

**SEPTEMBER 30 – OCTOBER 1, 2025**

# Charting the Future of AI in Virginia




# Virginia Academy of Science, Engineering, and Medicine

The Virginia Academy of Science, Engineering, and Medicine is a nonprofit organization consisting of members of the National Academies of Science, Engineering, and Medicine who reside or work in Virginia as well as Virginians who are leaders in these fields. Through its nonpartisan network of experts, the Virginia Academy provides rigorous analytical, technical, and scientific support to inform policy on issues critical to the Commonwealth.

The Virginia Academy also promotes research, fosters interchange among individuals and organizations, and recognizes and honors Virginians who have made major contributions to science, engineering, and medicine.

*Executive Director: Nat Draper*  
*Publishing and Writing: Charles Feigenoff*  
*Editing: Rosalind Hingeley*  
*Design: Roseberries*  
*Photography: Davic Kidd*

COVER AND INSIDE COVER PHOTOS:  
The Virginia Tech Innovation Campus  
in Alexandria, site of the 2025 VASEM Summit.



# Charting the Future of AI in Virginia

SEPTEMBER 30 – OCTOBER 1, 2025

## CONTENTS

### 2 Greetings

**Sen. Mark Warner**, U.S. Senator from Virginia

**Alfred Grasso**, President, Virginia Academy of Science, Engineering, and Medicine



### 4 Morning Session, September 30

**Keynote:** “What Is the Future of AI?”

Speaker: Arvind Narayanan, Professor of Computer Science and Director of the Center for Information Technology Policy, Princeton University

**Session One:** “AI Research across Virginia’s Universities.”

Chair: Scott Acton, University of Virginia

**Session Two:** “Building AI Infrastructure in Virginia.”

Chair: Yael Grushka-Cockayne, University of Virginia



### 10 Afternoon Session, September 30

**Keynote:** “Charting the Future: AI, Workforce Innovation, and the Public Policy Landscape.”

Speaker: Michelle Lopes Maldonado, Delegate, 20th District, Virginia House of Delegates

**Session Three:** “A TED Talk Approach to AI across the Commonwealth.”

Chair: Amarda Shehu, George Mason University

**Session Four:** “Making Virginia a National Leader in the AI Industry.”

Chair: Ryan Wright, University of Virginia



### 16 Morning Session, October 1

**Keynote:** “AI and Medicine: How Will Healthcare Change?”

Speaker: Mark Sendak, MD, Vega Health

**Session Five:** “AI Pandemic Modeling and Prevention.”

Chair: T. M. Murali, Virginia Tech

**Session Six:** “AI in Clinical Practice.”

Chair: Bob Phillips, MD, American Board of Family Medicine



### 23 Student Poster Presentations

### 25 Participant Biographies

### 28 VASEM Board of Directors

# GREETINGS



## Senator Mark Warner

Over a decade ago, I convened a meeting of National Academy members from Virginia and presidents of the Commonwealth's premier research institutions. Our purpose was to address the need for an independent body of experts to provide nonpartisan technical insight for state policymakers on complicated policy issues. These conversations led to the formation of the Virginia Academy.

I am proud that VASEM has done so much to advance public policy in Virginia, and these annual summits are an important part of that effort. This summit on artificial intelligence is particularly timely because additional investment in AI is critical to U.S. global technological leadership and, therefore, to our national security.

The summit coincides with bipartisan partnerships I have formed in the Senate to better understand the impact of AI on the workforce, protect children from harmful AI chatbots, enhance the security of the AI ecosystem, and limit the potential of AI-generated misinformation to disrupt elections. These are areas of national concern, but they are also ones in which Virginia has the expertise and infrastructure to take a leadership role.

I look forward to seeing where the conversation and ideas from this year's summit go, and, as always, look forward to continuing to support the efforts of the Virginia Academy.

I am proud that VASEM has done so much to advance public policy in Virginia, and these annual summits are an important part of that effort.

A handwritten signature in blue ink that reads "Mark R. Warner". The signature is fluid and cursive.

Mark R. Warner  
United States Senator

# Alfred Grasso

There have been hundreds of conferences about artificial intelligence over the past decade, but the Virginia Academy took a different approach for this year's summit. Our focus was not exclusively on the technology. Rather, it was on fostering the creation and dissemination of AI in Virginia in ways that promote economic growth and support the well-being of its citizens. This purpose is consistent with the Virginia Academy's mission of convening expertise and fostering collaboration to help policymakers address the most promising opportunities facing the Commonwealth.

For our 11th summit, we looked at pioneering research underway at Virginia's universities to advance artificial intelligence and to use it to expand and accelerate discovery in such fields as healthcare, cybersecurity, and defense. We heard from Del. Michelle Maldonado, who urged lawmakers to come together to promote innovation in AI while establishing guardrails to promote security, privacy, and transparency. And we learned how our leading businesses and healthcare providers are implementing AI today. Behind all these discussions was this underlying question: how can we catapult Virginia to the front rank of states that are leaders in this field?

I would like to thank the members of our organizing committee—Scott Acton, T. M. Murali, Bob Phillips, Naren Ramakrishnan, and Amarda Shehu—as well as our expert panelists and generous sponsors for making this summit such a success. I invite you to peruse this publication, which touches on the summit highlights, and join us at VASEM in helping to set the AI agenda for the Commonwealth.



Behind all these discussions was this underlying question: how can we catapult Virginia to the front rank of states that are leaders in this field?

A handwritten signature in blue ink, appearing to read 'Alfred Grasso', written in a cursive style.

Alfred Grasso  
President, Virginia Academy of  
Science, Engineering, and Medicine

# MORNING SESSION, SEPTEMBER 30

## KEYNOTE

# WHAT IS THE FUTURE OF AI?

Arvind Narayanan, PhD

For some people, the future of artificial intelligence (AI) is the stuff of nightmares. Daniel Kokotajlo and his colleagues at the AI Futures Project describe a doomsday scenario in their recently published manifesto, “AI 2027.” They warn that, in a matter of years, AI will rapidly attain superintelligence and effectively control all key human endeavors. In their view, AI’s capacity for disruption and destruction is on par with nuclear weapons.

Our keynote speaker, **Arvind Narayanan**, and his Princeton colleague Sayash Kapoor take a different tack. They argue in “AI as Normal Technology” that challenges similar to the ones technologies such as electricity faced will slow AI’s deployment and moderate its immediate transformational potential. From their perspective, AI is more accurately compared to nuclear power than nuclear weapons. “Thinking of AI and its future impact as normal technology,” Narayanan said, “is different from the view of AI as a tsunami.”

## PROGRESS AS INCREMENTAL CHANGE

In his keynote, Narayanan built on well-established rubrics for technology dissemination that trace the evolution of AI from invention to innovation to diffusion. Obstacles within and between each stage, he asserted, will give people and institutions time to adapt, effectively spreading its transformational economic and societal impact over decades. As applied to AI, *invention* is the development of new methods and capabilities that improve AI’s ability to carry out various tasks, *innovation* refers to the development of products and applications using AI that consumers and businesses can use, and *diffusion* includes both the processes by which AI products are adopted as well as changes that will allow society to take productive advantage of these innovations.

**“Thinking of AI and its future impact as normal technology is different from the view of AI as a tsunami.”**

Narayanan maintained that many commentators and companies promoting AI products have overestimated the current state of the technology, the effectiveness of resulting AI applications, and the readiness of users to adopt them. For instance, they have failed to appreciate that progress in each stage occurs over different timescales. Typically, the development of new methods and capabilities—invention—proceeds rapidly. Product development is slower and diffusion slower still.

In the first few years after ChatGPT was introduced, Narayanan noted, the AI industry assumed that advances in invention could quickly be transferred to the marketplace and that it could create software products for new domains simply by placing specialized wrappers around large language models. He cited a series of products based on this idea that immediately flopped.

One reason for developers’ misplaced confidence is that they failed to consider construct validity, the extent to which a test or measurement is appropriate for its purpose. Benchmarks that are useful to track invention may be irrelevant for innovation. Narayanan cited the example of legal software. “Hearing that AI aced the bar exam may seem like an endorsement for its use in legal products,” he said, “but a lawyer’s job does not consist of answering bar exam questions.” He and Kapoor have examined the capability benchmarks that many AI companies use and found that they often benchmark tasks that are both easier to measure and for AI to complete than tasks that require creativity and judgment.

In addition to developing benchmarking tasks that reflect relevant domain requirements, Narayanan argued



that AI product developers must focus on five essential criteria for product innovation. These are cost, reliability, privacy, safety and security, and the user interface. None of these matter much for invention. “The way you measure progress in one stage is different from the way you measure it in another,” Narayanan said.

