the LINK

THE MAGAZINE OF CARNEGIE MELLON UNIVERSITY'S SCHOOL OF COMPUTER SCIENCE

60 YEARS IN THE MAKING

CMU AI is Here

SUMMER 2017
ISSUE 11.1

theLINK

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.

Iain Mathews

Bhat, Matthews Win Academy Awards for Technical Achievement

School of Computer Science alumnus Kiran Bhat and former Robotics Institute faculty member Iain Matthews received Oscars on February 11, from the Academy of Motion Picture Arts and Science, for their work in capturing facial performances.

Bhat earned his doctorate in robotics in 2004, and helped design and develop the Industrial Light and Magic facial performance-capture solving system, which transfers facial performances from actors to digital characters in large-scale productions. The system was used in "Rogue One: A Star Wars Story" to resurrect the role of Grand Moff Tarkin, played by the late actor Peter Cushing, as well as to capture Mark Ruffalo's expressions for his character, the Hulk, in "The Avengers."

Matthews, a post-doctoral researcher and former faculty member in the Robotics Institute working on face modeling and vision-based tracking, was recognized along with his team for the design, engineering and development of the facial-performance capture and solving system at Weta Digital, known as FACETS. Matthews spent two years helping to develop the facial motion capture system for "Avatar" and "Tintin."

With Bhat's and Matthews' wins, Carnegie Mellon alumni and faculty have received nine Academy Awards to date.

The Link
Summer 2017 | Issue 11.1

Publisher
Andrew W. Moore

Editor
Kevin O'Connell

Contributing Writers
Susie Cribbs (DC 2000, 2006), Meghan Holohan, Nick Keppler, Kevin O'Connell, Aisha Rashid (DC 2019), Mark Roth, Linda K. Schmitmeyer, Byron Spice

Photography
Chandler Crowell, Rebecca Kiger, Howard Korn

Illustration
Curt Merlo

Design
Vicki Crowley (A 1996)

Office of the Dean
Gates Center 5113
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213
www.cs.cmu.edu

thelink@cs.cmu.edu

contents

Where We've Been and Where We're Going

Taking Stock

I'll celebrate my third anniversary as dean of the School of Computer Science in August, which I can barely believe. Wasn't it yesterday when they introduced me as the new dean? At the same time, though, so much has been going on that I'm not sure how it's only been three years. If you remember, in the fall of 2014 you all helped me out hugely in my onboarding "sponge" period, and then we wrote down the things we needed to work on. Since I'm more than halfway through my term, now seems like an appropriate time to take stock.

By the way, I don't say "we" accidentally. It's not a typo. Everything we do in SCS and all of our successes result from the incredible faculty, staff and students I have the privilege to work with every day.

Now let's get down to business.

I firmly believe that those of us in positions in advanced technology organizations have an obligation to create technology that improves people's lives, and our work in SCS over the past few years supports that vision. In this issue of The Link alone, you'll learn about how we're amassing our artificial intelligence resources on campus to create an initiative that will ensure we use AI for good. You'll read it in a Q&A, but you can also read about how AI helps improve our infrastructure and allows us to solve problems even when we don't have all the information we think we need.

During that sponge period I talked about, faculty members repeatedly told me that SCS lacks the big, collaborative, world-changing research projects that in the past resulted in breakthroughs like the Andrew wired and wireless networks. While times have changed and the sort of funding that enabled those big projects can be more difficult to find now, we've started an SCS Moonshots initiative that highlights projects we think will have the most wide-ranging implications for humanity. These projects receive the financial support and resources necessary for the risk-taking, all-hands approach we're known for at CMU. You read about one of them, FarmView, in our winter issue. In this issue, we introduce you to Revolutionizing the Creation of Assistive Devices — using 3-D printing to create more affordable, accessible artificial limbs for those who need them.

One challenge all CS programs face is diversity, and it's been a huge, multidecade initiative in SCS to create an environment where no one would ever think of SCS as a one-gender club. The class of 2020, who just finished its first year with us, was nearly half women, and the class of 2021 boasts 49 percent women as well. We hardly consider our efforts to achieve gender parity complete, though. To augment the work of organizations like SCS4All and Women@SCS, we now offer voluntary training for faculty and staff on recognizing and dealing with the unconscious bias we all carry with us. This BiasBusters program is something I'm incredibly excited about, and is featured in this issue.

But all is not well here. The world of computer science education also lacks economic diversity and the participation of underrepresented minorities. I don't think this is acceptable. We are SCS and we need to lead the rest of the world here, just as we are doing in so many other dimensions. One step we've taken in the right direction is establishing the Hopper-Dean Computer Science Participation Fund. Through the generosity of Heidi Hopper and Jeff Dean, the fund will support our efforts to reach out to students who are underrepresented in the field — whether socioeconomically or geographically — before they even get to college. Through community partnerships, and summer and school programs (including the Carnegie Mellon summer SAMS program, which is introducing a computer science track this year, and our partnership with the Microsoft TEALS program), we hope to identify and support students who have incredible CS potential, but may lack the precollege resources to help them reach that potential.

Using AI for good. Changing the world through collaborative, risk-taking research. Striving for gender parity, uncovering and dealing with bias, expanding CS to people underrepresented in the field. Maybe it doesn't seem like much, but it keeps me running day in and day out. Sometimes it even keeps me up at night. But I know that since I'm surrounded by the smartest people on the planet, changing the world for the better won't be nearly as hard as it seems.

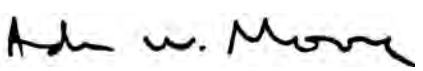

Andrew W. Moore
Dean, School of Computer Science

AI FOR GOOD, AI FOR EDUCATION

You might have noticed that SCS is making a full-court press on artificial intelligence these days. With the launch of CMU AI — a new initiative that marshals our AI work across departments and disciplines — we've created one of the largest and most experienced AI research groups in the world.

While AI research is important, we also know that the future of the field rests largely in the hands of students who haven't even reached college yet. As part of CMU AI, we're vigorously expanding our outreach efforts in programs that introduce these elementary and high school students to AI topics and encourage their computer science curiosity.

"We're teaching and engaging with those who will improve lives through technology, and who have taken responsibility for what happens in the rest of the century," SCS Dean Andrew Moore said. "Exposing these hugely talented human beings to the best AI resources and researchers is imperative for creating the technologies that will advance mankind."

At right is a reprint of a letter Moore, Associate Teaching Professor David Kosbie and Assistant Dean for Outreach Mark Stehlik wrote for the Harvard Business Review on how we can train the next generation of computer scientists.

HOW TO PREPARE THE NEXT GENERATION FOR JOBS IN THE AI ECONOMY

David Kosbie, Andrew W. Moore and Mark Stehlik
Originally appeared in the June 2017 Harvard Business Review.

MOST OF US regard self-driving cars, voice assistants and other artificially intelligent technologies as revolutionary. For the next generation, however, these wonders will have always existed. AI for them will be more than a tool; in many cases, AI will be their co-worker and a ubiquitous part of their lives.

If the next generation is to use AI and big data effectively — if they're to understand their inherent limitations, and build even better platforms and intelligent systems — we need to prepare them now. That will mean some adjustments in elementary education and some major, long-overdue upgrades in computer science instruction at the secondary level.

For example, consider how kids are currently interacting with AI and automated technologies: Right now, it might seem magical to tell Siri, "Show me photos of celebrities in orange dresses," and see a photo of Taylor Swift pop up on a smartphone less than a second later. But it's clearly not magic. People design AI systems by carefully decomposing a problem into lots of small problems, and enabling the solutions to the small problems to communicate with each other. In this example, the AI program divides the audio into chunks, sends them into the cloud, analyzes them to determine their probable meaning and translates the result into a set of search queries. Then millions of possible answers to those queries are sorted and ranked. Thanks to the scalability of the cloud, this takes just a few dozen milliseconds.

This isn't rocket science. But it requires a lot of components — waveform analysis to interpret the audio, machine learning to teach a machine how to recognize a dress, encryption to protect the information, etc. While many are standard components that are used and reused in any number of applications, it's not something a solitary genius cooks up in a garage. People who create this type of technology must be able to build teams, work in teams and integrate solutions created by other teams. These are the skills that we need to be teaching the next generation.

Also, with AI taking over routine information and manual tasks in the workplace, we need additional emphasis on qualities that differentiate human workers from AI — creativity, adaptability and interpersonal skills.

At the elementary level, that means that we need to emphasize exercises that encourage problem solving and teach children how to work cooperatively in teams. Happily, there is a lot of interest in inquiry-based or project-based learning at the K-8 level, though it's hard to know how many districts are pursuing this approach.

Ethics also deserves more attention at every educational level. AI technologies face ethical dilemmas all the time — for example, how to exclude racial, ethnic and gender prejudices from automated decisions; how a self-driving car balances the lives of its occupants with those of pedestrians, etc. — and we need people and programmers who can make well-thought-out contributions to those decision-making processes.

We're not obsessed about teaching coding at the elementary levels. It's fine to do so, especially if the kids enjoy it, and languages such as Snap! and Scratch are useful. But coding is something kids can pick up later on in their education. However, the notion that you don't need to worry at all about learning to program is misguided. With the world becoming increasingly digital, computer science is as vital in the arts and sciences as writing and math are. Whether a person chooses to become a computer scientist or not, coding is something that will help a person do more in whatever field they choose. That's why we believe a basic computer programming course should be required at the ninth grade level.

Only about 40 percent of U.S. schools now teach programming, and the quality and rigor of these courses vary widely.

Only about 40 percent of U.S. schools now teach programming, and the quality and rigor of these courses vary widely. The number of students taking Advanced Placement exams in computer science is growing dramatically, but the 58,000 students taking the AP Computer Science A (APCS-A) test last year still pales in comparison to the 308,000 who took the AP Calculus AB test. A third of our states don't even count computer science course credits toward graduation requirements.

The U.S. is woefully behind many of our peer nations. Israel notably has integrated computer science into its precollege curriculum. The UK has made good progress lately with its Computing at School program, and Germany and Russia have leapt ahead as well. President Obama's Computer Science for All initiative, announced in his 2016 State of the Union, was a belated step in the right direction, but may flounder amid budget cuts proposed by the Trump administration.

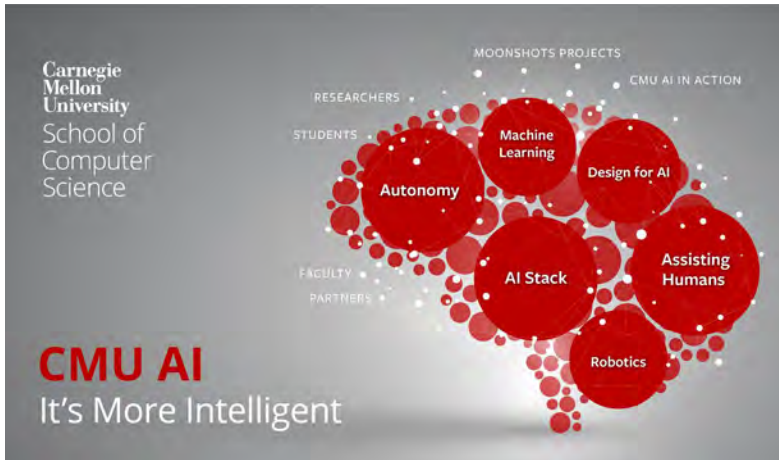

Expanding computer science at the high school level not only benefits the students, but could help the field of computer science by encouraging more students — and a more diverse group of students — to consider computer science as a career. Though we were thrilled last fall when almost half of our incoming first-year class at Carnegie Mellon was female, the field of computer science is still struggling to increase the number of women and minorities. Engineering intelligence into systems, and finding insights in a ubiquitous sea of data, is a task that cries out for a diverse workforce.

To be successful, however, it is critical that we update the way programming is taught. We're too often teaching programming as if it were still the '90s, when the details of coding (think Visual Basic) were considered the heart of computer science. If you can slog through programming language details, you might learn something, but it's still a slog — and it shouldn't be. Coding is a creative activity, so developing a programming course that is fun and exciting is eminently doable. In New York City, for instance, the Girl Scouts have a program that teaches girls to use JavaScript to create and enhance videos — an activity that kids already want to do because it's fun and relevant to their lives. Why can't our schools follow suit?

Beyond ninth grade, we believe schools should provide electives such as robotics, computational math and computational art to nurture students who have the interest and the talent to become computer scientists, or who will need computers to enhance their work in other fields. Few U.S. high schools now go beyond the core training necessary to prepare for the APCS-A exam, though we have a few stunning success stories — Stuyvesant High School in New York City, Thomas Jefferson High School for Science and Technology in Alexandria, Va., and TAG (The School for the Talented and Gifted) in Dallas, among others. These schools all boast committed faculty members who have a background or training in computer science.

