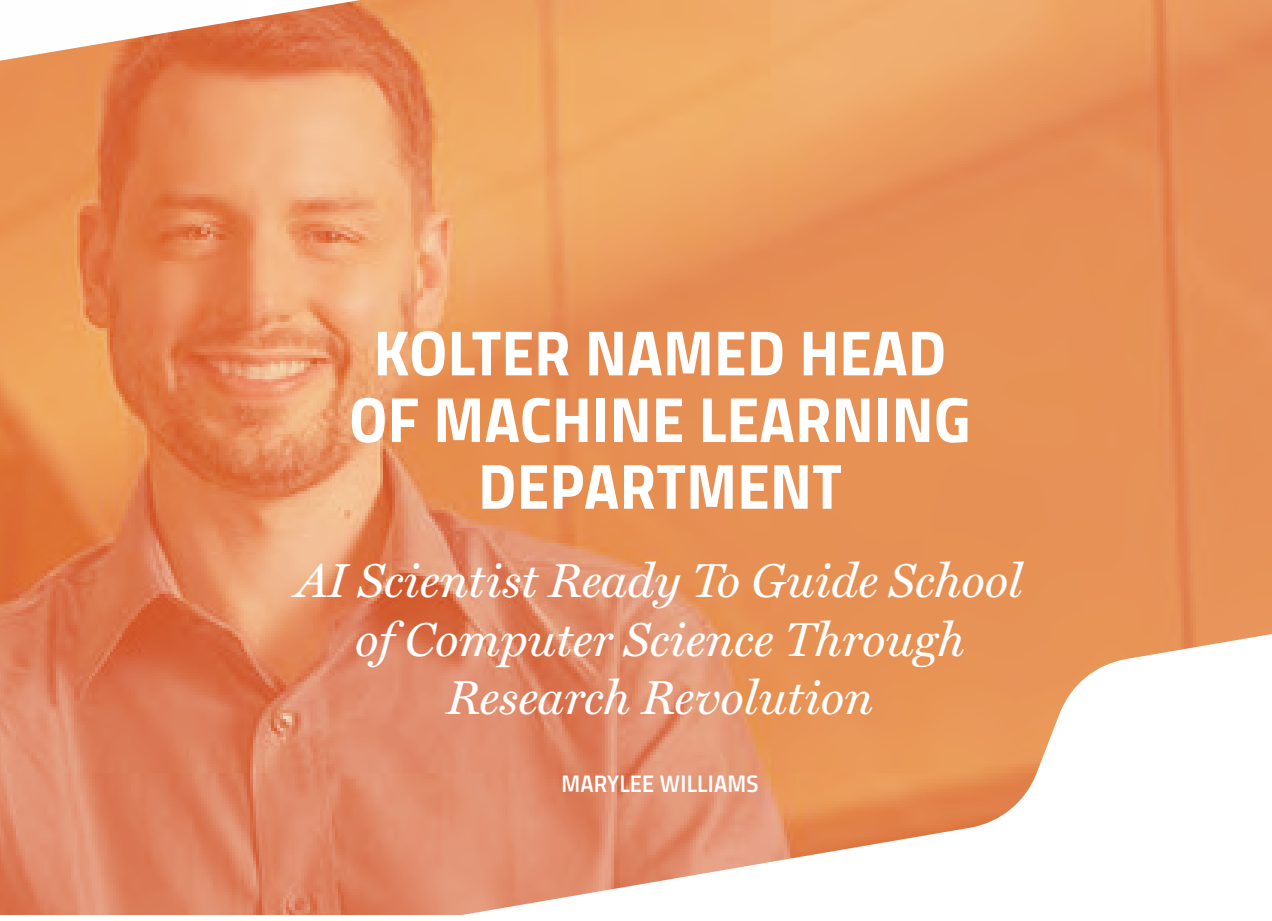

KOLTER NAMED HEAD OF MACHINE LEARNING DEPARTMENT

AI Scientist Ready To Guide School of Computer Science Through Research Revolution

MARYLEE WILLIAMS

"In this moment, the research happening in the Machine Learning Department is influencing and will continue to influence the direction of artificial intelligence," Kolter said.

Kolter joined SCS 11 years ago. He is currently an associate professor in the Computer Science Department with affiliations in the Software and Societal Systems Department, the Robotics Institute and the College of Engineering's Electrical and Computer Engineering Department.

While at CMU, he has also served in various roles in industry, including chief scientist of AI research at Bosch. Kolter earned his Ph.D. in computer science from Stanford University and then completed a postdoctoral fellowship at the Massachusetts Institute of Technology.