Similarly, adoption and diffusion pose their own unique challenges and occur at differing timescales. Early adopters, for instance, face a steep learning curve before they can use a new AI product productively. They will also need time to adjust their workflows and habits to take advantage of the benefits of new AI products and avoid their risks.

Diffusion also entails deep changes to organizations and business models. Narayanan cited the example of the Second Industrial Revolution, when electricity replaced steam power in factories. Manufacturers found that a one-to-one substitution failed to significantly reduce costs or boost productivity. It was only when they began to understand that electric cable gave them the flexibility to redesign workflows in ways that were impossible with steam pipes that they began to reap its benefits. They created the assembly line.

In addition to altering workflows, diffusion requires changes to institutions, laws, and norms. The need to make changes of this magnitude explains why AI has yet to provide a meaningful boost to productivity. “There is a myth out there that AI is the most rapidly adopted technology in history,” Narayanan said. “Technology may be changing rapidly, but the rate at which we can adapt cannot keep up,” he concluded. “Adoption happens at human speed.”

## APPLYING THE LESSONS

A basic understanding of the dissemination process for emerging technologies has implications for both AI researchers and companies trying to use AI. For instance, Narayanan argued that AI researchers should spend more time exploring new methods rather than exploiting existing ones. He noted that the development of general-purpose AI remains hyperfocused on transformer architecture, which was introduced in 2017. “Researchers focusing on alternative architectures have difficulty demonstrating their value because they don’t perform as well on benchmarks designed for transformer-based models,” he said. “This is a structural problem with how AI research is organized and incentivized.”

Narayanan noted that companies that are serious about adopting AI should be investing more on the human side of the equation. “Companies might have to allow employees to use 20 percent of their time to develop the skills for AI-enabled work—and do this for many years before they realize commensurate time-savings and productivity,” he said.

Narayanan’s application of diffusion theory to AI is part of the broader discussion that he and Kapoor lay out in “AI as Normal Technology”—which includes its effects on labor, approaches to risks it may create, and its implications for policy. All these discussions are predicated on the idea that AI is more like technologies that have emerged in the past and is not a singular creation capable of superseding all previous technologies.

## SESSION ONE

# AI RESEARCH ACROSS VIRGINIA'S UNIVERSITIES

The faculty of Virginia's colleges and universities have been quick to appreciate the potential of AI to accelerate their research as well as to address questions that exceed the reach of previous methodologies. They have also been active in pushing the technology forward, making AI more transparent and more accurate, and many have even reorganized their institutions to facilitate the diffusion of AI across the disciplines.

In the first session of the summit, moderator **Scott Acton**, chair of the Department of Electrical and Computer Engineering at the University of Virginia (UVA), AI advisor to the provost, and one of the organizers of the summit, asked faculty members from a cross-section of Virginia's universities to discuss their work on AI.

## COSMIC AI

We are in the midst of a golden age of astronomical observation. Powerful observatories on Earth and in space have generated vast amounts of data—the Hubble Space Telescope has accounted for more than 150 terabytes and counting—that if properly harnessed have the potential to enable transformative discoveries about the nature and origins of the universe. For UVA Associate Professor **Paul Torrey**, a computational astrophysicist and co-leader of the new AI Institute for Cosmic Origins, this data also provides a testbed for advances in the technology itself.

As Torrey pointed out, astronomical data meets all the prerequisites for tackling critical cosmological and technological challenges. The data is public, plentiful,



nonproprietary, and reasonably homogeneous—a great advantage for AI model-builders—but it is also complex, noisy, and high-dimensional, presenting significant challenges that, if addressed, could benefit researchers in other fields. He also noted that the open-source collaborative coding and publication ecosystem in astronomy makes it easy to pool resources to pursue research and develop and test AI methods.

“With the physics-based models we use now, we can produce simulations that we can compare to our observations,” he said. “With AI inference methods, the institute team can flip this approach on its head and infer the physics from our observations.” At the same time, institute members will be focusing on methods to improve AI trustworthiness, efficiency, interpretability, and robustness.

## GETTING UNDER THE HOOD OF LANGUAGE MODELS

Next to present was Assistant Professor **Ziyu Yao**, who codirects the Natural Language Processing Group at George Mason University with Antonios Anastasopoulos. One of its research thrusts is to shed light on the internal workings of large language models (LLMs), which are commonly perceived as black boxes. Understanding these underlying mechanisms would place users in a better position to evaluate LLM results and maximize their application for specific uses.

“We actually know very little about what happens under the hood of large language models,” Yao said. “Knowing more about how they succeed or fail at tasks provides a foundation for AI safety and trust and suggests directions for improvement.”

As an example, Yao discussed her team’s analysis of a three-operand arithmetic calculation, which most people would solve sequentially. While LLMs such as Llama 3 perform well on such tasks, Yao found that they solve it in a totally different way, completing all operations at once in the last token position. “This process helps explain why LLMs are so fragile and prone to error when tasks are more complicated and scaled up,” she said.

Yao is also the principal investigator for a National Science Foundation- and Microsoft Research-funded program called MathVC (Mathematics Virtual Classroom), which includes collaborators at William & Mary and colleagues at George Mason. She highlighted the team’s efforts to create a virtual AI peer that would help middle school students learn math concepts and practice math problem-solving skills.

## BUILDING MODELS ON BETTER DATA

No matter how sophisticated, an AI model is only as good as its data. In her presentation, Virginia Tech Assistant Professor **Ruoxi Jia** stressed that little is known about the data used to train popular AI models like Llama 3. “The builders of open-weight models are usually fully transparent about model architecture and even training procedures,” she observed, “but they provide basically no information about the data used to train them. We don’t know how they collect the data, where it comes from, or how they chose it.”

While much AI research has been focused on improving models, Jia’s lab concentrates on tackling the informational imbalance, creating tools that enable data users to assess data quality from limited samples that data providers sometimes supply and to understand how data influences model behaviors. The work that Jia and her colleagues are conducting also has implications for evaluating model performance, determining copyright issues, and ensuring AI safety and security.

## CROSS-CAMPUS AI

Over the past decade, a wave of AI-inspired research has swept the campus of William & Mary, accelerating innovation and inquiry in dozens of fields. To better support and coordinate these efforts, the college launched its first new school in 50 years, the School of Computing, Data Sciences & Physics (CDSP). Professor **Douglas C. Schmidt**, the school’s inaugural dean, highlighted some of the AI-work that is transforming research at William & Mary. “One of my priorities as dean is to figure out how we can inculcate and promulgate the use of effective and responsible AI throughout the research, teaching, and cocurricular activities we conduct at our school and beyond,” he said.

Noting that the school itself concentrates on advances in AI technology, Schmidt focused on applications of AI to other fields. Among the research projects he described were efforts to simulate the responses of human decision-makers in complex, high-stakes contexts, assess the health of the Chesapeake Bay and track animal migrations across vast scales, deploy large language models to make and accelerate discoveries in experimental and theoretical physics, and manage gait issues in Parkinson’s disease.

Schmidt also noted that William & Mary has launched an AI minor and plans to introduce a Bachelor of Arts in Applied AI. “In an institution like William & Mary, which prides itself on the liberal arts and social sciences, we are working to ensure that all our students can take advantage of AI,” he said.



## SESSION TWO

# BUILDING AI INFRASTRUCTURE IN VIRGINIA

For a discussion of the infrastructure needed to develop a flourishing AI sector in Virginia, the Virginia Academy assembled a panel of distinguished researchers and corporate executives to provide a range of perspectives. It consisted of the following experts:

- Professor **Lingjia Liu**, director of Wireless@Virginia Tech.
- Professor **Milos Manic**, director of the Virginia Commonwealth University Cybersecurity Center.
- **Nathan Mellis**, vice president of artificial intelligence at Maximus, a provider of government services.
- Professor **Manish Parashar**, chief AI officer at the University of Utah and executive director of its Scientific Computing and Imaging Institute.

Professor **Yael Grushka-Cockayne**, senior associate dean at the University of Virginia's Darden School of Business and special advisor to the provost on AI, moderated the session. She began by asking each panelist to reflect in general terms on the state of AI infrastructure in Virginia.

A specialist in wireless networks, Liu began by pointing out that AI has revolutionized virtually every field and is particularly important for the next-generation wireless networks he is helping to develop. "AI will be an enabling technology for 6G resource allocation, network slicing, and multiuser scheduling," he said. "It's something that is very much part of research we're conducting at Virginia Tech."

Manic turned the discussion to some of Virginia's attributes that set the stage for a flourishing AI sector. Some of these advantages are physical. The many transatlantic cables that make land in Virginia, he noted, have made the

Commonwealth the preeminent location for data centers in the world. Other advantages are institutional. Manic made a point of emphasizing Virginia's organizational infrastructure, including the Commonwealth Cyber Initiative (CCI) and the Commonwealth Center for Advanced Computing, as examples of Virginia's ability to coordinate resources and develop a foundation for technological innovation.