We also urge high school math departments to place less emphasis on continuous math, including advanced calculus, and more on the math that is directly relevant to computer science, such as statistics, probability, graph theory and logic. Those will be the most useful skills for tomorrow's data-driven workforce.

A major hurdle is that our schools face a severe shortage of teachers who are trained in computer science. This is where U.S. tech companies could help immensely. Microsoft, for instance, sponsors the TEALS program, which pairs computer professionals with high school teachers for a few hours a week. But we need thousands of educators teaching millions of students. Even greater commitments will be necessary going forward. On the academic side, The University of Texas at Austin's UTeach program is a model for preparing STEM teachers and has expanded to 44 universities in 21 states and the District of Columbia.

Much more is needed. As with science and math, we need governmental standards driving K–12 computer science education, along with textbooks, courses and ultimately a highly trained national cadre of computer science teachers that are tied to those standards. The Computer Science Teachers Association has been a leader in this area, promulgating a standards framework and an interim set of standards.

Investing in how the next generation understands and interacts with big data and AI is an investment that will pay off in the long run for all of us. ■

WHAT AI MEANS FOR LIFE, THE UNIVERSE AND EVERYTHING

an interview with Andrew Moore

Susie Cribbs (DC 2000, 2006)

Mention artificial intelligence in a crowd and no doubt some people will get twitchy. "Terminator"esque scenarios might come to mind, or visions of a world where humanity creates robots so smart that they rip our jobs out from under us, rendering us penniless and useless.

Andrew Moore is not one of those people.

As dean of the School of Computer Science — not to mention a former professor here, and a VP at Google Research Pittsburgh — Moore has witnessed AI's evolution at CMU and believes there's no better place in the world to harness its power for the good of both the planet and its people. He's so committed to this belief that he's one of the founders of a new initiative, CMU AI, that unites all of the AI resources in SCS and CMU to form the largest, most experienced cohort of AI experts in the world. We sat down with Dean Moore this spring to get a better understanding of why he's so dedicated to AI, and what he envisions for the future of the field — both at Carnegie Mellon, and around the world.

Why is artificial intelligence such a big deal?

I believe that, in the future, the world will realize that it's a safer place because of AI, and that people will enjoy interacting with intelligent systems. I want it to be the case that if someone is in trouble — they're getting sick, stressed out by their work or are about to make a bad decision — humanistic artificial intelligence can step in and help them like a friend. Not like an evil Big Brother with a red eye.

But why is CMU the place for this? Plenty of folks are working on AI research.

We're the right place to build intelligence that co-exists with humanity and makes our lives better and safer because we've spent decades building a culture where people care about solving the world's problems through technology.

What drew me personally to Carnegie Mellon is Allen Newell and Herb Simon's vision of a general problem solver for the human race. I think as the years have gone by, this vision has become a filter: people who are attracted to it come here. And now we have this large group of people determined to build things that actually help humans.

Would you say those are the factors that make us unique?

Not entirely. There are some other things going on here that you won't find anywhere else. First, we're building systems that matter. The second factor is our size and the local café culture. I think this is probably the only zip code or pair of zip codes in the world where if you're sitting in a bar, there's more than a two percent chance you'll hear someone talking about the Nash equilibrium or 3-D modeling of a thorax. We have a ridiculously high density of people who are doing this kind of work, making collaboration natural and easy.

Finally, our students are amazing. They're incredibly smart. It's much easier to be successful when you're surrounded by some of the smartest people in the world.

AWM on AI in DC
The AI race is on, and SCS Dean Andrew Moore wants to ensure the U.S. leads the charge. Moore took the AI show on the road last fall, when he addressed the U.S. Senate Committee on Commerce, Science and Transportation's Subcommittee on Space, Science and Competitiveness. During his testimony, Moore shared his view of what AI is and what it isn't, then focused on what the nation can do to ensure its competitiveness moving forward.

"The ability to perceive the world and the ability to search over really vast numbers of possibilities have both gone up wildly. This means there's now a land grab going on, where students, faculty and entrepreneurs around the world are finding new tasks for AI algorithms to tackle."

What are your favorite CMU AI projects?

That's a tough one, because of how the field has evolved. We've started to understand people better, first through modeling how they solve problems, then by modeling their language and speech. Now we're modeling what they look like and how their facial expressions change. Over time, we've increased our understanding of what it is to be human so that we could put that knowledge into real systems.

So if you ask me what my favorite thing is from the 1980s, it would be Raj Reddy's pioneering work on speech. The 1990s? Tom Mitchell and Manuela Veloso's research on artificial intelligences working together as teams. Going into the 2000s, it was the birth of machine learning as a huge new research area.

In retrospect, this trajectory makes sense. We first became interested in problem-solving and pioneered that. But we hit a natural limit, because no matter how smart your problem-solving algorithms are, they can't help you if they can't see what's going on in the world. From the mid 1990s until about 2010, we worked on creating and improving upon technologies that observed and reacted to the world around them. The thing that brought me back to CMU and that keeps me excited about my job is that since 2010, we've been able to return to Allen and Herb's original dream: AIs that know enough about the real world to solve real-world problems. Because now we have the technology we need to understand the world.

What are some practical ways AI will change how we live?

Here's a good example. Eight percent of search engine queries are people asking for help because they're worried about a medical issue. Imagine what the world would be like if the answers to those queries were vastly better than they are now. Suppose those rare cases when there's something serious going on were caught early without causing the use of unnecessary, expensive or dangerous therapy. AI will help us make better decisions about how to proceed in these situations — which can help people who don't have easy access to healthcare.

Do you like the term "artificial intelligence?" Will it soon just be "real intelligence?"

I like the term artificial. We are absolutely simulating human intelligence, but it's a really coarse simulation. It's actually more valid to compare airplanes to birds than it is to compare current AIs to humans.

Do you think AI and CS will soon be synonymous?

Yes. AI is often what we talk about when we look at the end application of an intelligent system. But that application lives on the computer science stack, which lives on the math stack, which lives on the physics stack. You cannot have a system that, for example, helps direct aid to refugees without having a strong machine learning algorithm behind it. That machine learning algorithm relies on a strong data processing algorithm to manage the data. That data processing algorithm requires strong operating systems work. That strong operating system needs the ingenuity of electrical engineers and physicists. So we're all part of this big AI stack.

Why should students come to CMU to study AI?

Because they want to work with practitioners — people who have built dozens of robots or are fielding large systems keeping us safe from radiologic threats. They want to be immersed in this world of two or three hundred AI experts who lead by example. Our students learn by doing, and they learn from the best in the business.

What do you say to folks terrified of AI?

I passionately believe AI will save billions of lives this century, but it's not without its dangers. That being said, AI is in better hands at CMU than the majority of places. We don't simply study AI, we think about the consequences of what we do. We're human-focused. And we're instilling those values into our students, who will build the AIs for the century.

The world is a scary place right now. All kinds of things make us fearful about the future. AI is one of the extra arrows we humans have in our quiver to potentially improve things dramatically. It's one of the things we've got going for us right now. That's why I believe it's so important. Those of us who have been lucky enough to be trained in advanced technology and computer science have a duty to use our talent and skills to make the world a safer place.

Poker AI Makes History

It was a history-making spring for Computer Science Professor Tuomas Sandholm and Ph.D. student Noam Brown, whose poker-playing AI agent beat four of the world's top poker professionals in the 20-day "Brains vs. Artificial Intelligence: Upping the Ante" competition.

Held at Pittsburgh's Rivers Casino, the competition pitted the AI, Libratus, against pros Dong Kim, Jimmy Chou, Daniel McAulay and Jason Les in 120,000 hands of Heads-up, No-Limit Texas Hold'em — an exceedingly complex game with more information sets than there are atoms in the universe. To further complicate matters, the AI has to make decisions without knowing all the cards in play, while trying to sniff out bluffing.

How did Libratus secure such a sound victory? Algorithms and the Pittsburgh Supercomputing Center's Bridges computer.

"After play ended each day, a meta-algorithm analyzed what holes the pros had identified and exploited in Libratus' strategy," Sandholm said. "It then prioritized the holes and algorithmically patched the top three using the supercomputer each night. This is very different than how learning has been used in the past in poker. Typically, researchers develop algorithms that try to exploit the opponent's weaknesses. In contrast, here the daily improvement is about algorithmically fixing holes in our own strategy."

Sandholm also said that Libratus' end-game strategy, computed live with Bridges for each hand, was a major advance.

"The end-game solver has a perfect analysis of the cards," he said.

But wait. There's more.

After Libratus' victory in Pittsburgh, Sandholm traveled to China with a version of Libratus named Lengpudashi or "cold poker master." In a five-day exhibition in Hainan, the AI beat six Chinese players by a total of $792,327 in virtual chips.

This new milestone in AI has implications for any realm where information is incomplete and opponents sow misinformation. Business negotiation, military strategy, cybersecurity and medical treatment planning could all benefit from automated decision-making that uses similar technology.

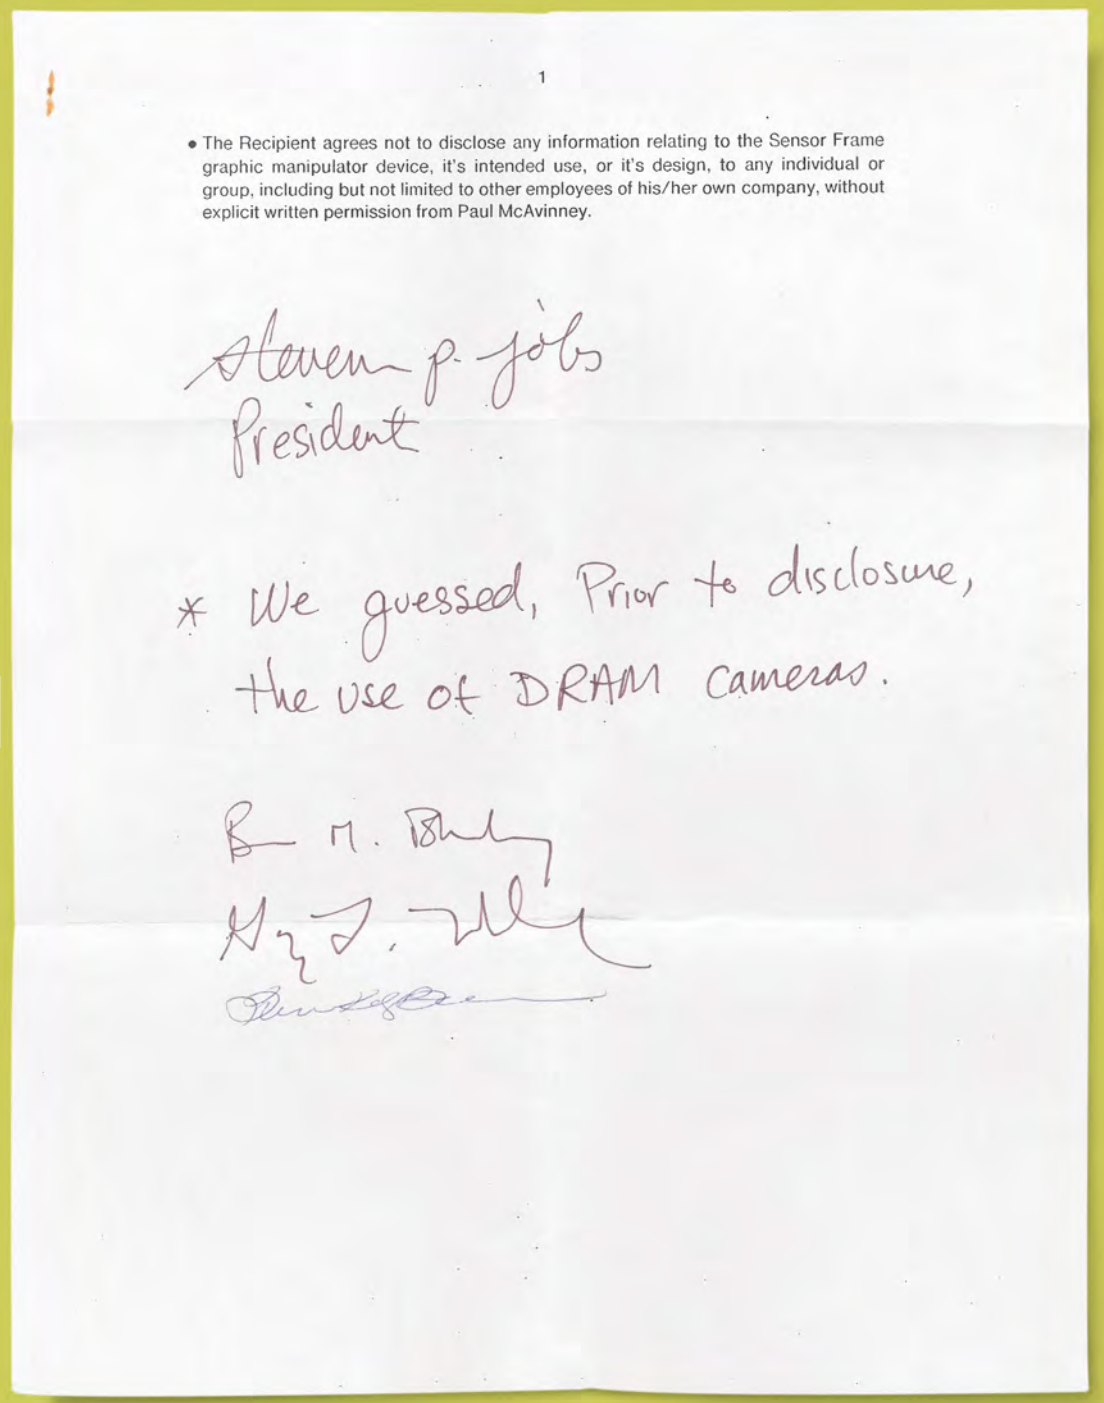

1

• The Recipient agrees not to disclose any information relating to the Sensor Frame graphic manipulator device, it's intended use, or it's design, to any individual or group, including but not limited to other employees of his/her own company, without explicit written permission from Paul McAvinney.

