

Nearly 50 years of AI at ASU

Before it became a household name, researchers at ASU were building the foundations of artificial intelligence as we know it today

By Lisa Robbins, ASU News
May 20, 2026

Editor's note: This story was featured in the [summer 2026](#) issue of ASU Thrive.

Story by Ellen Sheng

When Professor Subbarao Kambhampati started working in the field of artificial intelligence in 1983, his friends and colleagues felt sorry for him.

People asked, “Why can’t you do something more useful, such as databases, such as software engineering?” he recalls. At the time, “AI was an intellectual discipline that didn’t yet have perceivable benefits right away.”

Fast-forward nearly 50 years, and AI has become one of the biggest games in town. But the transformation didn’t happen overnight. Since the term AI was coined at a 1956 conference, sometimes called the “Constitutional Convention of AI,” the field has lurched between hype and “AI winters” when both funding and interest collapsed.

Throughout those ups and downs, researchers at ASU’s computer science department, which was founded in 1980, and even before it was a department that has now become part of the [School of Computing and Augmented Intelligence](#), have been steadily working on the fundamentals of AI: teaching a machine to absorb knowledge, make plans and move closer to human intelligence — long before its commercial value became obvious.

“We are on the front line,” says Regents Professor Huan Liu.

Fixing AI errors

In 1986, as a 20-year-old at a university in West Bengal, Professor Chitta Baral found his calling in a quest to teach computers to think.

“We humans make decisions on incomplete information,” Baral says, “and when we are given new knowledge, we are OK to change our minds.”

This is why using mathematical logic as the framework proves brittle — because human reasoning is flexible.

Reasoning this way can cause mistakes because machines lack underlying world models. To solve this problem, Baral spent decades developing a “knowledge calculus,” a new reasoning methodology that allows machines to update understanding when receiving new information.

That work, which continued after he joined ASU in 1999, is now recognized as foundational to one of AI’s central unsolved problems, and for helping solve it in certain situations. Current large language models may be fluent in English, but, at times, they still make world model errors, confidently producing answers that ignore basic facts about how things work.

While ChatGPT or Claude can give these human-like answers that sound authoritative, as discussed in AI circles, some models sometimes make amusing mistakes. In one example, Baral recalls hearing a system suggest someone walk 200 meters to the car wash instead of drive because it’s so close anyway.

In 2021, in collaboration with researchers at the Allen Institute for AI, Baral and his ASU doctoral student Swaroop Mishra, '23 PhD in computer engineering (computer systems), now at Microsoft, formerly at Google DeepMind, published one of the earliest papers demonstrating that large language models could be trained to follow human instructions, a technique now called instruction tuning. OpenAI later scaled this approach into ChatGPT.

We are on the front line.

—

Huan Liu

Regents Professor in the School of Computing and Augmented Intelligence

Teaching machines to identify relevant data

These days, Professor Huan Liu is best known as a pioneer in social media data mining, which finds trends in massive social media databases. But when he started the early phases of AI and machine learning in the 1990s, the landscape looked very different.

ASU students are learning AI and transforming industries

Across campus and across degree programs throughout the university, students and researchers are using AI to reimagine their fields:

Research experimental work was done on desktop computers and researchers relied on small, well-curated datasets — like the Iris dataset, which was just 150 rows and four columns describing three varieties of flower — to develop and validate new algorithms.

There were no textbooks. Instead, Liu compiled his own course materials from research papers.

The constraints didn't dampen his enthusiasm. In an emerging field like he was helping to create, he noted, there's a freedom not possible in established disciplines.

"There were a lot of ideas you could pursue," he says, "instead of having to follow previous researchers."

As the data grew from 150-row datasets to millions of data points, the question of what to pay attention to became urgent. Liu's contribution is feature selection, which teaches machines to identify what information actually matters for a given task and what doesn't. Without it, models either drown in irrelevant data or become confused and inefficient.