SCS Dean Martial Hebert shared his excitement for Kolter stepping into this role.

"I know that his energy and research leadership will be instrumental in leading the Machine Learning Department through this critical period in its history," Hebert said.

Founded in 2006 as the world's first academic department of its kind, MLD evolved from CMU's Center for Automated Learning and Discovery (CALD). Kolter will be its sixth department head, succeeding Roni Rosenfeld, who has held the role since 2018.

Tom Mitchell, SCS Founders University Professor and MLD's founding department head, led the search committee. He said Kolter has the right mix of skills to lead the advancement of both the department and field.

"Zico is a great choice for the head of the Machine Learning Department, especially at this time of major progress and change in AI and machine learning," Mitchell said. "He's a star researcher with a real understanding of how the field is evolving and how the department should help lead this evolution. He has plenty of management experience, lots of ideas for the future and he's fun to work with. I look forward to supporting him as he leads the department to the next level."

> **I WANT TO ENSURE THAT OUR BIGGEST RESOURCE, OUR PEOPLE, HAVE THE INFLUENCE ON THE FIELD THAT WE CAN AND SHOULD HAVE.**
> — NICO KOLTER

Kolter's research focuses on machine learning, optimization and control, with much of the work centered on making deep learning algorithms safer, more robust and more modular. Students working with Kolter undertake an immense breadth of AI research. Kolter wants to put researchers at the vanguard of developing the field's core methodologies and assessing the potential risks or concerns that arise from AI.

"This is an area where I would argue that the Machine Learning Department excels and will continue to excel," Kolter said. "We have people working at the forefront of ethics, safety, transparency and explainability. These areas are going to be equally crucial as research and as policy to enable the safe and responsible uptake of AI."

Last year, Kolter's work on attacking large language models (LLM) showed there were ways to circumvent their safeguards. The research team created a system that automatically generated suffixes that, when attached to a query, caused the LLMs to answer dangerous questions, such as "How to hotwire a car?" The researchers hoped their work would clarify the risk these automated attacks pose to LLMs like ChatGPT.

Today's LLMs have changed dramatically from past models. Kolter said his lab has worked on understanding how data impacts the way these models work and has helped define memorization in these models, addressing concerns about permissible data usage.

"We don't even know the data we use to train models in many cases, yet we're still using them. We need to understand how data impacts these models — how different design choices made during training can lead to vastly different downstream behavior," Kolter said. "It is vital to understand how to use these tools safely."

Across all the work his lab has done, Kolter stresses how impressive it is that students conduct a variety of AI research. As the department head, he hopes to empower and facilitate what he sees as the most valuable resource at MLD: its people. This includes students, faculty and staff, Kolter said. His primary goal is to ensure that the people in and outside MLD have what they need to succeed.

"I want to ensure that our biggest resource, our people, have the influence on the field that we can and should have. This includes getting them the access to the models, data and compute that they need," Kolter said.

Kolter became the MLD department head on June 15. Rosenfeld, an MLD professor with affiliations in the Language Technologies Institute, the Ray and Stephanie Lane Computational Biology Department, and CMU's Heinz College of Information Systems and Public Policy, announced in January that he was stepping down as department head to focus on the work of the Delphi Group, of which he is a co-founder and co-lead. ■

FACULTY AUTHORS

ROBOTS: EXPLORE THE WORLD OF ROBOTICS AND AI

by Henny Admoni, Assistant Professor in RI

From pizza-making robots and slithering snake-bots to robots that help keep humans safe in disaster zones and robotic rovers already exploring space, robots are no longer science fiction.

In her new children's book, "Robots: Explore the World of Robotics and AI," Admoni explains what robots are, and shows how they are made and some of the amazing ways we can use them. Along with beautiful illustrations by Amy Grimes, Admoni answers some key questions for kids about robotics and AI, emphasizing the positive impact robots have on us all.

Available from Macmillan Publishers

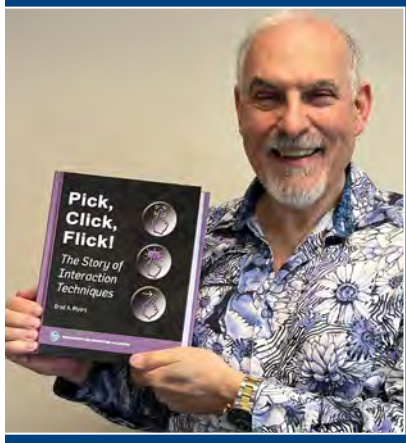

PICK, CLICK, FLICK! THE STORY OF INTERACTIVE TECHNIQUES

by Brad A. Myers, Director of the HCII

A new book from Brad Myers explores the history, current and future design of interaction techniques. The book, "Pick, Click, Flick! The Story of Interaction Techniques," digs deep into the details of interaction techniques and how they work, often with information from the techniques' inventors.