Steven P. Jobs
President

* We guessed, Prior to disclosure, the use of DRAM cameras.

B. M. Buxton

Bill Gates

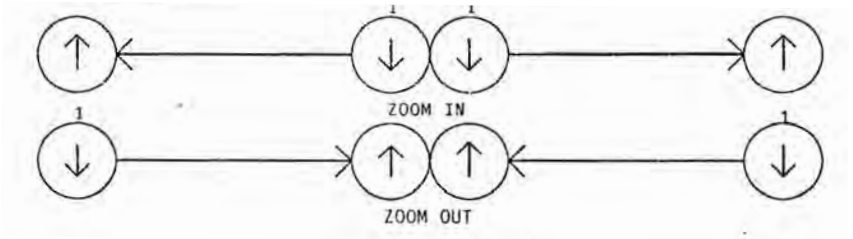

THE UNTOLD HISTORY OF MULTITOUCH

STEVE JOBS RELUCTANTLY SIGNED THE NON-DISCLOSURE AGREEMENT sitting on the desk in Carnegie Mellon's Music Lab. The people in the room say he didn't want to, but he and his team needed to know what was happening inside the lab. So he signed it. It was 1985, and the world as we know it now — teeming with multitouch tablets and smartphones — remained a dream. But inside that lab, three Carnegie Mellon researchers were developing technologies that not only compelled Jobs to get a glimpse inside, but also changed computing as we know it.

Kevin O'Connell

ROGER DANNENBERG (SCS 1983), first met Paul McAvinney in 1983 when the two discussed creating a computer music group on campus. McAvinney, an independent inventor, did not work at CMU at the time, but their conversation formed the impetus for creating the Music Lab within SCS.

After the conversation with Dannenberg, McAvinney realized developing computer music would require an innovative approach and technologies not available at the time. "It occurred to me that in order to control music in real time, you really need to be able to change things at the tens-of-milliseconds level," said McAvinney. "I wondered if I could use a touch screen to do that." He needed a screen that could respond to multiple points of contact and allow for the manipulation of objects. The essence of multitouch.

In today's computing landscape, we use multiple fingers to interact with computer screens all the time. But single-touch screens were all that was available in 1983. McAvinney contacted a screen manufacturer with the notion that they might help him develop a multitouch interface, but they showed little interest. So he set off on his own to create one.

McAvinney called what he created the Sensor Frame — an open frame of infrared lights that fit over a computer monitor. The prototype used infrared-sensing dynamic random access memory (DRAM) chips that detected the obstruction of light when fingers entered the frame. This optical approach differs from capacitive touch technology used in smartphones and tablets today.

He showed his work to Dannenberg, and piqued his interest.

"[The team] invited Raj Reddy, then head of the Robotics Institute, to see what I had done," said McAvinney. His work was well received. Dannenberg then invited McAvinney to move his operation on campus. Newly minted graduate student Dean Rubine (CS 1991) joined the team in the fall of 1983, primarily developing Sensor Frame's software.

Thus the Music Lab — and its multitouch research — was born.

In this lab, McAvinney began perfecting how computers sense multiple fingers simultaneously. But it took time. By 1985 he had published the paper, "The Sensor Frame™: A Device for the Manipulation of Graphic Objects." Though it would take the team until 1991 to get Sensor Frame to market, the document remains an important milestone in the development of multitouch and gesture recognition.

PINCH-TO-ZOOM AND PATENTS

McAvinney's Sensor Frame was the first multitouch screen available to the public, but his 1985 document also noted a gesture for zooming that used a "pinch-like" motion similar to today's commonly used pinch-to-zoom technique. His two-fingered gestures, "zoom in" and "zoom out," controlled just that — zooming the entire frame.

Manipulating graphic objects on the screen would emerge slightly later with the work of Dean Rubine's 1991 thesis, "The Automatic Recognition of Gestures."

"I had something very similar to pinch-to-zoom in my thesis called 'Rotate-Scale-Translate,'" Rubine said. Though the thesis centered on getting computers to automatically recognize new gestures created by an individual user, within its pages lay a small gesture-recognition chart detailing how a user could use a pinching motion to expand and contract an object.

"The lawyers are concerned with this thing called 'pinch-to-zoom,' which was not a term I used," Rubine said. McAvinney and Rubine's seminal gesture work has proven close enough to pinch-to-zoom, however, that they both frequently serve as consultants for patent lawyers and testify in court cases to determine who created the technology and can patent it.

At the iPhone's unveiling 10 years ago, Jobs famously said, "… It's super smart. You can do multifinger gestures on it. And boy, have we patented it."

But Apple hadn't patented pinch-to-zoom. No one had. Not even Rubine or McAvinney. Because the U.S. Patent Office had yet to issue patents on abstract computational ideas. Patent law on software didn't begin to take hold until after the mid-1990s.

"Apple wanted to claim to invent certain things including pinch-to-zoom," said McAvinney, "and the lawyers asked me 'Do you have prior art on this?' And I said that I did."

The courts did award Apple a limited patent on pinch-to-zoom in 2010, which was then rejected in 2013 by the U.S. Patent Office. The battle continues today, as courts attempt to define which types of ideas and innovations are eligible for patent, and who arrived at them first.

But the question remains — had Jobs been aware of the CMU team's work? His visit to the Music Lab in 1985 included touring the Sensor Frame offices and labs, and seeing its documentation. McAvinney owns the nondisclosure agreement Jobs reluctantly signed. "He signed it in October 1985," McAvinney recalled, "which is a few months after I published that little thing with the 'zoom-in and zoom-out' gestures. That was important to the patent lawyers."

Though the history is murky at best, McAvinney and Rubine quickly note that other researchers worked on multitouch and gesturing during the same time period. One such researcher was Bob Boie of Bell Labs, who often garners credit for developing the first multitouch screen. Another researcher, Wayne Westerman, wrote his Ph.D. thesis at the University of Maryland and went on to work for Apple doing gesture recognition for multitouch. McAvinney says the latter's work was certainly influenced by the early efforts at CMU.

"Apple wanted to claim to invent certain things including pinch-to-zoom," said McAvinney, "and the lawyers asked me 'Do you have prior art on this?' And I said that I did"

"[Westerman] mentioned my work and spent a lot of time talking about Dean's thesis. He really based a lot of his work on Dean and sort of picked up where Dean left off."

Westerman holds the title of multitouch architect at Apple.

McAvinney remains clear on the point that even though these researchers worked independently and tackled problems from varying perspectives, the similarities in their efforts became clear and obvious years later — once they started paying closer attention to each others' work. And while Boie and others receive due credit for the development of the capacitive touch technologies we use today, it took the advent of smartphones and touchpads to catapult multitouch to the forefront by giving it a suitable platform for use. Once that platform emerged, McAvinney said, the gestures became more or less obvious, as evidenced by the similarities in the pinching work of McAvinney Rubine, Westerman and others.

No one really knows what Jobs may or may not have seen that day in 1985, but the CMU research team's influence remains undeniable. And with certainty, ahead of its time. ■

BiasBusters aims to tackle biases we are not aware we hold

BUSTING UNCONSCIOUS BIAS

Meghan Holohan

Imagine two students — Lakisha and Steve — competing for the same assistant position in a lab. Three researchers have met with the students, read their resumes and worked closely with them. The researchers now have to decide which student is best for the position. When it comes to skills and experience, the students are evenly matched.

The first researcher, Bob, shares his thoughts.

"Lakisha is really talented, but she's abrasive. Steve is talented, too, but just needs to learn to be more patient," he says.

"I don't find Lakisha abrasive. She's confident," says Jane, another researcher.

"She's brassy and comes on too strong. Steve is the one who exudes confidence," says Bob.

To the casual observer, this interaction might seem harmless. Two researchers are discussing the merits and disadvantages of the two candidates for the position. But there's coded language underlying Bob's evaluation of Lakisha's traits. "Abrasive," "brassy" and "comes on too strong" can often be gender-based criticisms of women.

But these are subtle biases. Bob doesn't dislike women. On the contrary, he believes that Jane is a valued colleague. So how can people overcome their biases when they're not even aware they hold them?

BiasBusters hopes to help solve this complex problem.

"What is most surprising is how big of a change people can make. They are surprised in themselves."

— Diana Marculescu
Professor and Associate Department Head for Academic Affairs in the Electrical and Computer Engineering Department

BiasBusters aims to tackle hidden biases, and has been engaging the Carnegie Mellon University community since fall 2015.

"We think role-play is central to what we do because it is one way to see something from different points of view," says Carol Frieze, director of Women@SCS and SCS4All, who runs the BiasBusters workshops. "The role-play within sessions are based in real-life experience."

When groups discuss the scenarios, they often end up challenging their own biases.

"It's never the same until you experience it," says Diana Marculescu, professor and associate department head for academic affairs in the Electrical and Computer Engineering Department, who also runs BiasBusters.

Google started the program in 2013, and so far more than 2,000 employees have taken BiasBusters workshops, according to a USA Today article. Additionally, half of its 56,000 employees took part in the 90-minute introductory class. The intro class and BiasBusters workshops tried to reduce unconscious biases. Men still dominate the field at companies like Google, but the company wants to be more inclusive and reduce bias. Google leadership feels the best way to reduce these subtle nagging biases is to hold its own programs aimed at eradicating them.

Since Frieze and Marculescu introduced BiasBusters to campus last year, they've held 72 sessions where they trained 58 facilitators, who can be faculty, staff or students. They estimate that more than 1,395 CMU community members have participated in the program.

Frieze has been interested in tackling unconscious bias since 2014. At the same time, Marculescu was hoping to start a similar program. When Dean Moore arrived at SCS from Google and talked about its internal program, it was kismet. Carnegie Mellon became the first university to host the Google program.

Role-play, which can make some people feel uncomfortable, seems particularly effective, Frieze and Marculescu say. Even though they were more aware of unconscious biases in their own lives, the program has changed them.

"I noticed a lot of improvement — improvement in myself," Marculescu says.

Even though role playing can cause people to confront their own beliefs and biases, few people shut down when learning about unconscious discrimination and bias. Most of the BiasBusters participants join the workshops voluntarily, so they want to improve themselves and remain open to the experience.

"I can honestly say I've been really impressed. It's all voluntary. People come in and they are usually on board," says Frieze.

Most people in the workshops feel as if they have a better understanding of the subtle behaviors and cues that can be perceived as biased. This understanding helps many participants make tangible transformations.

"What is most surprising is how big of a change people can make. They are surprised in themselves," Marculescu says.

Another important element of the program is that many facilitators share their own experiences when people acted in a biased manner toward them — or even when they have used it.

Frieze believes that people learn a great deal from hearing personal stories of bias, and those same people often tell her how moving that aspect of the program is for them.

Even though BiasBusters has helped many people change their views, the leaders say that the goal of the workshop is much more practical.