From a business point of view, Mellis observed, Virginia's network of research universities is a critical part of its AI infrastructure. "We keep a close eye on what's happening in academia," he said. "We are interested in determining how we can apply breakthroughs there to real world use cases." And while he conceded that when people think of advances in AI they think in terms of GPUs and servers, Maximus takes a more holistic view. "The real infrastructure for AI is the proper application of people and process as well as technology," he said. "At this point, Virginia is doing about as well in these three areas as any state."

Parashar echoed Mellis' observations. He drew on his experience working with legislators and community members in Utah to launch the One-U Responsible AI Initiative, a \$100 million university investment that leverages AI to solve such regional and statewide challenges as air quality and mental health. "To be successful, we need to attract talent, build infrastructure and support structures, and engage the community," he said.



## LEVERAGING EXISTING STRENGTHS. BUILDING NEW ONES

Grushka-Cockayne then asked the panelists to take a deeper dive into Virginia’s advantages as a center for AI development, as well as the challenges it will face translating these advantages into leadership.

Liu provided a detailed list of strengths, asserting that “Virginia has many of the prerequisites necessary for AI to be deployed successfully.” In addition to Virginia’s proximity to Washington and the federal government and its position as a global data center, Liu noted the space and defense facilities in the Hampton Roads area, the state’s excellent transportation system, its talented workforce, and an established ecosystem of collaboration as epitomized by the CCI. “All these characteristics make Virginia good if not better than other states in readiness for AI and a leader in a number of niches,” he said.

Parashar emphasized Liu’s point about cooperation. In Utah, he said, industry, academia, and state government, working together, established a state Office of Artificial Intelligence Policy in record time, going from initial conversations about AI regulation to hiring an executive director in about a year. “The give and take that characterized the creation of this office is extending to the regulations it creates,” he said. “This has been a win-win for everyone.”

## AI IS NOT ALWAYS THE RIGHT ANSWER

As AI grows in capability, the pressure to move applications of AI into the marketplace will only increase. For Grushka-Cockayne, this raises an interesting question: how do you decide when a new and emerging technology is ready for primetime? How do you know when to release it?

There is the technical side. Liu, who worked for Samsung before joining Virginia Tech, noted several practical questions that must be addressed. Is it reliable and robust? Will it work under real-life conditions and produce

as few errors as possible? Is it interoperable with existing systems? And do the benefits it provides outweigh the costs of implementing it? This is particularly an issue with AI, which is especially power hungry.

Mellis joined Liu in maintaining that even when AI products perform flawlessly, they might not be the best solution for a particular customer or application. Mellis noted that many of Maximus’ federal clients are interested in bringing AI into their agencies or programs. The company sometimes finds, however, that after looking at the issue they are trying to solve, an established technology like a rules engine or a basic linear regression model would serve them better. “The first step in evaluating a new technology is determining whether it solves your problem,” Mellis said. “The second is determining if it represents a better solution than existing technologies.”

## INTO THE UNKNOWN

We do not know what the future of AI has in store for society. When Grushka-Cockayne asked the panelists what about AI keeps them up at night, uncertainty about its consequences was a common thread in their answers. As an expert in cybersecurity, Manic is professionally suspicious. “When you show cybersecurity specialists a new capability, the first thing they think about is what can go wrong,” he said. The question of intent figures prominently in his thinking. “AI is a powerful tool, and like any powerful tool, it can be used for noble or nefarious purposes,” he said.

Parashar added that AI is still in its formative stages—and it has the potential to evolve in ways that are both beneficial and detrimental to society. Building a broad, informed community, he said, is the best way to shape the desired outcome. “Our challenge is to build partnerships that can help steer its development and application,” he added.



## AFTERNOON SESSION, SEPTEMBER 30

### KEYNOTE

# CHARTING THE FUTURE: AI, WORKFORCE INNOVATION, AND THE PUBLIC POLICY LANDSCAPE

Del. Michelle Maldonado

**D**elégate **Michelle Maldonado** has come by her tech credentials honestly. A self-described “recovering tech lawyer,” she joined America Online as a legal counsel in 2000, shortly after the dot.com bubble burst, and stayed with the company through its rebranding as AOL later in the decade. Although she has subsequently been devoted to leadership transformation, launching her firm, Lucenscia, 10 years ago, she has nonetheless emerged as the foremost authority in the General Assembly on the impacts of AI and other emerging

technologies. She founded and chairs the General Assembly’s bipartisan and bicameral Technology & Innovation Caucus and is a member of the Joint Commission on Science and Technology and its AI Subcommittee.

Maldonado is enthusiastic about the potential of AI to improve people’s lives. She cited two examples of healthcare applications to illustrate its power. One was the use of AI to assist radiologists in interpreting mammograms. While it is not yet standard treatment, numerous hospitals and health systems are offering AI-powered software as an option that

may identify small early-stage cancers while reducing the false positives that lead to unnecessary follow-up tests and biopsies.

She also mentioned the use of AI to replicate voices of people who have illnesses that will eventually deprive them of their ability to speak. A few hours of recording can provide sufficient data for generative AI to replicate their voices, complete with characteristic tone, cadence, and inflection. When they lose their voice, patients simply type into their device what they want to say—and their unique voice is resurrected. “It’s an amazing way for people to retain their personality and presence,” she said.

But Maldonado can imagine situations where criminals could use the same AI system to steal a person’s voice and use it for their own ends. The potential of using AI for harmful purposes is one reason statistics show that more people are anxious than excited about AI—and why more than 50% believe it should be regulated in some way.

## THE AI WORKFORCE

One reason for this unease, Maldonado noted, is that the pace of technological change has accelerated sharply since the first Industrial Revolution. As a result, the urgency of managing the disruption caused by change has increased commensurately. Among issues that have received most attention is the impact of AI on the workforce.

Maldonado echoed other speakers in believing that, over the long term, there will be more jobs created by AI than lost. At the same time, she insisted that it would be disingenuous to maintain that AI will not replace some positions. She cited the example of the automotive industry. When robotics were introduced, workers were displaced. “We must look at the kind of jobs and communities that are affected,” she said. “The overall picture is exciting, but we are called to look beneath the numbers.”

It is therefore imperative, she maintained, that we develop programs to upskill and reskill displaced workers, seeing the challenges as an opportunity for augmentation rather than replacement. “If you are well-equipped and know how to use AI tools, you’re going to make yourself more productive and your organization more competitive,” she said. Fortunately, she pointed out, Virginia has an extensive network of colleges and universities ready to meet this challenge.

## BALANCING INNOVATION AND PROTECTION

A second challenge raised by the AI revolution is what Maldonado called the “natural tension” between promoting innovation and economic development and ensuring

privacy, security, safety, and transparency. Related to this issue are limits on the role of AI. “There is a growing conversation about the line between decisions that are acceptable for AI to make and those that are not,” she said. “Making this distinction is not as simple as deciding to keep AI away from the nuclear button. As the technology becomes more sophisticated and agentic becomes more common, we need to determine where AI intersects with accountability.”

Because data centers provide the infrastructure for AI, making decisions about the proper boundaries of AI also requires decisions about how and where we want data centers to be placed. Right now, Northern Virginia is the data center capital of the world. Can we accommodate the exponential leap in supply that AI requires? “We have to think about data center placement and proliferation as well as their impact on the energy grid and water supply,” she said. And while data center operators are adopting technologies such as closed loop water systems that limit the drain on our resources, Maldonado insisted that a discussion about proliferation must still take place, one that includes the everyday users whose devices and lifestyle, collectively, contribute to a growing demand for data centers.

## THE IMPORTANCE OF COLLABORATION AND PARTNERSHIP

Given the pace of AI development, it is imperative, Maldonado maintained, to find ways to resolve these issues sooner rather than later. The only way these priorities can be balanced over the long term, she asserted, is through collaboration and partnership. “Our job when we differ is to lean into these conversations, not step away,” she said. “We have to bring together people who have different approaches to figure out how to move forward.”

This will also require coordination on the federal and state levels. The 2025 Federal AI Action Plan has three pillars: accelerating AI innovation, building AI infrastructure, and leading in AI diplomacy and security by expanding American exports and countering foreign influence. “This is a start,” Maldonado said, “but we need the federal government to emphasize issues like privacy and protection.” State governments also have a role to play.

Maldonado explained that we are not faced with either/or decisions about the future of AI. The options, she stressed, are not choosing either to lean in fully on innovation and competitiveness or to focus on issues like data privacy and transparency. “These are not separate conversations,” she said. “They must be happening at the same time.”



## SESSION THREE

# A TED TALK APPROACH TO AI ACROSS THE COMMONWEALTH

For our first afternoon session, moderator **Amarda Shehu**, George Mason University's (GMU's) inaugural vice president and chief AI officer and one of the summit's organizers, chose to break things up by asking her panelists to deliver short TED-style talks on AI. TED talks are designed to be entertaining as well as informative, and our speakers more than amply rose to the challenge.