This is Myers' fourth book and available in hardcover, paperback and digital forms from the publisher ACM Books.

C. Gordon Bell

BUILT THE FOUNDATION FOR MODERN COMPUTING

MATTHEW WEIN

C. Gordon Bell, a visionary designer of computer systems and former Carnegie Mellon University faculty member whose work helped shrink computers from room-filling mainframes to more compact, affordable and practical machines, died May 17, 2024, at his home in Coronado, California. He was 89.

"Most people consider Gordon the father of computer architecture," said Daniel Siewiorek, the Buhl University Professor of Electrical and Computer Engineering and Computer Science at CMU and Bell's close friend and colleague.

"Gordon laid the foundation for modern computing and transformed how people interact with technology," said CMU President Farnam Jahanian. "His infectious curiosity will continue to inspire — both in the world at large and here at Carnegie Mellon University, through the alumni who benefited from his wisdom during his tenure on our faculty and through his generous support of our computer science students and faculty today. Our thoughts are with Gordon's family and loved ones as we remember his remarkable legacy and celebrate his many contributions to society."

Bell had already established himself as one of the preeminent computer architects in the world by designing minicomputers for the Digital Equipment Corporation when, in 1966, he went looking for a change of scenery and a new project. He found it at CMU.

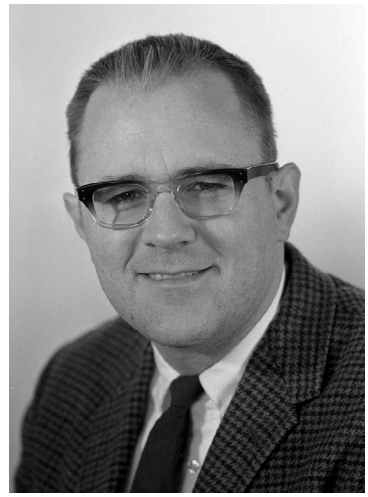

C. Gordon Bell, a visionary designer of computer systems and former Carnegie Mellon University faculty member, died May 17. He was 89.

Photo courtesy of the Marketing and Communications Photograph Collection, Carnegie Mellon University Archives.

"Gordon had great admiration for the scientific mind of Allen Newell, and it was Newell who convinced him to go to Carnegie Mellon," said Edward Feigenbaum (ENG 1956, TPR 1960), Professor Emeritus of Computer Science at Stanford University and a friend and colleague of Bell. "It was a hotbed of new ideas, and Newell agreed to collaborate with him."

Professorship was a jump for Bell, who had only a handful of scholarly publications and no Ph.D. when he arrived on campus. He had started doctoral work at the Massachusetts Institute of Technology but abandoned it soon after.

"I didn't really want a Ph.D.," Bell said in a 1991 interview. "I wanted to build things."

"The computers that Gordon built were his publications," Siewiorek said. "CMU was very broad in recognizing contributions and things that can have an impact."

"Gordon laid the foundation for modern computing and transformed how people interact with technology."

— CMU PRESIDENT FARNAM JAHANIAN

At CMU, Bell and Newell teamed up to write "Computer Structures: Readings and Examples," a nearly 700-page tome that established a taxonomy for computer architecture and became an essential text in the field.

"It basically legitimized the study of computer structures," Siewiorek said. "It really taxonomized and organized the field and had a lot more value than just citations in the publication."

Chester Gordon Bell was born in Kirksville, Missouri, on Aug. 19, 1934. His mother, Lola, was a schoolteacher and his father, Chester, was an electrician. Bell showed an early predilection toward technology, spending hours in his father's shop taking apart and reassembling various home appliances.

After earning both a bachelor's and master's degree in electrical engineering from MIT, Bell went on to the University of New South Wales as a Fulbright scholar. In 1960, he began designing minicomputers for the Massachusetts-based Digital Equipment Corporation (DEC). He was the primary architect of the PDP-8, a 12-bit machine that became the first commercially successful minicomputer.

"Gordon was a master of systems, but he also had an aesthetic about systems," Feigenbaum said. "They had to be not just functional, but beautiful."