"We're not about blame and victims, but about building allies. We also say we don't try to change minds or bias. We are giving them tools," Frieze says. "We have had success with cultural change and we know it can happen." ■

Make4All endeavors to level the field for people with disabilities

Revolutionizing Assistive Devices

Linda K. Schmitmeyer

Music has always been an important part of young Wilbur's life: When only four, he had the lead in the annual Christmas musical at the Russian orphanage where he lived. And, as a fourth grader at the Waldorf School of Pittsburgh, he played cello in the school's spring recital — using a prosthetic bow holder made on a 3-D printer and designed by researchers in Carnegie Mellon University's Human-Computer Interaction Institute.

"When I push my arm down really hard, I can move the 'wrist' up," says Wilbur, pointing to a red plastic section between the forearm cover and the holder. "It helps me make a really big sound." Wilbur, who has congenital amputations in both arms below the elbows, likes playing the cello because, as he says, "it has a deep sound that calms me when I'm stressed."

But getting the "wrist" to flex so he could make that sound wasn't easy. It took months of back-and-forth design adjustments between him and his adoptive mother and the Human-Computer Interaction Institute (HCII) researchers, who are part of Make4All, a multidisciplinary School of Computer Science effort to improve the way assistive technologies are made and distributed.

The SCS Moonshots initiative continues our rich history of boundary-breaking innovation. While research projects abound in all seven SCS academic units, those selected as Moonshots have been presented to an interdisciplinary team of faculty and staff who have deemed them the most likely to significantly impact the world in years to come.

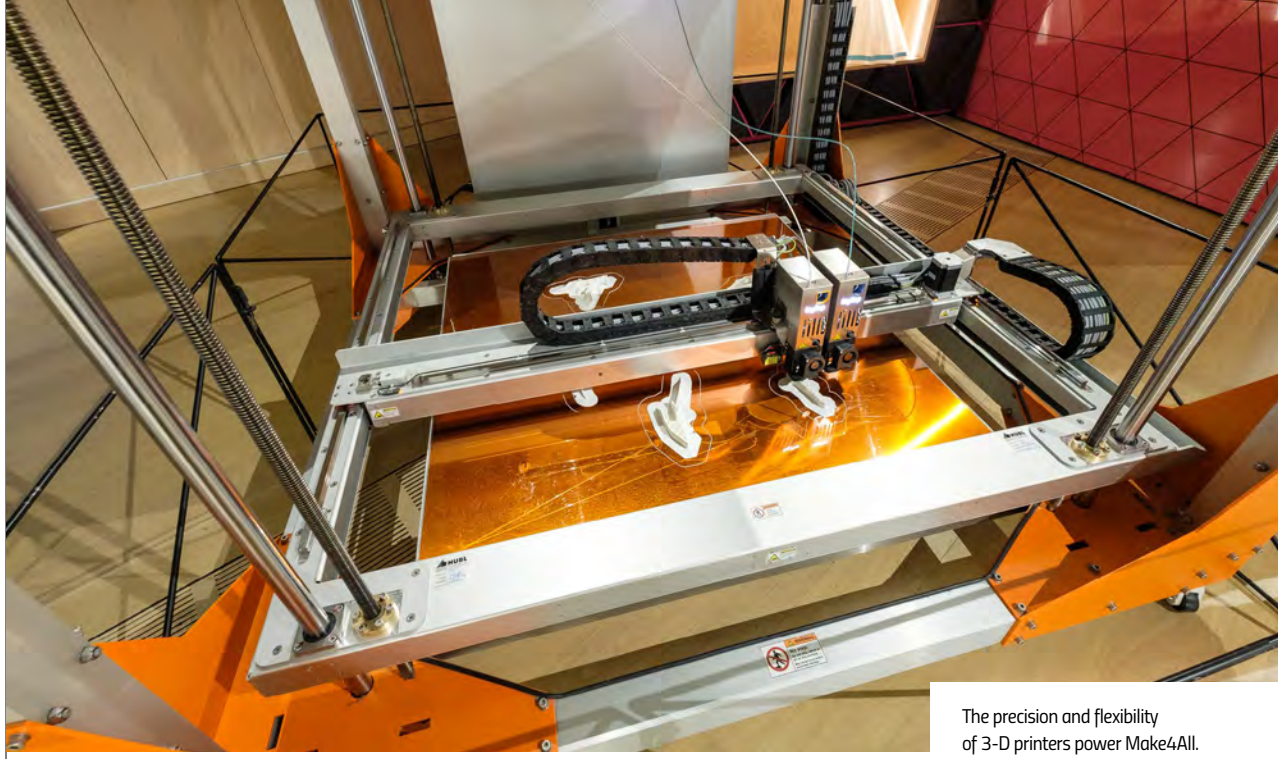

The precision and flexibility of 3-D printers power Make4All.

With one in four working Americans facing a disability sometime during his or her career, and seven out of 10 current workers identifying as more than 80 percent disabled, we need a better way to create customized solutions for people with disabilities for such everyday tasks as holding a spoon or opening a jar — or, for Wilbur, playing the cello. Yet creating effective solutions with today's technology hasn't been easy.

"The needs of people with disabilities are very diverse, and if you give them a prosthetic that doesn't fit their needs, you see abandonment rates of 50 percent or higher," says HCII professor Jennifer Mankoff, a Make4All project co-leader whose research includes developing computer toolkits and processes that enable the creation of such effective applications as assistive devices. "We are trying to change the 3-D modeling process so it will be easier for people who are not highly trained in computer science to achieve high-quality results."

To revolutionize the process, CMU researchers have teamed with the Enable Community Foundation, a global network of volunteer makers, who use 3-D printing technology to produce simple mechanical prosthetic hands for underserved populations around the world. Since its formation several years ago, the foundation's grassroots movement has responded to nearly 2,000 requests for prosthetic hands. Researchers from Carnegie Mellon and several other universities are working together to find ways to streamline the process even further and to develop a more effective way for volunteer makers to customize, produce and distribute these do-it-yourself assistive technologies.

But improving the way makers create computationally customizable physical devices requires significant advances in computer software.

"The CAD tools available today design shape, not function," says Scott Hudson, a Make4All project co-leader and HCII professor whose research interests cover a range of topics involving user interface software and technology. "Currently we have to rely on the user to get from form to function, and that's a huge gap that most people can't easily cross. We need to develop tools that design function as well as form."

To accomplish this, CMU researchers are developing algorithms that will allow software to understand function.

"It will require back-and-forth exchanges between the software program and demonstration videos, design databases and input from the end users in order to specifically create what a user needs," Mankoff says.

A comfortable and flexible fit for Wilbur's prosthetic wrist allows him to play the cello better.

"The goal is to improve the efficiency of building a custom prosthetic device for a particular recipient by allowing multiple people to contribute different skills in measuring, building and fitting the device, while still presenting a single point of contact for recipients and their caregivers."

Building effective teams

Challenges extend beyond fabricating physical devices. To meet the demand for prosthetic hands, the Enable Community Foundation depends on a diverse group of makers from around the world who have a range of skills and time to devote to the project. To make their work more scalable — to help connect more applicants with the devices they need or to expand the devices beyond hands — CMU faculty will create virtual service teams that apply research in machine intelligence, crowdsourcing and collaboration to solve complex problems like those inherent in designing do-it-yourself assistive technologies.

"The goal is to improve the efficiency of building a custom prosthetic device for a particular recipient by allowing multiple people to contribute different skills in measuring, building and fitting the device, while still presenting a single point of contact for recipients and their caregivers," says team member Robert Kraut, the Herbert A. Simon Professor of Human-Computer Interaction. His research includes how open-source communities use information to improve processes for the maker community.

In addition, researchers will place sensors on the prostheses and conduct follow-up surveys with recipients to determine how the devices are being used, and whether they can predict why some devices are abandoned.

Finding ways to make innovation easier

To revolutionize assistive devices, the research has centered on creating real-world objects that serve real-world needs. In addition to Wilbur's prosthetic bow holder, the researchers have designed a prosthetic device that allows a woman with a similar disability to use her hand cycle, a bicycle powered by a hand crank, and adapted an existing prosthesis to hold a knife so another individual can cut his own food.

Although creating these new technologies will be particularly important for customizing devices for people with disabilities, CMU's research has other, far-reaching implications.

"In a deeper way, this research is about creating a different way of working so that people with expertise in many different domains will be able to use these new technologies," Hudson says. "It will help people deliver goods in ways they couldn't have delivered them before." ■

For more information about SCS Moonshots projects, email **moonshots@cs.cmu.edu**.

An aerial drone photographs more easily and efficiently than human assessors.

Drones Assess Aging Infrastructure

Nick Keppler

America's infrastructure is rotting. That's the opinion of the American Society of Civil Engineers (ASCE), which every four years issues an assessment of the roads, railways, dams, seaports, airports, water systems and means of waste disposal that Americans take for granted every day. Its most recent report, from 2013, details "a pressing need for modernization" and "backlog of overdue maintenance across our infrastructure systems." Taking the form of a report card, the 2013 survey graded various kinds of infrastructure and gave the U.S. an overall grade of D+.

"Our roads and bridges are crumbling, our airports are out of date and the vast majority of our seaports are in danger of becoming obsolete," said Steve Kroft in a report for CBS's "60 Minutes" one year later. "All the result of decades of neglect. None of this is really in dispute."

The infrastructure of the U.S. was built mostly in two spurts: one in New Deal-era stimulus programs and the other during the post-World War II economic boom. Since then, the government has been stingy in updating or replacing it. The ASCE's prescribed improvements would cost $3.6 trillion, making it all but impossible they will be implemented by its recommended date of 2020.

A group at the Robotics Institute is developing a tool that will at least help governments and property owners prioritize repairs. AIR Lab's Aerial Robotic Infrastructure Analyst (ARIA) combines 3-D imaging, computer self-navigation and aerial robots (or "drones") to scan aging pieces of infrastructure.

It's a much needed technology for a country sitting on eroding steel and decaying concrete, says Sebastian Scherer, systems scientist for the Robotics Institute. "A lot of infrastructure built in the '50s has reached the end of its useful life," he says. "We need to learn what to prioritize, particularly when it comes to bridges."

There are several devices in use that allow for 3-D scanning for flaws within large structures. Most of these are trucked or boated to the site and moved around manually. The airborne ARIA can fly over, alongside and under structures, giving it the potential to create more detailed images of the cracks and cavities within and leaving no bit unscrutinized.

More than creating a device that can fly, Scherer and his team are trying to create a device that can fly itself. They are looking to combine systems of GPS and independent location recognition to allow the ARIA to take off and scan a structure without human control. Inspectors could merely take it to the site and unbox it. Through either GPS or a downloaded map, ARIA would know its way around the structure. This, Scherer says, is the greatest technical challenge of the project, as the 3-D modeling and flight capacities necessary for the end goal have already been developed.

The Aerial Robotic Infrastructure Analyst (ARIA) rapidly creates comprehensive, high-resolution, semantically rich 3-D models of infrastructure — an interactive assistant for infrastructure inspection.

Perhaps the most anxiety-inducing part of the ASCE's report was finding that nearly 70,000 bridges (one in nine) is structurally deficient. With permission from the Pennsylvania Department of Transportation, which appoints an official to a seat on the project's advisory board, Scherer and his team tested the ARIA on a bridge in Big Run, Pa., in May.

"The main goal was to learn about localization," recalls Weikun Zhen, a graduate student working on the project who manually flew the robot for most of the day. "We did learn some valuable things about how the robot might localize itself beneath beams."

Scherer and his team have used several commercially available aerial robots for the various early incarnations of ARIA, outfitting them all with a main computer and LIDAR, a surveying system that illuminates a target with laser light to measure distance. Working in the High Bay, a warehouse-like machine shop on the lower floors of Newell-Simon Hall, they are currently developing a version with greater battery life and more advanced imaging. Zhen says the next ARIA will be built on a Matrice 600, a larger, steadier aerial robot with six propellers

Burcu Akinci, a professor of civil and environmental engineering, volunteered for the board of the project to add an inspector's perspective. She says the value in a device like ARIA is streamlining inspection. "I think service companies and civil highway and dam inspectors would be able to gather accurate data in a much quicker way," she says. Self-operating inspection robots would be able to clear away backlogs of inspections.

She adds that a complete and accurate picture of all the cracks and cavities might be enough to create the political impetus to repair a bridge or dam. "We can push this [issue] as much as we can with data, but putting a picture before the public is a great way to force the issue." ■

CMU faculty formed Girls of Steel in 2010 with the goal of attracting girls to robotics, a field that has a staggering gender imbalance.

THESE ARE THE GIRLS OF STEEL

Nick Keppler

While the campus of California University of Pennsylvania is silent over spring break, within the Convocation Center, overhead speakers blare out arena rock standards, like "The Boys Are Back in Town," "What I Like About You" and the requisite "We Will Rock You." Squads of teenagers walk around in protective eyewear and t-shirts with team colors. This is For Inspiration and Recognition of Science and Technology's (FIRST) Greater Pittsburgh Regional Conference.