## NOT READY FOR PRIMETIME

The first speaker was GMU Professor **Missy Cummings**, director of the Mason Autonomy and Robotics Center. As Shehu pointed out, her colleague is also a former U.S. Navy fighter pilot. It is thanks to her naval career that Cummings gained her first exposure to what she referred to as “baby forms of AI”: “Working with Predator drones, I spent a ton of time looking at what could go wrong when putting AI into safety-critical systems.”

This skepticism informs her assessment of AI. “There is a lot of fear pushed by the media, the tech companies, and academicians about the power of AI,” she said. “Fear drives dollars.” Cummings is committed to debunking myths that underly this fear. She asserted that AI is not as capable as some want the public to believe: “AI doesn’t know anything. It doesn’t understand anything. It is simply linear algebra on steroids.” She further noted that AI doesn’t reason especially well under uncertainty, is not proficient at generating code, and, while it can do a better job than people at recognizing

patterns, it cannot make sense of some images—such as a stop sign partially blocked by leaves—that are intuitively obvious to people. “Right now, AI needs human babysitting,” she concluded.

## FOR COLLABORATION, NOT COMPETITION

For **Luiz DaSilva**, the Bradley Professor of Cybersecurity at Virginia Tech and executive director of the Commonwealth Cyber Initiative (CCI), the kind of partnership that CCI embodies will drive innovation. “We have 47 institutions of higher education rowing in the same direction rather than competing,” he said. “This is the reason for CCI’s success since it was founded five years ago—and why it has already spawned many AI-inspired research projects and start-ups.”

CCI’s view—there is no cybersecurity without AI and no AI without cybersecurity—explains why AI represents a significant portion of its portfolio. AI can play a critical role in making cyber-systems less vulnerable to attack. And having a firm cybersecurity foundation is critical if AI is to be widely trusted. “The two are tightly connected,” he said.

Each year, CCI provides seed grants for promising research undertaken by teams of researchers from multiple institutions. In April, it awarded \$1.6 million for 18 projects, each of which approached AI and cybersecurity from a different perspective. Many of the 12 spinouts that have benefitted from CCI support have also involved novel uses of AI. For instance, WiSights Lab is pioneering AI and LLM solutions for advanced telecom networks.



These efforts have already made their mark on the Commonwealth's economy. CCI researchers have received an estimated \$384 million in external funding since the program started, and CCI efforts have generated 4,600 new jobs and \$370 million in labor income and added \$526 million to Virginia's GDP.

## NUISANCE FLOODING

As sea levels rise, coastal communities have increasingly been subjected to sunny day flooding during high tides. The problem is particularly acute in Norfolk, home to Old Dominion University, because the ground is subsiding. During these events, many streets become impassable. "Our goal," said Professor **Khan Iftekharuddin**, director of the Old Dominion University Vision Laboratory and inaugural director of its Institute of Data Science, "is to automatically assess this flooding and provide guidance to drivers, so that they can avoid flooded areas."

Although traffic-flow cameras provide information about the extent of flooding, they do not reveal any information about its depth. Iftekharuddin and his colleagues realized that they could apply machine learning to estimate depth by using such reference points as automobile tires.

To implement this idea, they had to overcome several obstacles in training their model. The first was the dearth of labeled flooding videos. In response, Iftekharuddin's team developed synthetic videos of flooding at different water depths. "To the best of our knowledge, this is the first time that anyone has generated synthetic flooding videos with physically relevant features such as moving ripples and small waves when cars move on flooded streets," he said. Another was the difficulty interpreting the size of vehicles in these datasets because they were recorded at differing scales and perspectives. The team developed a method

to normalize these images so they could be analyzed to develop learning models.

When they tested their machine learning model on a dataset of real-world images, it estimated depth with 75% accuracy. Iftekharuddin plans to build on this success by customizing the model for specific locations like Norfolk by incorporating weather data and local infrastructure.

## CREATIVITY AND AI

"The promise of generative AI is that it doesn't just recognize the world," said Assistant Professor **Pinar Yanardag-Delul**, leader of Virginia Tech's Generative Modeling Lab (GEMLAB), "it imagines new ones—and it places that power in the hands of anyone. We've gone from AI that says, 'This is a cat,' to AI that invents entire worlds filled with cats that never existed."

Yanardag-Delul has long been fascinated by using AI as a tool to foster creativity. Before joining Virginia Tech, she was CEO of AI Fiction, a creative design studio, and cofounder of GLITCH, the world's first generative AI clothing line.

The dilemma for creators—as well as anyone who uses generative AI—is that these models are so complex that it is virtually impossible to understand how they produce their output and whether their results reflect built-in biases. Yanardag-Delul noted that ChatGPT 4 is often estimated to have around 2 trillion parameters, and an image-generation model like Stable Diffusion 3.5 has more than 8 billion parameters. "I started GEMLAB with the hope of making these giant models more explainable and transparent and to give humans the ability to control them," she said.

As an example of her research, she highlighted her group's discovery of semantic dials within pretrained models like Stable Diffusion. "These are the building blocks of AI imagination," she said. "Once we uncover them, we can start steering the kind of content it produces."

## SESSION FOUR

# MAKING VIRGINIA A NATIONAL LEADER IN THE AI INDUSTRY

**P**rofessor **Ryan Wright**, currently the co-director of AI Research@UVA and moderator of the fourth session, had been tracking the day's discussions. "By my count, we've mentioned fourteen different research projects and six AI initiatives thus far in the summit and made 18 references to data centers," he said.

Wright went on to cite several other factors that position Virginia for leadership in AI. These include the 11-story, 300,000-square-foot academic building on the Virginia Tech Innovation Campus where the summit was being held, the Commonwealth's persistently high performance in CNBC's Top State for Business ranking, and the presence in Virginia of companies at the forefront of AI, including Amazon, Boeing, Google, Lockheed Martin, and MITRE.

The last two had representatives on the panel:

**Mark Maybury**, vice president of commercialization, technology, and strategic innovation at Lockheed Martin, and **Charles Clancy**, senior vice president, chief technology officer, and head of MITRE Labs. They were joined by **Mehul Sanghani**, founder and former CEO of Octo, which provides emerging technology and IT modernization services, including artificial intelligence, to the federal government. He and his wife, Hema, helped fund the Sanghani Center for Artificial Intelligence and Data Analytics, headquartered at the Innovation Campus.

Rather than focus on the past, Wright explained that the session was about the future: "We are here to discuss steps Virginia needs to take if it is to emerge as the AI Commonwealth," he said.

## A SAMPLING OF AI-FORWARD INITIATIVES

Wright kicked off the conversation by asking the panelists to discuss AI-related initiatives their organizations have on the table. Clancy discussed MITRE's efforts to support the federal government's campaign to incorporate AI into its systems. As part of this initiative, it created the AI Assurance and Discovery Lab, a configurable space it is using to develop methods to identify, assess, and manage risk in AI-enabled systems. "Right now, we are very focused on helping the U.S. government on its AI journey and accelerate adoption, but we're doing so in a way that has assurance baked in ahead of time," Clancy said. Partnering with NVIDIA, MITRE has also created a \$30 million supercomputer for its Federal AI Sandbox that will enable federal agencies to experiment with prototypes and advanced AI deployments.

Maybury highlighted the launching of Lockheed Martin's new Astris AI subsidiary, which makes the AI infrastructure that Lockheed Martin developed for internal use—the LM AI Factory—available to commercial customers. Maybury recalled that back in 2018, Lockheed Martin's leadership team, realizing that AI was becoming a significant part of the weapons systems it was developing, decided to centralize AI across the company under the banner of the Lockheed Martin AI Center. The company subsequently added machine learning operations and generative AI software platforms, with security and compliance required for entities working in highly regulated and classified environments. "Through Astris AI, we are making these capabilities available to companies seeking to develop reliable, high-assurance AI solutions while protecting their intellectual property," he said.

Sanghani talked about his decision to endow the Sanghani Center and his support for the Innovation Campus. "When I was considering endowing the center, I did so with the idea of making a real difference and having a tangible business impact," he said. "The support we've received from Boeing and Northrop Grumman speaks to that vision." The goal of the Innovation Campus is to produce leaders who focus not on demos but on building and shipping actual products. "That's a differentiator for the state," he said.

## PUTTING UP POINTS

Wright then asked the panelists what they thought it would take to foster AI leadership in Virginia. Sanghani referred to the challenge as a scoreboard issue. Virginia already has a unique bundle of advantages that is helping it put points



on the board. The first is proximity to Washington. “The federal government is the largest funder of R&D and the largest services buyer in the world—and they’re right on our doorstep,” he said.