After six years at CMU, Bell returned to DEC where he remained until 1983, when a heart attack prompted him to step away from the company.

Bell went on to found a pair of startups and invest in more than 100 others. In 1987, he joined the National Science Foundation to help build the National Research and Education Network — a supercomputer networking project that produced an early iteration of the internet. That same year, he sponsored the creation of the ACM Gordon Bell Prize, which is given annually to recognize achievement in high-performance computing.

Bell served as an adviser to Microsoft in the early '90s before joining the company's Silicon Valley Research Lab in 1995, where he worked until 2012. He returned to Microsoft in an emeritus role in 2015.

An avid collector of historical pieces of computer hardware, Bell and his first wife, Gwen, co-founded the Digital Computer Museum, now known as the Computer History Museum.

"He started collecting items in his house, and when it started overflowing, he got DEC to offer up some space," Siewiorek said.

Once the museum outgrew the DEC lobby, it moved to a dedicated space in the Museum Wharf in Boston. The museum is now located in Mountain View, California.

In 2010, CMU recognized Bell with an honorary Doctor of Science and Technology degree. Additionally, the fifth-floor conference room in the Gates-Hillman Centers is named in his honor.

"He was just amazing," said Raj Reddy, the Moza Bint Nasser University Professor of Computer Science and Robotics at CMU. "Gordon was a master organizer, but he was also incredibly friendly, collegial and supportive."

"He was always encouraging. His attitude and approach really rubbed off," Siewiorek said. "He inspired and enabled a lot of people."

Bell is survived by his second wife, Sheridan Sinclaire-Bell; a son, Brigham, and daughter, Laura, both from his first marriage; a stepdaughter, Logan Forbes; a sister, Sharon Smith; and four grandchildren. ■

Names in the **News**

Mary Shaw has received the **IEEE Computer Society's Technical Community on Software Engineering Lifetime Achievement Award** for her pioneering and lifetime contributions to software engineering.

SCS faculty members Nathan Beckmann, Aaditya Ramdas, Justine Sherry and Virginia Smith (from upper left to lower right) have been named **2024 Sloan Research Fellows**.

Lenore Blum, a foundational researcher in computer science at CMU and a tireless advocate for women in math and science, has been elected to the **American Academy of Arts & Sciences**.

Matthew Mason, professor emeritus and a former Robotics Institute (RI) director, has been elected a fellow of the **American Association for the Advancement of Science**.

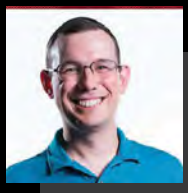

Faculty members Joshua Sunshine (upper left) and Steven Wu (lower left) have received **NSF Faculty Early Career Development Program (CAREER)** awards totaling more than $1.5 million.

CSD Ph.D. Student Kaiyang Zhao was selected for a **North American Qualcomm Innovation Fellowship**.

CSD's Rashmi Vinayak and Juncheng Yang won the **Community Award** at the 2024 USENIX Symposium on Networked Systems Design and Implementation.

SCS student Claire Jin was one of CMU's three **2024 Goldwater Scholars**, one of the most prestigious STEM scholarships for undergraduates.

Andrea Bajcsy, Motahhare Eslami, Ken Holstein, Aditi Raghunathan, Andrej Risteski and Hong Shen (from upper left to lower right) earned **2024 Google Research Scholar Awards**.

Gabriele Farina, who earned a Ph.D. from the Computer Science Department last year, received the **ACM Special Interest Group on Economics and Computation Dissertation Award**.

Three members of the HCII earned **2024 SIGCHI Awards**, including (from left) Jodi Forlizzi for Lifetime Research, Amy Ogan for Societal Impact and Karan Ahuja for Outstanding Dissertation.

CSD Ph.D. student Zoë Marschner earned a **2024 Hertz Fellowship**, which provides five years of funding for students in applied science, engineering and mathematics.

Guy Blelloch was elevated to the rank of **University Professor**, the highest distinction a faculty member can receive at CMU.

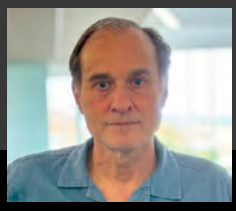

Founders University Professor Takeo Kanade received the **Banco Bilbao Vizcaya Argentaria Foundation's Frontiers of Knowledge Award** in Information and Communication Technologies.

Carnegie Mellon University
School of Computer Science

OFFICE OF THE DEAN
5000 FORBES AVENUE
PITTSBURGH PA 15213-3890

SCSatCMU

cs.cmu.edu
magazine@cs.cmu.edu