The idea behind FIRST Robotics Competitions is to make the kids who are into computers and engineering teen champions for a day. Because of this, the games borrow a lot of atmospheric cues from sports.

When the match starts, the six-wheeled, 120-pound robots roll onto a field that takes up half the basketball court to the sound of bugle horns. Teens on the sidelines armed with remote controllers control some of their motions. (Others are automated.) The robots — designed, built and

Among students participating in the DARPA Robotics Challenge, perhaps the most important college competition, just 5.2 percent were women. Just like the original Rosies who took manufacturing jobs for the war effort, women in today's tech jobs are broaching into extremely male-dominated fields.

programmed by the students with the help of mentors — must complete three tasks for this year's game: shoot plastic balls into a tall object with a horn-shaped opening on top, pass an oversized novelty gear to a robotics team member standing atop a platform and climb a rope onto the same platform. Adult referees, dressed like sports referees, enforce rules, like no robot-on-robot contact.

Each match requires only a handful of team members (and matches are played all day, until a ranking is compiled). The others await their turn and cheer from the stands.

The crowd, as loud as that of any high school basketball game, is a patchwork of colors: the pink of Pittsburgh's City Charter High School, the black of Chartiers High School, the red of Chambersburg's schools. One team, whose members have taken the front and center rows, wasn't organized at a high school. Instead of school colors, they wear polka-dotted red bandanas and blue flannel — the iconic costume of the World War II "We Can Do It!" Rosie the Riveter poster. The team is all female.

These are the Girls of Steel (GOS). The team was formed in 2010 by Carnegie Mellon University faculty with the goal of attracting girls to a field that has a staggering gender imbalance.

Only 14 percent of American engineers are women, according to a 2012 congressional report. Robotics might be even more male-dominated; in 2015, The Washington Post reported that, among students participating in the DARPA Robotics Challenge, perhaps the most important college competition, just 5.2 percent were women. Just like the original Rosies who took manufacturing jobs for the war effort, women in today's tech jobs are broaching into extremely male-dominated fields.

This comes not from aptitude but from childhood social cues, said George Kantor, a senior systems scientist at CMU's Robotics Institute, who co-founded Girls of Steel. "Girls often play with Barbies and boys play with hammers," said Kantor. "Before getting onto our team, they may not have had any [technology] role models."

Robotics Institute Outreach Manager Patti Rote, the other co-founder, said the team is part of CMU's efforts to engage promising technology students of all ages and genders. "It seemed like a lot of girls hadn't been asked to do something like this," she said, "so we decided to see what would happen if we asked." She gathered donations from robotics professionals and started recruiting home-schooled girls and students from the Ellis School, Pittsburgh's longstanding all-girls prep academy. Rote then moved on to local high schools that lacked their own robotics team.

Today, GOS had a budget of about $120,000, from sponsorships and donations, and has a set roster of 60 girls in grades eight through 12.

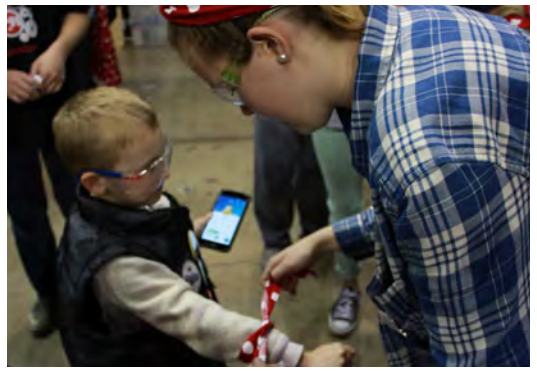

GOS is active all year, holding demonstrations and visiting middle and elementary schools to teach robotics to younger students.

For about six months, they are focused on FIRST competitions. In October, they hold their initial meetings and form squads for design, mechanics, programming and electronics. In January, FIRST sends the tasks outlined in this year's game, and GOS, just like more than 50,000 robotics teams across the country, creates a robot that can complete the tasks. They meet mostly in the National Robotics and Engineering Center, an off-campus arm of the Robotics Institute.

Corinne Hartman, a junior at the Ellis School, said at peak time she dedicates 10 to 20 hours a week to the sport. She's on the programming squad. "We just tested the programming again and again," she said.

GOS's robot for this competition is named Mo, after former first lady Michelle Obama, "because we really like what she did for STEM and equality in education," said Girl of Steel Alexis Selwood, a junior at Fox Chapel Area High School.

After six weeks of work on the robot, Mo was bagged until competition time, as per FIRST rules. The team practiced using a copy called Bo (named after the former president). At the regional conference, the Girls of Steel were ranked ninth out of 39 teams.

Cheyenne Meyers, a senior at Pittsburgh's Brashear High School, has participated in GOS since she was a freshman. Meyers, who wants to be an engineering educator, said girls in technology are often underestimated. She is one of three girls in Brashear's shop class, an elective with about 30 students total. "I feel that girls are put down," she said. "In class, the boys will ask, 'Do you want me to drill that for you?'"

Molly Urbina, a former Girl of Steel who is now an undergrad at the Illinois Institute of Technology, said the group taught her not just technical knowledge, but also teamwork and social skills. Urbina had been home-schooled until grade nine and began on the team at the same time she started at the private Winchester Thurston School.

"I was very timid," she recalled. "But a lot of that changed once we had to work on presentations together." At competitions, her teammates "dragged" her to talk to other teams, teaching her how to ask for advice from peers when completing an engineering task. Urbina said that at Illinois Institute of Technology she has learned that much of excelling in a technology field comes from the ability to work within a team. "Without Girls of Steel, I don't think I would have that." ■

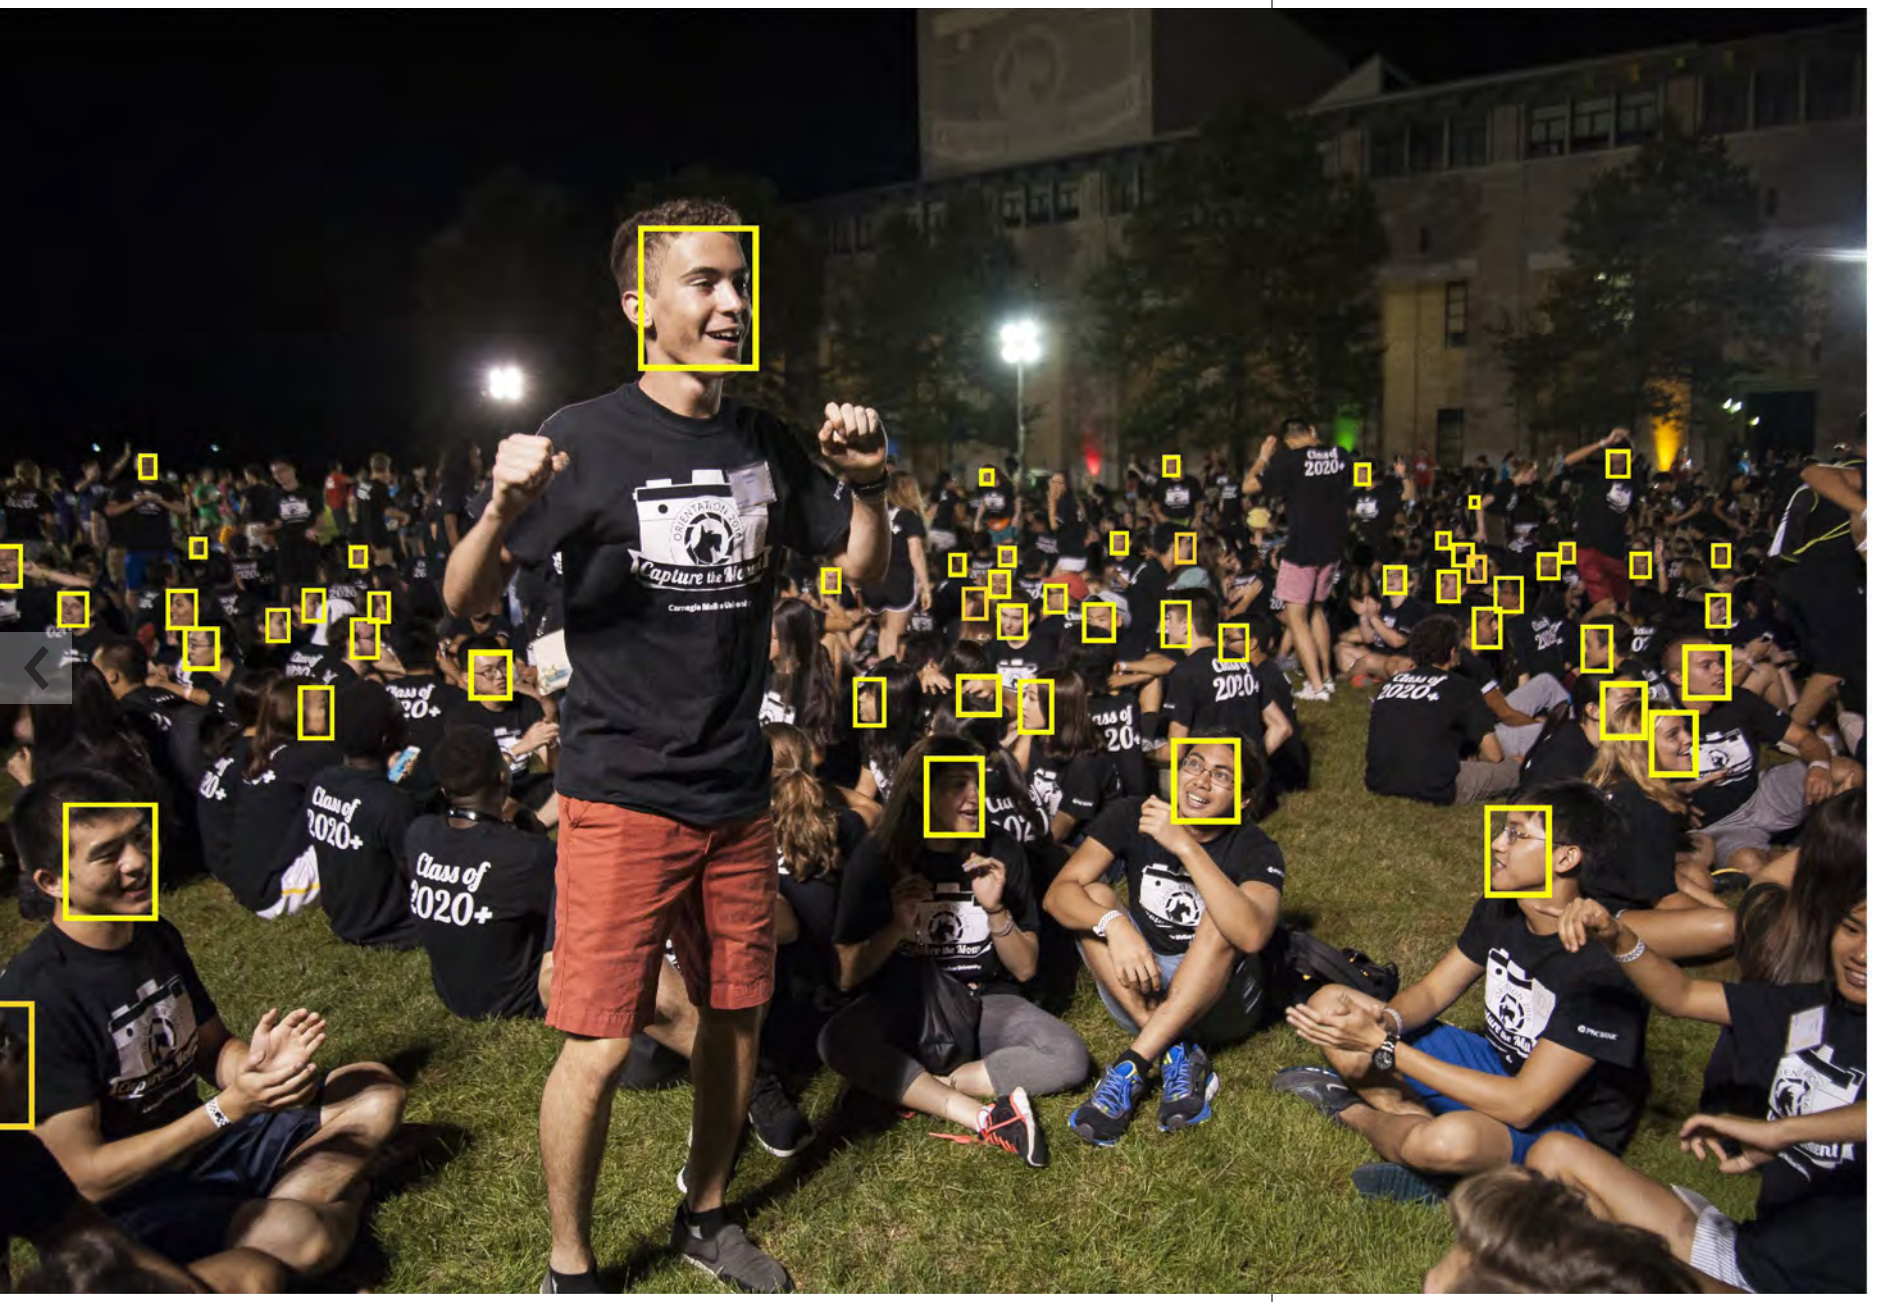

Finding Faces in a Crowd

Using computer vision systems to spot a face in a crowd, or recognizing any small or distant object within a large image, can be a major challenge. The trick to finding tiny objects, say researchers at Carnegie Mellon, is to look for larger things associated with them.