Sanghani also pointed out that the major aerospace, intelligence, and defense firms all have a significant presence in Virginia and that the Commonwealth is home to some of the world’s leading research institutions. And finally, there is Virginia’s data center infrastructure. “Thirteen percent of the world’s data center capacity runs through two counties in Northern Virginia,” he said. Total this all up, and Virginia is ahead of most state players.

Clancy believes that Virginia could score even more points by providing additional support for start-ups. He noted that over the past decade, the federal government has gone from building its own bespoke technology to increasingly relying on commercial companies for everything it does and that the current administration is leaning heavily on the world of dual-use defense tech innovation. “I think we could be a gateway to innovation if we could develop an investor ecosystem to drive start-ups here,” he said. Clancy pointed out that a similar conversation about cybersecurity 10 years ago led to the creation of the Commonwealth Cyber Initiative (CCI). An organization based on this model could provide funding for start-ups based on Virginia’s AI research, he claimed.

Maybury also underscored the importance of investing in start-ups; he pointed out that Lockheed Martin has provided \$400 million in venture capital to start-ups, many of which leverage AI. And he highlighted the importance of efforts like CCI that foster collaboration. But he also noted two other elements essential to making Virginia a business leader in AI: access to research and talent. “The connection among research, talent, and investment is extremely tight,” he observed. “If you remove one of the pieces, the puzzle falls apart.”

Sanghani agreed, emphasizing the importance of developing a highly qualified AI workforce. “Talent is the be-all and end-all,” he said, and the need for talent is large and growing. Sanghani cited a statement by Virginia’s secretary of labor that there were 31,000 statewide AI job postings in July; he used the Innovation Campus as an example of the state’s extraordinary efforts to meet employer requirements. “The Innovation Campus was designed to produce 400 to 500 master’s students and PhDs in engineering and computer science each year,” he said. “The average university program produces about one fourth of that. That difference provides an idea of the scale of the investment we are making.”

# MORNING SESSION, OCTOBER 1

## KEYNOTE

# AI AND MEDICINE: HOW WILL HEALTHCARE CHANGE?

Mark Sendak, MD

**O**n the face of it, innovation in healthcare may seem like an oxymoron. Hospitals and health systems are inherently conservative, prioritizing the mastery of best practices over the development of new ones. And many healthcare organizations are siloed, hierarchical, and under-resourced, all of which makes it difficult for them to innovate.

But as **Mark Sendak** stressed in the course of delivering the summit's healthcare keynote, AI and machine learning have the potential to help clinicians deliver a significantly higher standard of care while reducing the administrative burden that takes up so much of their time. "Not only does AI hold the promise of improving patient care, but it can also make healthcare organizations better places to work," he said.

For his keynote, Sendak drew on the lessons he learned during his tenure as the population health and data science lead at the Duke Institute for Health Innovation (DIHI) and as a member of the leadership team of the Health AI Partnership (HAIP). He focused on the institutional barriers to technology development and dissemination found in even the most forward-looking health systems and the challenges of transferring that technology to other healthcare providers. He also spoke about the issues that led him to found Vega Health, which brings validated AI solutions to healthcare organizations to help them achieve their strategic goals.

## AN ENGINE FOR HEALTHCARE INNOVATION

DIHI was created as the in-house innovation engine for Duke Health. Its goal is to stimulate and implement transformative health and healthcare innovations that

**“Our role is to clear away the obstacles and smooth the way for the organization to move forward.”**

improve clinical outcomes, lower costs, increase access, enhance provider well-being, and advance population health. It does so by providing the technical expertise and resources needed to develop solutions that address problems proposed by frontline staff and to scale these solutions across Duke Health. DIHI also is dedicated to cultivating a community of entrepreneurial thinkers. Since it was launched in 2012, the institute has delivered over 70 projects including the groundbreaking Sepsis Watch™, which combines a deep learning model with electronic health record data to predict this potentially deadly infection five critical hours before patients are typically diagnosed.

DIHI realized early on that organizational barriers to innovation are as formidable as technical ones. Sendak devoted a significant portion of his presentation to detailing the lessons it learned. The first is for healthcare providers to make sure they are working on the right problem. Sendak describes DIHI's role as laying tracks ahead of the slow-moving, high-inertia train that is healthcare. "The direction of that train is already set—and it's not our job to try to change that," he said. "In fact, we are not on the train at all. Our role is to clear away the obstacles and smooth the way for the organization to move forward."

DIHI also safeguards its resources by maintaining a well-defined focus, working only on transformational projects. Sendak observed that most large health systems have teams



devoted to LEAN, Six Sigma, and other waste reduction and quality improvement processes. “If the goal of a project is to optimize existing processes for current patients, there are other resources clinicians can tap,” Sendak said. “DIHI concentrates exclusively on projects that introduce new products or service lines to new patient populations.”

A crucial lesson that DIHI learned is that it is essential to align the individual levels in the organizational hierarchy to remove diversions that can sidetrack a project. One way that DIHI does this is through its annual request for applications. The C-suite identifies strategic priorities for the funding cycle. Business unit leaders review applications that meet these requirements with frontline workers and evaluate them before sending them back up to the C-suite for final selection. Once they are approved, DIHI works with business unit leaders and frontline workers to scope out the problem, design a solution, and conduct the initial implementation. “To successfully sustain an innovation within an organization, you have to break down information and communication silos and secure buy-in across the organization, especially at the business unit level,” Sendak says. “Business unit leaders are often the ones holding the purse strings.”

## SPREADING THE WORD

During 2021, Sendak was part of a team advising the Gordon and Betty Moore Foundation on its investments in AI and healthcare. Over the course of their meetings, representatives from the foundation became intrigued about

what DIHI had accomplished and challenged it to scale its capabilities to other health systems. With funding from the foundation, Duke launched HAIIP. HAIIP is mobilizing expertise from Duke and other large health systems to empower healthcare organizations around the country to use AI safely, effectively, and ethically.

Introducing AI solutions to other groups requires a different model of change than the one DIHI had adopted at Duke. Sendak likened the difference to the distinction between an island and a bridge. “At Duke, we’re working on an island. We know what works and what doesn’t work,” he said. “It takes a different mindset to build a bridge to another organization.”

HAIIP takes a pragmatic approach. Rather than focus on premarket controls—requirements describing what AI vendors must do to participate in the healthcare market—HAIIP concentrates on postmarket supports, on assisting healthcare delivery organizations to assess and make the best use of the growing array of available AI products.

HAIIP soon realized that it was not enough just to distill best practices and publish free resources online. It had to take a hands-on approach. In response, HAIIP launched a Practice Network program that includes peer learning and technical assistance. The program connects AI experts from well-established systems with leaders and clinicians from smaller institutions. The inaugural cohort of the Practice Network program concluded in September 2025.

## CURATING TECHNOLOGY

Sendak was inspired to create his own company, Vega Health, to provide a parallel service for AI technology. “Most healthcare organizations interested in AI don’t have the internal capacity to evaluate, implement, and monitor products,” he said. “At the same time, there are a host of great technologies developed at the larger healthcare organizations that are not diffusing.”

Sendak cited the example of the Sepsis Watch, which took 10 years to achieve market penetration of just two sites despite its superiority over widely adopted solutions offered by electronic health record providers. “At Vega Health, our goal is to aggregate the best validated AI products coming out of healthcare and place them in the hands of community hospitals, while helping them develop expertise to build basic systems of their own,” he said.

## SESSION FIVE

# AI PANDEMIC MODELING AND PREVENTION

The world discovered five years ago just how difficult it is to defeat a pandemic-causing pathogen once it makes the transition from an animal host to human beings. Assessing the danger of potential pathogens and preventing their spread, however, is an enormous challenge fraught with uncertainty, as Virginia Tech Professor **T. M. Murali**, one of the organizers of the summit and moderator of its fifth session, well knows. Murali directs the NSF-funded COMPASS Center, which is exploring the use of predictive models and computational approaches to understand and prevent the emergence and spread of pandemics.

Murali used his vantage point to set the stage for this session on pandemic modeling and prevention. “Imagine a situation where we discover an unknown pathogen infecting an animal that lives in close proximity to human beings,” he proposed. “This situation immediately raises a number of questions for which there are no easy answers.” Is this pathogen capable of jumping to human beings? What changes would the pathogen need to undergo to enable it to infect new species? How might the pathogen replicate inside a host? What are the factors that allow a pathogen to persist in the environment? “These are some of the unknowns that inspired us to form the COMPASS Center,” he said. “Our hope is to move towards a world where we have the tools to foresee pandemics and act proactively to minimize their impact.”

The summit organizers chose the session’s panelists, Murali said, for the light they can shed on the challenges addressing these different questions and the progress that has been made so far.