An improved method for coding the necessary context from an image has enabled Associate Professor of Robotics Deva Ramanan and robotics Ph.D. student Peiyun Hu to demonstrate a significant advance in detecting tiny faces.

When applied to benchmarked datasets of faces, their method reduced error by a factor of two, and 81 percent of the faces found using their methods proved to be actual faces — compared with 29 to 64 percent for prior methods.

"It's like spotting a toothpick in someone's hand," Ramanan said. "The toothpick is easier to see when you have hints that someone might be using a toothpick. For that, the orientation of the fingers and the motion and position of the hand are major clues."

Finding a face that may be only a few pixels in size requires first looking for a body within the larger image, or realizing the image contains a crowd of people.

The method that Ramanan and Hu developed uses foveal descriptors to encode context in a way similar to how human vision is structured. Just as the center of the human field of vision is focused on the retina's fovea, where visual acuity is highest, the foveal descriptor provides sharp detail for a small patch of the image, with the surrounding area shown as more of a blur.

By blurring the peripheral image, the foveal descriptor provides enough context to help understand the patch shown in high focus, but not so much that the computer becomes overwhelmed. This technique allows Hu and Ramanan's system to make use of pixels that are relatively far away from the patch when deciding if it contains a tiny face.

Spotting tiny faces could have applications such as doing headcounts to calculate the size of crowds. Detecting small items in general will become increasingly important as self-driving cars move at faster speeds and must monitor and evaluate traffic conditions in the distance. ■

DIRECTOR'S MESSAGE
OFFICE OF ENGAGEMENT AND ANNUAL GIVING

ROBOTS ROCK, AND PEOPLE ARE AWESOME, TOO

When some people think about computer science at Carnegie Mellon, they think about robots. A few times a year, visitors show up at our office door and ask where we keep the robots and if they can see them. The media call and ask if they can use our robots in the news. Groups of high school students tour campus and want to see the robots.

Don't get me wrong — we love our robot reputation. We love coming to work in the Gates and Hillman Centers and watching a CoBot bustling through the building with a basket of candy on its back or attempting to board an elevator. We love seeing CMU creations in the news. We love that interacting with robots can inspire kids to become computer scientists when it otherwise may not have crossed their minds as a career option.

Niccole Atwell, Senior Associate Director of Engagement and Annual Giving and Ashley Patton

But we're so much more than robots in SCS. Behind everything we do — robots and other artificial intelligences included — are teams of amazingly talented individuals dedicated to improving how we live and work. You've already met some of them in this issue of The Link. They're creating drones that inspect bridges, encouraging women in STEM, helping us become aware of our unconscious bias, connecting people who need artificial hands with the makers creating them via 3-D printing. They're leaving their impressive marks on the world, and they're what keep us excited about what we do, day in and day out.

In the pages that follow, you'll learn more about our extraordinary SCS people — why one of our alums gives back, names that have landed in the news and exciting research from people around the school. Plus we'd be remiss if we didn't include some photos welcoming the class of 2017 into our alumni fold.

Robots rock, we don't deny it. But our people are awesome, too, and we're happy to celebrate them in this issue of The Link. ■

Ashley Patton
Director of Engagement
and Annual Giving
School of Computer Science
Department of Electrical
and Computer Engineering
awpatton@cs.cmu.edu

David Murray (CS 2006)

Giving Back to SCS

Unquestionably, I would not be where I am today in life, my career and otherwise, without CMU. I give back each year because I want to make sure future generations have the opportunities I had during my time here. I genuinely believe that every dollar I give further perpetuates all of this awesomeness.

There are so many other universities that lack the culture we have. With all our quirkiness, our hardcore work ethic and our amazing cross-disciplinary approach, we are one of a kind in such a wonderful way! Keep in mind that we don't have the kind of crazy endowments that other top-tier universities have had for hundreds of years, so if we can be where we are without that head start, imagine where we could be if every alum gave!

Some years I have more to give than others, but I have a reminder on my calendar each year to give, and it feels great every time I do. I highly recommend you take a moment right now to put that recurring reminder on your calendar — I guarantee you'll feel happier afterwards! ■

Rong Yan

Mark Roth

Rong Yan, 37, is the director of product engineering for Snap Inc., the company behind Snapchat. He graduated in 2006 with a Ph.D. from the School of Computer Science in language and information technologies. The native of China also earned his master's degree from CMU in 2004.

What do you think makes Snap a unique company?

Photo and video messaging was the starting point for Snapchat, but now it has transformed itself into a camera company. We want the Snapchat camera to be used by millions of people every day throughout the world. Snapchat users already are taking billions of photo and video Snaps every day. Another key part of the company is its development of creative tools so people can have fun and truly enjoy image creation and image sharing.

What are some of the company's most popular features?

Lenses are one of our most popular tools. Lenses decorate your face and the world around you in 2-D and 3-D animated experiences. Even President Obama has used Lens. It's an augmented reality feature, and augmented reality is essentially what Snapchat has been working on for years already.

What is the difference between augmented reality and virtual reality?

I rank augmented reality higher than virtual reality. Augmented reality is imposing digital content on top of the existing reality that you see and experience. Virtual reality, on the other hand, is creating a completely virtual experience digitally. The problem with VR is that it separates people from reality and has to be done with a device and in a space that takes them away from their everyday experiences. But augmented reality can be applied to your daily life. You can make it mobile, lightweight, and allow it to impose information or other filters over top of the scenes you see before your eyes.

How did the School of Computer Science prepare you for your present career?

I joined CMU in 2001, and spent five-and-a-half years there. I worked on the Informedia project for retrieving multimedia information from digital archives, and my advisor was research professor Alex Hauptmann. My experience at Carnegie Mellon taught me everything about video processing and image searching. Just as importantly, CMU taught me the mindset of always turning research insights into real-world applications for people to use.

You began your career working for the IBM T.J. Watson Research Center before moving to Facebook, Square and now Snap. Why did you choose to go into industry research?

It has turned out to be the right fit for me. I think research has a very different meaning in 2017 compared with the 1980s. Not as many people are doing research anymore just for the sake of getting peer recognition for writing papers — they're doing research to achieve innovations. Today, that goal of knowledge transfer is woven into research activities from the start.

Snapchat developed its reputation as a service that allowed people to share messages that would not be stored permanently. Is that still a guiding principle?

Yes. But the privacy aspect is only one of the reasons for that feature. What the temporary nature of the images does is allow you to create a fun experience without worrying about it staying there for a long time, and that gives people greater creative liberty.

What are examples of that?

Our community has been super creative. Snapchatters can decorate their photo and video Snaps with doodle cartoons or other types of art, sometimes adding things like captions, Stickers or Geofilters. One of our most recent creative tools is Custom Stickers, which let you use touch to cut out an area of a photo to use as a sticker in other Snaps and Chats.

It seems a major goal of Snapchat is to enhance users' ability to have fun with images. How do you see that?

We want people to be entertained. Snapchat has a younger demographic than some other social media platforms, but the company believes that by continually innovating with this entertainment goal in mind, it will be able to keep all those customers as they age. Our belief is that having a fun experience is never a younger-generation-only experience.

What makes you enjoy coming into work each day?

At Snapchat, it's innovation. We're constantly learning from our community and thinking six months or even years ahead about what new features we can develop that will please our community. ■

Maitreyee Joshi

Aisha Rashid (DC 2019)

When rising senior Maitreyee Joshi first arrived at the School of Computer Science, she assumed her life would be all CS all the time. That the time had come to bid adieu to her passion for things like art and acting. She was happy to discover that she was wrong and that a CS education at Carnegie Mellon allows people to pursue other passions.

Joshi originally chose to study computer science because of her desire to use technology for social good. Previous projects include developing an indoor navigation system app for the blind, improving prosthetic systems for amputees, and a biclustering algorithm that detects cancer-causing mutations.

This summer, Joshi adds another passion to the list of those CMU has allowed her to pursue as she undertakes an internship with Talent Beyond Boundaries, a D.C. based nonprofit organization founded through the 2014 Advanced Leadership Initiative at Harvard University. There, Joshi will assume a dual role — research policy intern and tech intern — to help connect refugees to the international labor market. She will explore the real-life applications of technology, from understanding how labor markets work in other parts of the world and the best ways to connect refugees to those markets, to building the technical infrastructure that creates those connections.

She chose to pursue this policy internship because she believes that she can make an even bigger social impact on people's lives through policy. Specifically, her passions lie in protecting human rights and studying international human rights law. Joshi wants to affect policy that guarantees the world's most vulnerable populations are afforded the basic rights and freedoms that all humans are entitled to.

Joshi took a risk when she rejected offers from tech companies to find a more policy-related opportunity. She notes that making such a decision — not unheard of among CS majors, but certainly rare — would have been impossible without the support of CMU's Career and Professional Development Center, her computer science advisors, SCS dean Andrew Moore and the CMU alumni network.

Joshi believes her computer science studies will allow her to bring a fresh perspective to nonprofit work that other interns might not have. "Technology is vital and helps make things more efficient, scalable and accessible," she said. She also recognizes that much of nonprofit work involves interfacing with corporations, and thinks her previous work at big tech corporations like Microsoft will benefit her and provide a unique perspective in the nonprofit world.

Joshi hasn't quite decided what lies beyond CMU yet, but law school may be on the horizon. This summer's internship could help her determine if that is the path she would like to take. Law school or not, she is determined to get the most from her experience.

"The skills I learn this summer will be invaluable to me, no matter where I go," Joshi said. "And at the end of the day, they're something tech companies will also value." ■

On May 21, the School of Computer Science proudly welcomed the Class of 2017 into the fold of CMU Alumni. Congratulations to all graduates, families and friends.

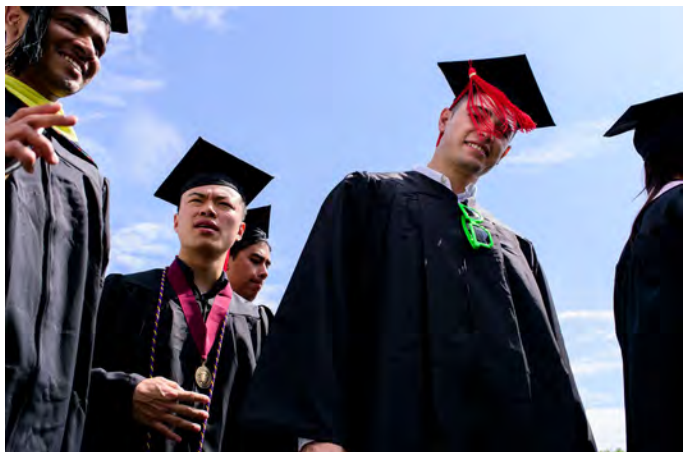

Self-Driving Buggies Reach Historic Milestone

During this year's Spring Carnival Sweepstakes, SCS made innovative history within long-standing CMU Buggy tradition. For the first time in over 90 years, two self-driving buggies participated in the Buggy race. The race is a unified effort of mechanics, pushers and drivers, who work to build an un-motorized vehicle, push it uphill during the first half of the course, and drive it downhill for the second half at speeds of up to 35 mph.

Team RoboBuggy and the Carnegie Involvement Association (CIA) took teamwork and ingenuity to the next level, showcasing their interdisciplinary talent and applying robotics engineering to change Buggy tradition.

The RoboBuggy project sprung from a conversation between Matt White (CS 1996), and Mark Stehlik, the associate dean for outreach in SCS. They decided to use technology from CMU's Field Robotics Center's NavLab project and adapt it for Buggy, an idea that White later explored in his independent study.

The project had a recent revival in 2013 within the CMU Robotics Club, in addition to CIA's work on robotic buggies as a part of the Atlas Project.

"It was surreal, honestly, watching it steer itself after having failed every other time," said Benjamin Warwick, a junior in mechanical engineering and the founder of the Atlas Project. "What I'm hoping this will do is encourage other people to start robotic buggies of their own."

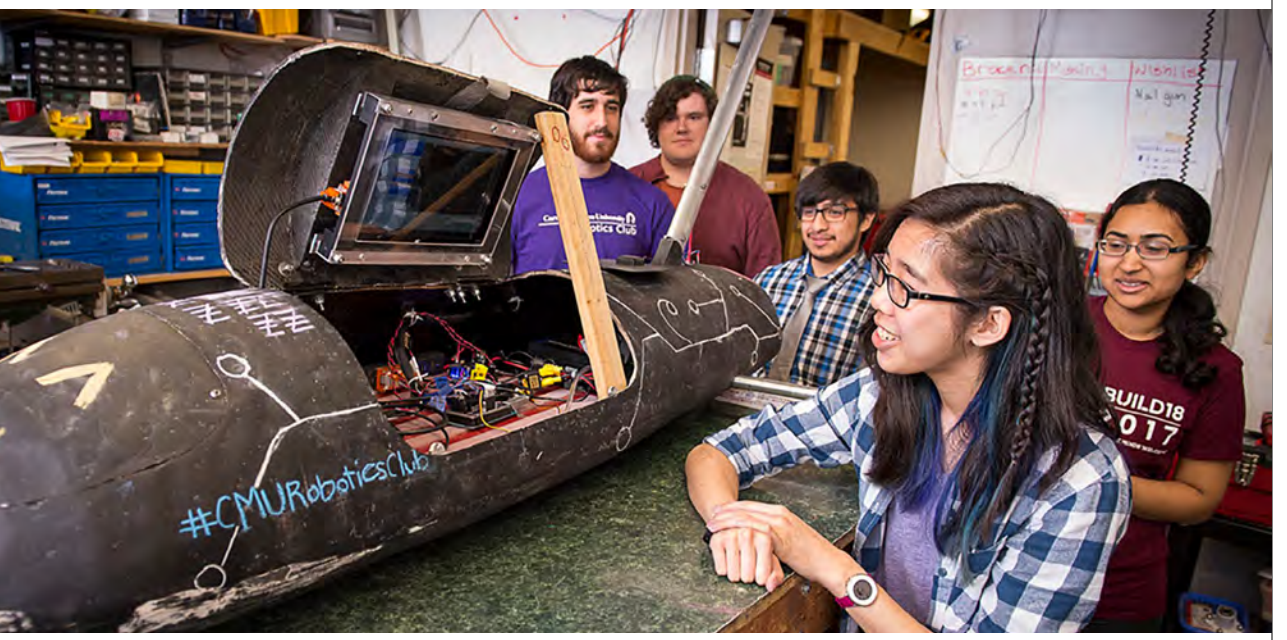

Search Engine Uses AI to Find Lawyers

At the click of a button, search engines provide internet users with whatever resources and contacts they need. Justice Toolbox Inc., a Maryland-based startup launched this past November, does exactly this by helping everyday people find lawyers. The search engine, founded by Bryant Lee (CS 2008), uses data mined from official state court records to compute and display how many cases a specific lawyer has won and lost, as well as their approximate win rate.

Justice Toolbox allows users to obtain data from close to five million state court records. Free of charge, the search engine provides information on more than 70,000 lawyers and allows searching for 180 case types, a refined search feature not found in any other product on the market. Justice Toolbox uses a custom-designed artificial intelligence (AI) program that analyzes court records used by lawyers, or "dockets." The AI reads and understands these dockets just as a lawyer would to determine the case outcome, case type and the attorneys involved.

"Choosing counsel is one of the most important decisions a person can make in their legal case," Lee said. "People should have the benefit of seeing an attorney's public court records before making that decision."

Delphi Group Launches 2016–2017 Flu Forecasting Models

CMU's Delphi research group used AI and the wisdom of crowds to guide its efforts to forecast 2016–2017 flu activity. The team, comprising faculty and students from the machine learning, statistics, computer science and computational biology departments, is part of a research initiative with the U.S. Centers for Disease Control and Prevention to develop methods of accurately forecasting flu activity.

During the 2015–2016 flu season, three CMU forecasting systems proved the most accurate, besting 11 competing systems fielded in the initiative. The group's top-ranked system was a model that uses machine learning to make predictions based on past patterns and input from the CDC's domestic influenza surveillance system. Its second-ranked system — which topped the 2014–2015 forecasting systems — uses weekly human predictions that reflect the wisdom of crowds.

Forecasting the flu is just the beginning. The group hopes to apply their forecasting tools to such diseases and conditions as HIV, drug resistance, Ebola, Zika and Chikungunya.

Researchers Seek Solution to Selfie-Related Deaths

Most of us are guilty of taking selfies, but the seemingly harmless activity can actually prove fatal. A growing number of people die each year while snapping selfies on cliffs, railroad tracks and other hazardous spots. Researchers in Pittsburgh and India are looking for ways to reduce this risk.

In a new study, researchers at CMU and the Indraprastha Institute of Information Technology in Delhi (IIITD) scoured public records to compile a list of 127 deaths associated with selfies worldwide between March 2014 and September 2016. After analyzing those selfie deaths, they designed a system that uses location, image and text to classify whether a selfie was dangerous or not.

One of the primary goals of the research, launched by Ponnurangam Kumaraguru, an associate professor at IIITD who earned his Ph.D. in computer science at CMU in 2009, is to develop an app that could warn users when they're preparing to take a hazardous selfie or even temporarily disable a smartphone's selfie function. Another goal is to use this data to inform public policy, perhaps establishing "no selfie zones" in hazardous areas or creating public education programs to help people recognize dangerous selfie practices.

Norman Sadeh

App Behavior May Appear at Odds With Privacy Policies

How an app says it will collect or share your personal information may not be the same as what it actually does, according to a new automated analysis system developed by Institute for Software Research professor Norman Sadeh and post-doc Sebastian Zimmeck. Their system analyzes the text of privacy policies, then examines the app's code to see whether its behavior suggests it shares personal information — indicating it should have a privacy policy. It also checks to see if the app's data collection and sharing behavior is consistent with an existing privacy policy.

After analyzing almost 18,000 popular free apps from the Google Play store, the system found that nearly half lacked privacy policies, even though 71 percent of those appeared to be processing personally identifiable information and would be required to explain how they do so under the laws of some states.

"Overall, each app appears to exhibit a mean of 1.83 possible inconsistencies and that's a huge number," Sadeh said. The number of discrepancies is not necessarily surprising to privacy researchers, he added, "but if you're talking to anyone else, they're likely to say 'My goodness!'"

RoboVote Uses AI To Help Groups Make Decisions

Voting is on the brain these days. But a new online service has nothing to do with the presidential election. Instead, it enables anyone to use state-of-the-art voting methods to make optimal group decisions.

RoboVote, a project of CMU and Harvard researchers, doesn't just tabulate votes. Rather, it uses artificial intelligence and decades of social choice research on how opinions, preferences and interests can best be combined to reach a collective decision. RoboVote can be used for anything from helping friends determine the best place for pizza to aiding a product development team as it decides which idea is best suited for commercialization.

"We're leveraging the latest work in optimization and AI to help people make decisions in their daily lives," said Ariel Proccacia, assistant professor of computer science. "We have taken what years of research have proven to be the best algorithms for making collective decisions and made them available with an interface that anyone can use."

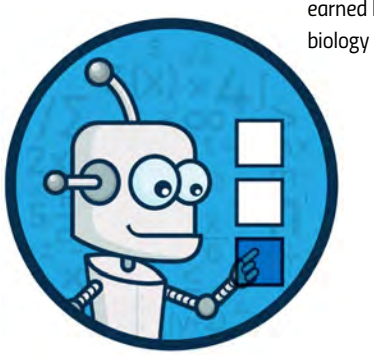

Molecular Conductors Help Plants Respond to Drought

When plants need water their leaves droop and they start to look dry. But what's happening on a molecular level? Researchers at the Salk Institute for Biological Studies and the SCS Computational Biology Department have found answers that could be critical to helping agriculture adapt to drought and other climate-related stressors.

The new research suggests that in the face of environmental hardship, plants employ a small group of proteins that act as conductors to manage their complex responses to stress. The results, detailed in Science magazine, may help in developing new technologies to optimize water use in plants.

The research team was led by Joseph Ecker, director of Salk's Genomic Analysis Laboratory, and included Ziv Bar-Joseph, professor of computational biology and machine learning, and Aaron Wise, who earned his Ph.D. in computational biology at CMU in 2015.

Carnegie Mellon Will Lead $253 Million Advanced Robotics Venture

Pittsburgh will be a hub for advanced robotic manufacturing, thanks to a new $253 million grant from the Department of Defense. The CMU-founded independent institute, American Robotics, includes 220 partners in industry, academia, government and the nonprofit sector nationwide. The institute will receive $80 million from the DOD, and an additional $173 million from partner organizations.

Howie Choset, a professor in the Robotics Institute, played a large role in conceiving of and establishing the institute. He said its four-pronged mission is to empower American workers to compete with low-wage workers abroad; create and sustain new jobs to secure U.S. national prosperity; lower the technical, operational and economic barriers for small- and medium-sized enterprises as well as large companies to adopt robotics technologies; and assert U.S. leadership in advanced manufacturing.

"This work has the power to benefit society broadly, and to benefit many, many potential workers and their families," Choset said.

Researchers Unleash Robots That Learn Like Babies

Babies learn about the world by pushing and poking objects, putting them in their mouths and throwing them. CMU scientists are taking a similar approach to teach robots how to recognize and grasp the objects around them.

Manipulation remains a major challenge for robots and has become a bottleneck for many applications. But Robotics Institute researchers have shown that by allowing robots to spend hundreds of hours poking, grabbing and otherwise physically interacting with a variety of objects, the robots can teach themselves how to pick those objects up.

The researchers, led by Abhinav Gupta, assistant professor of robotics, are scaling up this approach with help from a three-year, $1.5 million Focused Research Award from Google.

"We will use dozens of different robots, including one- and two-armed robots and even drones, to learn about the world and actions that can be performed in the world," Gupta said. "The cost of robots has come down significantly in recent years, enabling us to unleash lots of robots to collect an unprecedented amount of data on physical interactions."

CMU Among First Universities in New Alexa Fund Fellowship Program

Amazon has announced a new Alexa Fund Fellowship to support universities and researchers focused on transformative speech technologies, and Carnegie Mellon is among its first four participants. Ran Zhao, a Ph.D. student in the Language Technologies Institute, has been named CMU's first Alexa Fund fellow.

The Alexa Fund Fellowship, named for Amazon's Alexa voice service, supports researchers working in areas such as text-to-speech, natural language understanding, automatic speech recognition and conversational artificial intelligence. As part of the fellowship program, Amazon provided Alexa-enabled devices for use this spring in an LTI project course taught by Research Professor Emeritus Alex Rudnicky and Professor Alan Black, with assistance from Zhao.

"It's a moment in time when voice interaction is a matter of increased interest," Rudnicky said. With a number of voice-enabled systems now on the market, researchers are interested in pushing the technology as far as possible, while companies look for both new applications and new graduates who are accustomed to developing dialogue systems, he and Black noted.

CMU-Q, Robotics Institute Put Negotiation Under the Microscope

Researchers at Carnegie Mellon in Qatar and Pittsburgh have joined forces to explore the behavior of successful negotiators — very, very closely.

In the study, 10 students from CMU-Q engaged in a mock trading exercise in the Robotics Institute's Panoptic Studio, a two-story geodesic dome fitted with nearly 500 synchronized video cameras, to capture the fine details of human interaction — eye movements, eye contact, hand gestures and subtle body language. The hundreds of cameras in the dome captured the negotiators' every movement, generating a half terabyte of data from the cameras and audio feeds each minute.

"If we can understand what behaviors lead to success, we can train business students to be better negotiators," said John O'Brien, associate dean at CMU-Q and associate professor of accounting. "Down the road, these insights could lead to augmented reality applications."

New Computational Method Makes Gene Expression Analyses More Accurate

The gene expression analyses used to diagnose and monitor cancer have become more accurate thanks to a tool from Carnegie Mellon researchers.

Working with colleagues from Stony Brook University and the Dana-Farber Cancer Institute, Associate Professor of Computational Biology Carl Kingsford has developed Salmon, a computational method that can correct for the technical biases known to occur during RNA sequencing (RNA-seq) — the leading method for estimating gene expression. The report on their findings appeared in Nature Methods.