## IS IT OR ISN'T IT?

Determining whether an organism is a pathogen is not as easy as it might seem at first glance. As Professor **Mihai Pop**, co-director of the University of Maryland Center of Excellence in Microbiome Sciences, noted, there are three well-established methods, each with its own advantages and disadvantages. Researchers can culture the organism in a Petri dish, use the polymerase chain reaction technique to determine if it contains DNA sequences that are characteristic of known pathogens, or sequence its entire genome. Pop focuses on sequencing because he finds the process ideal for working with unknown pathogens. “You can start developing an understanding of how the organism works even as you are trying to determine what you are looking at,” he said.

Machine learning is the most efficient way to examine these sequences to ascertain if they are pathogenic. But as Pop went on to explain, there are pitfalls in using the technique. He cited the example of an unknown bacterium discovered in the air filter of the International Space Station to emphasize his view that stretches of DNA sequences must be assembled into a whole before applying machine learning. “It turned out to be nearly identical to the 5 million-letter anthrax genome,” he said. “If you compare the individual sequences of this unknown organism to their counterparts in anthrax, you will conclude, based on the preponderance of evidence, that you are looking at anthrax,” he said. “It is only when the sequences are assembled that it becomes clear that the organism in question is not anthrax.”

## JUMPING ACROSS SPECIES

Murali’s colleague at Virginia Tech, Associate Professor **Anuj Karpatne**, is a member of the COMPASS team who focuses on understanding zoonotic jumps, mutations in viruses that enable them to move from one host to another. His goal is to look at the sequence of amino acids in a protein from an unknown virus and determine whether it can infect a particular host. To do this, he is applying large language models (LLMs), which use a sequence of words or code elements to produce an outcome. “By using a similar format for protein sequences, we want to see if LLMs can be leveraged to understand the language of viral sequences for host prediction,” he said.

Before this goal can be achieved, difficulties must be overcome. A significant issue is the amount of skew in relevant data sets. More than 60 percent of the protein sequences in databases of hosts are human, while there



are 1,000 species that each account for less than 1 percent of database entries. This imbalance makes it very difficult to predict if one of these hosts is vulnerable to a specific pathogen. In addition, foundation models must be able to work with extremely long protein sequences, which can contain as many as 13,000 amino acids.

Karpatne and Murali addressed these problems by developing HAVEN, which Karpatne described at the first foundation model for host prediction using viral protein sequences that can generalize over unseen viruses and hosts. “We have demonstrated that HAVEN is comparable with state-of-the-art protein language models despite having much fewer internal parameters,” he said. “We also showed that it was useful for host prediction in highly skewed databases.”

## BIG PICTURE MODELING

Professor **Madhav Marathe**, executive director of the University of Virginia’s Biocomplexity Institute, has been part of the support team for every pandemic in the past 20 years, a job that has become increasingly demanding. “Early on, the government approach was limited to developing plans that would be followed in the event of a pandemic,” he said, adding, “As the COVID-19 pandemic demonstrated, we now expect a real-time operational response that gives decision-makers the information they need when they need it.”

But speed is not the only criterion. Data scientists like Marathe must have the capacity to address new, often unexpected questions as a pandemic evolves. “You need

to develop models to synthesize the data that is available, provide a range of interpretations, evaluate responses and actions, monitor intervention responses, coordinate understanding among diverse stakeholders, and make it usable to analysts who are not computing experts,” he said.

Marathe’s affinity for large-scale biological, information, social, and technical systems has allowed him to play an important role at the federal and state levels. During the COVID-19 pandemic, his team supported the Center for Disease Control’s Scenario Modeling and Forecast Hubs. It was also the lead analytical group for the Commonwealth, producing more than 250 weekly updates between 2020 and 2024 and generating more 3,000 slides. The team also supported the Defense Threat Reduction Agency as well as a number of international agencies, including European CDC and state governments in India.

The foundation for Marathe’s contribution is his ability to generate live digital similars that could incorporate new datasets as the pandemic evolved, contextualizing the digital similar with ground conditions. For instance, during the pandemic, he and his colleagues built a digital similar of the U.S social contact network. Now that the pandemic has subsided, Marathe and his colleagues are preparing for the next pandemic by creating a livestock digital similar that can be used to track multihost pandemic spread.

## SESSION SIX

# AI IN CLINICAL PRACTICE

**A**I has the potential to transform healthcare, improving patient outcomes through earlier diagnosis and personalized treatments, increasing operational efficiency by reducing administrative burden, and accelerating medical research and drug discovery. In this session, moderator **Robert Phillips**, executive director of the Center for Professionalism and Value at the American Board of Family Medicine, asked his panel to describe some of the practical issues that health systems face as they deploy AI.

## GETTING THE DATA RIGHT

The first panelist, **Nathaniel Hendrix**, is a researcher and data scientist with the American Board of Family Medicine. Hendrix focused on the dearth of high-quality data needed to train AI models for use in healthcare. “We all understand that AI is only as good as the data we put into it,” he noted.

The problem, Hendrix pointed out, is that data used to train clinical AI models—for instance, to diagnose a disease—is invariably collected for some other purpose. “It might be collected for billing, clinical decision-making, liability protection, or other purposes,” he said. “That mismatch diminishes our faith in the results.”

Focusing just on diagnostic data, Hendrix underscored several sources of inconsistency. They include diagnostic error, cognitive biases, and degree of diagnostic skill. Hendrix also pointed out that diagnostic requirements change over time. “In the early days of the COVID-19 pandemic, doctors diagnosed the disease only after confirmation by a PCR test. Later, an at-home test or a combination of symptoms and known exposure was sufficient for the same diagnosis.”

There are also differences in documentation. For instance, he noted that when he and colleagues from Stanford reviewed Paxlovid prescriptions submitted for



reimbursement, they found that about a third were not accompanied by a COVID-19 diagnosis—even though Paxlovid is a COVID-specific antiviral. Furthermore, Asian, Black, and older patients were more likely to have a missing diagnosis. “This is a good example of the messiness of clinical data,” he said. “AI trained on such data may not support reliable clinical decisions.”

Hendrix’ solution is to start with “good enough AI.” “It might smooth away some of the rough edges in the data collection process and eliminate a little bit of the unjustified heterogeneity we see,” he said. “Once our data becomes clearer, we can build a better system.”

## TAKING A DELIBERATE APPROACH TO AI ADOPTION

The second panelist, **Jon McManus**, is the chief data and AI officer at Inova, a nonprofit health system in Northern Virginia. McManus discussed the adoption of AI from an organizational point of view. He began by noting that Inova



has deployed close to 70 different AI-based systems over the past few years and that adopting each one has required substantial time and effort.

This expenditure is noteworthy, he maintained, because Inova anticipates doubling its AI deployment several times in the near future. “How are we going to scale our capacity to implement safe and responsible AI-based tools given the sheer diversity and breadth of features that this rate of adoption entails?” he asked.

McManus enumerated the elements that he believes any business or enterprise needs to have in place to incorporate AI wisely. The first is having a charter that delineates the organization’s position on AI. For hospitals, this will include expectations about consent and transparency. It is also important to have an inventory of the AI-based products they use, as well as their relationship to these products. “At Inova, we divide the AI we use into three buckets,” McManus said: “The products we take, the products we shape, and the products we make.” The first group accounts for 10 percent of Inova’s portfolio, the second 80 percent, and the last 10 percent. Among other advantages,

understanding Inova’s relationship to its AI products helps it master its risk and think about the talent and skills it needs. “Because 90 percent of the AI we source at Inova is through partnerships, we don’t need a lot of data scientists,” he said.

To successfully scale AI deployment, McManus also urges organizations to have in place what he calls “powering infrastructure,” tools like universal orchestration or universal monitoring that will enable them to manage a diverse portfolio of AI capabilities and provide continuity as they move forward.



Finally, McManus stressed the importance of investing in workforce education and AI literacy. “When it comes to adopting AI, people should not feel that they are on an island making a dangerous decision,” he said. “They need to understand what’s in front them, develop a basic understanding of how it works, and have the support to be successful.”

## PRIORITIZING AI PRODUCTS THAT MAKE A MEASURABLE DIFFERENCE

**Andrew Miner**, Inova’s chief health informatics officer, offered a clinical perspective. Miner began by describing how lucky he feels: “I’m in an incredible position in an incredible field during an amazing time,” he said. When he started training, he recalled, clinicians wrote their notes and orders on paper, and he remembers the skepticism and frustration that greeted the first electronic health record system (EHR). “That’s all behind us,” Miner said. “At this point, we are all connected to the technology. It is an integral part of what we do.”

To reinforce his point, Miner noted that the reception of AI has been generally positive. “I was at a conference of chief medical officers recently where the phrase *love-letter metric* was used,” he said. “The idea is that the success of the technology we introduce can be measured by user messages telling us how much they appreciated it.” By this indicator, the ambient AI program that Inova introduced that summarizes patient-physician conversations and generates notes has been extraordinarily effective. Although providers still review and edit them, it is saving them about two hours a week.