"Salmon provides a much richer model of the RNA-seq experiment and of the possible biases that are known to occur during sequencing," Kingsford said. This is important, he added, because the technique is increasingly used for classifying diseases and their subtypes, understanding gene expression changes during development, and tracking the progression of cancer.

Carl Kingsford

Names in the News

Allison Wang

Serena Wang

Enxu Yan, Akhil Chaturvedi , Adarsh Prasad

Brad Myers

Scott Hudson

Lorrie Faith Cranor

Harry Shum

Takeo Kanade

Jeff Bigham, Carol Frieze

Faciometrics

Daisy, a hexapod robot built by CMU spinoff Hebi Robotics, made its debut as part of the Robot Revolution exhibit May 11 at Chicago's Museum of Science and Industry.

SCS undergraduates Allison Wang and Serena Wang are among 54 students nationwide named 2017 Kleiner Perkins Caufield and Byers Engineering fellows.

Two SCS students outcoded more than a hundred of their peers at the 2017 Global CodeCon Finals. Raymond Kang, a junior studying computer science and mathematical sciences, earned ninth place; and Gabriele Farina, a Ph.D. student in computer science, came in 13th. Kang took first place in last year's competition.

Machine Learning Ph.D. students Adarsh Prasad and Enxu Yan were part of a team that won the $25K prize in the CMU Department of Statistic's Citadel/Citadel Securities Datathon. Hui Han Chin, a graduate student in computer science, was part of the runner-up team.

Three SCS faculty members have been honored by the Association for Computing Machinery's Special Interest Group on Computer-Human Interaction. Scott Hudson and Brad Myers, both professors in the Human-Computer Interaction Institute, will receive the organization's Lifetime Service Award and Lifetime Research Award, respectively. Lorrie Faith Cranor, professor of computer science and engineering and public policy, has been named to the CHI Academy.

Harry Shum, executive vice president of Microsoft's Artificial Intelligence and Research group, has been elected as a foreign member of the prestigious National Academy of Engineering (NAE). Shum earned a Ph.D. in robotics at Carnegie Mellon in 1996.

Takeo Kanade, the U. A. and Helen Whitaker Professor of Robotics and Computer Science, has been named the 2017 recipient of the IEEE Founder's Medal — one of IEEE's highest honors. The medal recognizes Kanade "for pioneering and seminal contributions to computer vision and robotics for automotive safety, facial recognition, virtual reality and medical robotics."

A team of CMU students and alumni took first place in the 2016 Facebook Global Hackathon. Team members included Abdelwahab Bourai, who earned his bachelor's degree in computer science this past May and is now a master's student in computational data science; Jean Haddad, a junior in chemical engineering; Sebastian Guerrero Cardenas, a junior majoring in information systems and human-computer interaction; and SCS alumnus Vivek Krishnan.

SCS alumni Matt Humphrey and Xi Chen are among Forbes magazine's 30 Under 30 for 2017. Humphrey, who earned a bachelor's degree in computer science as well as an MBA at CMU, is a serial entrepreneur who founded LendingHome. Chen earned a Ph.D. in machine learning and is an assistant professor of information, operations and management sciences at NYU.

Jeff Bigham and Carol Frieze were both recognized as recipients of the 2016 AccessComputing Capacity Building Award, which honors collaborators who work to advance students with disabilities in computing fields. Bigham is an associate professor in the Human-Computer Interaction Institute, and Frieze directs Women@SCS and SCS4All.

Facebook has agreed to acquire Faciometrics, a spinoff from CMU's Robotics Institute that develops facial analysis software for mobile applications. Fernando De la Torre, associate research professor of robotics and head of the Human Sensing Laboratory, is the company's founder and CEO.

then & NOW

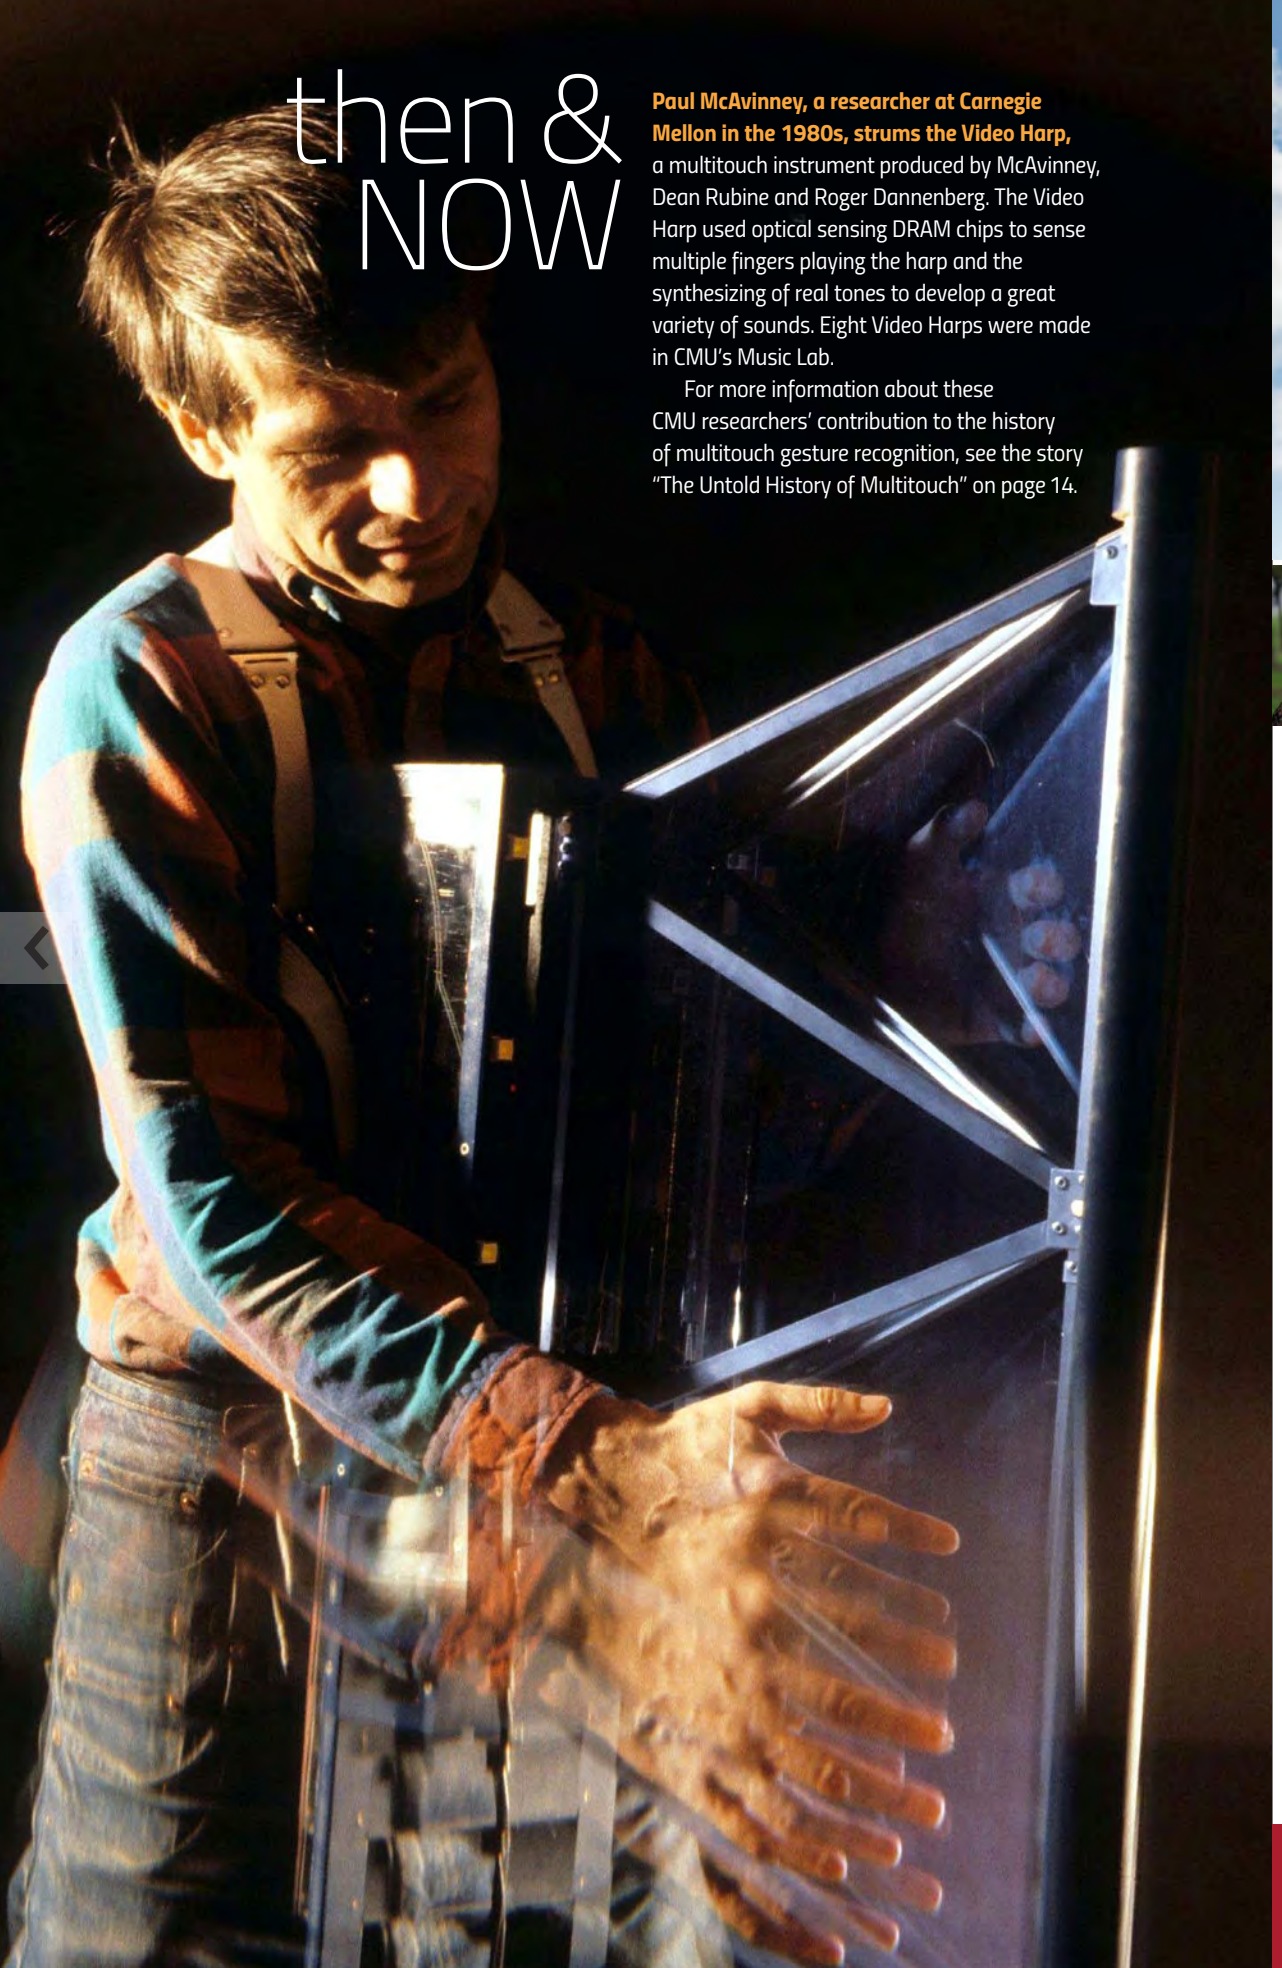

Paul McAvinney, a researcher at Carnegie Mellon in the 1980s, strums the Video Harp, a multitouch instrument produced by McAvinney, Dean Rubine and Roger Dannenberg. The Video Harp used optical sensing DRAM chips to sense multiple fingers playing the harp and the synthesizing of real tones to develop a great variety of sounds. Eight Video Harps were made in CMU's Music Lab.

For more information about these CMU researchers' contribution to the history of multitouch gesture recognition, see the story "The Untold History of Multitouch" on page 14.

Carnegie Mellon University
School of Computer Science

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

SCSatCMU

www.cs.cmu.edu

calendar of events

August 15
Summer Session
Final Grades Due

August 19
First-Year Orientation
and Move-In Day

August 28
Semester & Mini-1 Classes Begin

September 4
Labor Day; No Classes

October 13–15
Family Weekend 2017

November 11
Carnegie Mellon University
Homecoming