Miner pointed out that we’re just seeing the beginning of the AI revolution. There are many areas where AI would represent a significant improvement over existing products. He mentioned that the state-of-the-art software program that checks pharmacy orders for allergic reactions and harmful interactions has a 90 percent override rate. “AI could make a big difference,” he said. “I’m really excited about products that will help us make better clinical decisions and practice medicine more effectively.”



# STUDENT POSTER SESSIONS

*The future of AI innovation was on ample display at our poster session. Students from many of Virginia's best universities and high schools showcased their research.*

## **Navya Annapareddy**

*University of Virginia*

Accountable Labeling and Provenance Frameworks for Generative AI

## **Elisha Barnes**

*Virginia State University*

AI-Driven Grid Management: Deep Reinforcement Learning for Mitigating Reverse Power Flow in Distributed Photovoltaic Networks

## **Tugce Burcu**

## **Julia Hsin-Ping Hsu**

*George Mason University*

An AI-Based Framework for Understanding Occupational Injuries across Virginia

## **Victor Cobilean**

## **Swagat Das**

*Virginia Commonwealth University*

AI in Cybersecurity of Critical Infrastructure

## **Sanmay Das**

## **Shafkat Farabi**

## **Gaurab Pokharel**

*Virginia Tech*

Street-Level AI: Are Large Language Models Ready for Real-World Judgements

## **Sanmay Das**

## **Gaurab Pokharel**

*Virginia Tech*

Discretionary Trees: Understanding Street-Level Bureaucracy via Machine Learning

## **Ramashish Gaurav**

## **Daniel Rosen**

*Virginia Tech*

Neuromorphic Computing with Spiking Neural Networks

## **Soumee Guha**

## **Yi Huang**

*University of Virginia*

PSF-SRDN: Point Spread Function-Aware Speckle Reducing Diffusion Network

## **Heesang Han**

## **Sandesh Jain**

*Virginia Tech*

Sparse-FS3D and ADRIT for Reliable Few-Shot 3D Perception and Tracking in Outdoor Scenes

## **Chongyu He**

## **Nitya Khamar**

## **Jonathan McGee**

*University of Virginia*

Multi-Modal Speaker Classification in Educational Settings

## **Chongyu He**

## **Jonathan McGee**

*University of Virginia*

Semantic Box: Bounding Box-Guided Caption Enhanced Action Recognition for Instructional Videos

## **Gie Myung Lee**

*George Mason University*

The Clearinghouse Project

## **Srishanth Ravi**

*Thomas Jefferson High School*

MSallNet: A Deep Learning Approach for Multiple Sclerosis Lesion Analysis

## **Trevor Schonbrun**

*James Madison University*

Deep Metric Learning for Ancient Coin Identification

## **Bhavna Sheth**

*Radford University*

Navigating Early Onset Dementia: A Scoping Review of AI Applications

## **Samarth Singh**

*University of Virginia*

An Interpretable Medical Image Analysis Framework for MR-Guided FUS Thalamotomy

## **Ryan Stuck**

*Radford University*

Customizing Generative AI Systems for the Classroom

## **Maasai Thornton**

## **Trevor Tran**

*Hampton University*

Bias in Skin Cancer Detection: Improving Accuracy for Darker Skin Tones Using Artificial Intelligence

## **Maryna Veksler**

*Virginia Commonwealth University*

Practitioner-Driven Framework for AI Adoption in Digital Forensics

## **Murong Yue**

*George Mason University*

MATHVC: An LLM-Simulated Multi-Persona Virtual Classroom for Mathematics Education

## **Xiaomeng Wang**

*Virginia Tech*

Chiplets Design Flow with AI

## **Min Zhang**

*Virginia Tech*

Don't Go to Extremes: Revealing the Excessive Sensitivity and Calibration Limitations of LLMs in Implicit Hate Speech Detection



# PARTICIPANT BIOGRAPHIES

**Scott Acton** was an organizer of the summit. He is the American Telephone and Telegraph Company Professor and chair of the Charles L. Brown Department of Electrical and Computer Engineering at the University of Virginia and serves as the AI advisor to the provost. He specializes in biological/biomedical image analysis and leads the Virginia Image and Video Analysis Laboratory. He is a fellow of the Institute of Electrical and Electronics Engineers.

**Charles Clancy** is a senior vice president at MITRE and head of MITRE Labs. He oversees the work of more than 4,000 scientists, technologists, and engineers who provide sophisticated technical capabilities and solutions to the six federally funded R&D centers MITRE operates on behalf of the U.S. government. He is a fellow of the Institute of Electrical and Electronics Engineers.

**Missy Cummings** is a professor at George Mason University and the director of its Autonomy and Robotics Center. She holds faculty appointments in the Mechanical Engineering, Electrical and Computer Engineering, and Computer Science Departments. Her research interests include human-unmanned vehicle interactions, human-autonomous system collaboration, human systems engineering, the public policy implications of unmanned vehicles, and the ethical and social impact of technology.

**Luiz DaSilva** is the inaugural executive director of the Commonwealth Cyber Initiative and the Bradley Professor of Cybersecurity at Virginia Tech's Bradley Department of Electrical and Computer Engineering. His research focuses on

adaptive resource management in wireless networks, including cognitive networks and the application of game theory to wireless networks. He is a fellow of the Institute of Electrical and Electronics Engineers.

**Alfred Grasso** is the president of the Virginia Academy. He is the former president and CEO of The MITRE Corporation, a not-for-profit company operating federally funded research and development organizations. During his tenure, he led the corporation's overall strategic and business operations, overseeing 8,200 employees, two major campuses, and more than 60 domestic and international locations.

**Yael Grushka-Cockayne** is vice dean and professor of business administration at UVA's Darden School of Business and codirector of the Collaboratory for the Advancement of Education and Data Science. Her research and teaching focus on data science, forecasting, project management, and behavioral decision-making. Grushka-Cockayne is also a frequent speaker at international conferences on decision analysis, project management, and management science.

**Nathaniel Hendrix** is a researcher and data scientist with the American Board of Family Medicine and its Center for Professionalism and Value in Health Care. His research focuses on natural language processing for clinical notes, epidemiology, and artificial intelligence for decision support. Among other questions, he is interested in how the heterogeneity of diagnoses by different clinicians can be reduced.

**Khan Iftekharuddin** is a professor in the Department of Electrical and Computer Engineering at Old Dominion University and director of its Vision Lab, which focuses on computer vision, artificial intelligence, and machine learning. His research interests include computational image analysis, machine learning, biomedical imaging and signal analysis, big data analytics, pattern recognition, and computer vision modeling, with application to real-world problem-solving.

**Ruoxi Jia** is an assistant professor in the Bradley Department of Electrical and Computer Engineering at Virginia Tech and director of the ReDS (Responsible Data Science) Lab. Her research interests include machine learning, security, privacy, and cyberphysical systems. Recently, she has been focusing on data-centric AI and trustworthy machine learning.

**Anuj Karpatne** is an associate professor in the Department of Computer Science at Virginia Tech, where he leads the Knowledge-Guided Machine Learning (KGML) Lab. In addition to KGML, his research interests include knowledge-based scientific applications of AI and physics-informed machine learning.

**Lingjia Liu** is professor in the Bradley Department of Electrical and Computer Engineering at Virginia Tech, a member of its Institute for Advanced Computing, and director of Wireless@Virginia Tech. His research interests include 6G networks, machine learning for wireless applications, open radio access networks, and dynamic spectrum access/sharing.

**Michelle Maldonado** was a keynote speaker at the summit. She represents Virginia's 20th House District in Northern Virginia. She is a former technology lawyer and tech leader who founded and chairs the General Assembly's bipartisan and bicameral Technology & Innovation Caucus, chairs its AI work group subcommittee, and is a steering committee member for a national AI work group composed of state legislators from around the country.

**Milos Manic** is a professor in the Computer Science Department at Virginia Commonwealth University and is the director of its Cybersecurity Center. In addition to cybersecurity and the protection of critical infrastructure, his research interests include data mining and machine learning applied to energy optimization, resilient control, and human-machine interaction. He is a fellow of the Institute of Electrical and Electronics Engineers.

**Madhav Marathe** is an endowed distinguished professor of biocomplexity, executive director of the Biocomplexity Institute, and a professor of computer science at the University of Virginia. His areas of expertise include digital twins, network science, artificial intelligence, multi-agent systems, high-performance computing, computational epidemiology, biological and socially coupled systems, and data analytics. He is a fellow of the Institute of Electrical and Electronics Engineers, the Association for Computing Machinery,



the American Association for the Advancement of Science, and the Society for Industrial and Applied Mathematics.

**Mark Maybury** is vice president of commercialization, technology, and strategic innovation for Lockheed Martin, responsible for leading efforts to commercialize dual-use products and services across the corporation. He is also a member of the Board of Directors of its Astris AI subsidiary. He is a fellow of the Institute of Electrical and Electronics Engineers and the Association for the Advancement of Artificial Intelligence.

**Jon McManus** is vice president and chief data and AI officer at Inova, a nonprofit healthcare provider in Northern Virginia. He has expertise in directing technical teams and managing data and analytics capabilities in complex healthcare environments and has held leadership positions in enterprise data governance, application management, and IT systems at several healthcare organizations.

**Nathan Mellis** is vice president of AI for Maximus, a provider of government services. He is an expert in artificial intelligence, machine learning, data science, predictive modeling, automation, and decision analytics.

**Andrew Miner** is the chief health informatics officer and an emergency medicine physician at Inova, a nonprofit healthcare provider in Northern Virginia.

**T. M. Murali** was an organizer of the summit. He is a professor and associate department head for research in the Department of Computer Science at Virginia Tech. He directs the NSF COMPASS Center and the Virginia Tech Pandemic Prediction and Prevention Destination Area. Murali's research group develops phenomenological and predictive models dealing with

the function, behavior, and properties of large-scale molecular interaction networks in the cell.

**Arvind Narayanan** was a keynote speaker at the summit. He is a professor of computer science at Princeton University and the director of the Center for Information Technology Policy. He studies the societal impact of digital technologies, especially AI, and was among the first to show how machine learning reflects cultural stereotypes.

**Manish Parashar** is the inaugural chief AI officer for the University of Utah and presidential professor in its Kahlert School of Computing. He is the executive director of the Scientific Computing and Imaging Institute and leads the university's One-U Responsible Artificial Intelligence Initiative. He is a fellow of the Institute of Electrical and Electronics Engineers, the Association for Computing Machinery, and the American Association for the Advancement of Science.

**Robert Phillips** was an organizer of the summit. He is executive director of the Center for Professionalism and Value in Health Care at the American Board of Family Medicine and is professor of family medicine at Georgetown University and Virginia Commonwealth University. He practices one day each week in a community-based residency program in Fairfax, Virginia. He is a member of the National Academy of Medicine.

**Mihai Pop** is a professor of computer science and codirector of the University of Maryland Center of Excellence in Microbiome Sciences. He develops computational approaches to analyzing microbial communities, particularly for characterizing their strain-level diversity. Other interests include biological databases, antibiotic resistance, and software testing. He is a fellow of the

Association for Computing Machinery and the International Society for Computational Biology.

**Naren Ramakrishnan** was an organizer of the summit. He is the Thomas L. Phillips Professor of Engineering in the Department of Computer Science at Virginia Tech and directs the Sanghani Center for AI and Data Analytics and the Amazon-Virginia Tech Initiative for Efficient and Robust Machine Learning. He is a fellow of the Association for Computing Machinery, the American Association for the Advancement of Science, and the Institute of Electrical and Electronics Engineers.

**Mehul Sanghani** is the founder and former CEO of Octo, which IBM purchased in 2022 from Arlington Capital Partners for \$1.3 billion. Octo provides IT services, including Agile, DevSecOps, artificial intelligence, and cloud solutions, to the US federal government. A \$10 million gift from Sanghani and his wife, Hema, supports the Sanghani Center for Artificial Intelligence and Data Analytics at Virginia Tech's Innovation Campus in Alexandria and scholarships for minority students studying AI.

**Douglas C. Schmidt** is the inaugural dean of William & Mary's School of Computing, Data Sciences & Physics. His research covers a range of software-related topics including patterns, optimization techniques, and quality assurance of frameworks and model-driven engineering tools that facilitate the development of mission-critical middleware for cyber-physical systems and intelligent mobile cloud computing applications.

**Mark Sendak** was a keynote speaker at the summit. He is the former population health and data science lead at the Duke Institute for Health Innovation



at Duke University, a current member of the leadership team of the Health AI Partnership, and founder of Vega Health. He launched the company to help health systems purchase the best value artificial intelligence for their patients and solve the persistent issue of commercialization for AI inventors.

**Amarda Shehu** was an organizer of the summit. She is a professor in the George Mason University Department of Computer Science and its inaugural vice president and chief AI officer. Her research interests include artificial intelligence, stochastic optimization, machine learning, deep learning, optimization for deep learning, generative models, language models, bioinformatics, and computational biophysics. She is a fellow of the American Institute for Medical and Biological Engineering.

**Paul Torrey** is an associate professor in the University of Virginia's Department of Astronomy. A computational astrophysicist, he focuses on the formation and evolution of the cosmos. Torrey's research group builds, runs, and analyzes large-scale cosmological simulations that allow it to examine in a new way our assumption of the universe.

**Ryan Wright** is the Rolls-Royce Commonwealth Professor of Commerce at the University of Virginia's McIntire School of Commerce, co-director of AI Research @ UVA, and special advisor to the provost on AI research. He leads university-wide and McIntire-based strategic initiatives in artificial intelligence research and academics, oversees collaborative AI efforts across disciplines, and advises on the advancement of AI strategy and governance at UVA.

**Pinar Yanardag-Delul** is an assistant professor in the Department of Computer Science at Virginia Tech and director of its GEMLAB. Her research concentrates on the development of generative AI methods for image generation and editing, video and motion editing, and personalized compositions.

**Ziyu Yao** is an assistant professor in the Department of Computer Science at George Mason University and co-leads its Natural Language Processing Group. Her research focuses on the use of large language models (LLMs) for reasoning and planning, actionable interpretability of LLMs, and human-AI/LLM interaction.

# VASEM BOARD OF DIRECTORS

## **JAMES AYLOR, PhD**

*Immediate Past President*

Dean Emeritus, School of Engineering and Applied Science and the Louis T. Rader Emeritus Professor of Electrical and Computer Engineering, University of Virginia

## **ANTHONY BEASLEY, PhD**

Director, National Radio Astronomy Observatory

## **BARBARA BOYAN, PhD**

Former Dean, College of Engineering; Executive Director, Institute for Engineering and Medicine; and the Alice T. and William H. Goodwin Chair in Biomedical Engineering, Virginia Commonwealth University  
Member, National Academy of Engineering

## **ROBERT CAREY, MD**

Dean Emeritus, School of Medicine and the David A. Harrison III Distinguished Professor of Medicine, University of Virginia  
Member, National Academy of Medicine

## **PATRICIA DOVE, PhD**

University Distinguished Professor and the C.P. Miles Professor of Science, Virginia Tech  
Member, National Academy of Sciences

## **ANTONIO ELIAS, PhD**

Former Executive Vice President and Chief Technical Officer, Orbital ATK  
Member, National Academy of Engineering

## **ALFRED GRASSO**

*President*

Former President and CEO, MITRE Corporation

## **ANITA JONES, PhD**

Professor of Computer Science Emerita, University of Virginia  
Former Director of Defense Research and Engineering, U.S. Department of Defense  
Member, National Academy of Engineering

## **ROBERT KAHN, PhD**

*Board Member*

President & CEO, Corporation for National Research Initiatives  
Member, National Academy of Engineering, National Academy of Sciences

## **ALEX KRIST, MD**

Inaugural Sheldon M. Retchin Professor in Healthcare Innovation, Virginia Commonwealth University

## **CHEN-CHING LIU, PhD**

The American Electric Power Research Professor and Director Emeritus of the Power and Energy Center, Virginia Tech  
Member, National Academy of Engineering

## **GENERAL LESTER L. LYLES (Ret.)**

Former Commander, Air Force Materiel Command  
Member, National Academy of Engineering

## **X. J. MENG, MD, PhD**

*Vice President*

University Distinguished Professor of Molecular Virology, Virginia Tech  
Member, National Academy of Sciences

## **ROBERT PHILLIPS, MD**

Vice President of Research and Policy, American Board of Family Medicine  
Member, National Academy of Medicine

## **DAVID ROOP**

*Vice President*

Former Director of Electric Transmission, Dominion Energy  
Member, National Academy of Engineering


## **JENNIFER WEST, PhD**

*Secretary / Treasurer*

Dean, School of Engineering and Applied Science and Nancy and Neal Wade Professor of Engineering and Applied Science, University of Virginia  
Member, National Academy of Engineering, National Academy of Medicine

## **A. THOMAS YOUNG**

Former Executive Vice President, Lockheed Martin  
Former Director, Goddard Space Flight Center, NASA  
Member, National Academy of Engineering





# VASEM

Virginia Academy of Science, Engineering, and Medicine

Gateway Plaza  
800 East Canal Street, 11th Floor  
Richmond, VA 23219  
[www.VASEM.org](http://www.VASEM.org)  
[info@vasem.org](mailto:info@vasem.org)