He helped build workable mathematical foundations for how AI systems learn to focus, developing methods now embedded in how machines process language, behavior and social interaction. The methods he developed for finding patterns in noisy, human-generated data helped lay the groundwork for how AI systems now interpret human communication and behavior.

He has graduated 47 PhD students at ASU thus far, many of whom carried these methods into the next generation of AI research at leading companies such as General Dynamics Information Technology, Google, Microsoft and LinkedIn.

Helping AI plan

The Grand Canyon, with its cliffs and intense summer heat, could be deadly, but most people who visit survive because they avoid walking along the edges, study trail maps and bring water.

"People don't just show up, close their eyes and wander around," says Kambhampati, who arrived at ASU in 1991 but started in AI in 1983 with his undergraduate thesis. "But in essence, that's what computers do."

Accelerating medical discovery by using AI to analyze large sets of biological data and identify protein targets linked to side effects in cancer treatments

Making roads safer by training AI systems in simulation and testing them on physical models to improve how technologies like traffic control and self-driving cars perform in real-world conditions

Using AI to analyze digital evidence and uncover patterns that help investigators track cybercriminal activity and respond more effectively

Enabling students to create an "agentic self" — an AI persona reflecting their values, voice and goals — through a class taught by musician and innovator will.i.am on the EDU.FYI platform, a collaboration between ASU and FYI.AI

In addition, ASU offers a variety of flexible undergraduate and graduate degrees, including a doctoral degree in artificial intelligence for the next generation of leaders and innovators.

Learn more at ai.asu.edu.

Teaching machines to navigate the world, a critical skill for AI to work, has been Kambhampati's mission.

His focus is automated planning, ensuring an intelligent agent doesn't "step on its own toes" or make a mistake from which it cannot recover, like a Mars rover falling into a ditch.

"Planning is needed to make sure that even though the world is unforgiving, you can still live in it without catastrophe," he says.

While today's large language models draft everything from travel itineraries to wedding plans, Kambhampati's research proved that today's models still often make errors in planning because they lack underlying reasoning capabilities.

In one well-known experiment, Mystery Blocks World, he simply renamed the objects in a test. AI models that were working well suddenly collapsed to near-zero accuracy, revealing they had been pattern-matching against training data rather than reasoning.

"They will give you a great wedding plan," he says, "but they will likely miss the critical logistics, like having a caterer in place."

To bridge this gap, starting in 2022, he developed the LLM-Modulo framework. AI-generated plans are run through an external "verifier" to catch errors before any action is taken. It's much like a student checking answers against the back of a textbook.

This process of ensuring agents act safely in a world of irreversible consequences is what the field now calls AI safety, a discipline Kambhampati was defining decades before the term became a buzzword.

**There are only three fundamental questions facing humanity:
the origin of the universe, the origin of life and the nature of
intelligence. AI tries to have something to do with the last part.
... So, how could you not be interested?**

Subbarao Kambhampati

Professor in the School of Computing and Augmented Intelligence

From the lab to real world

Meanwhile, Ross Maciejewski, the current director of the School of Computing and Augmented Intelligence, starting back in the 2000s, focused on a different challenge: how to make the mountains of data useful for the people responsible for public safety.

Beginning in the mid-2000s, he worked with the Department of Homeland Security to analyze emergency room records. Every 12 hours, a new batch of records would come in and needed to be analyzed for anomalies or signs of potential disease outbreaks. In one memorable instance, his team's system flagged a cluster of patients from a single address complaining of headaches. They eventually discovered a carbon monoxide leak in the apartment building.

"Given the volume of records, you're not going to do that manually. So, we would think about how to use novel AI techniques and build new AI techniques to support the support systems," Maciejewski says.

He built on that early work and applied game theory to defense, helping the Coast Guard and Transportation Security Administration develop randomized patrol routes to make it harder for contraband to slip through ports and airports.

Today, his research centers on human-AI teaming, studying how humans, such as airport security officers, should interact with AI alerts. How much should they trust a machine's suggestion, and how do we ensure humans remain responsible for the final, high-risk decisions? His work helps ensure that as AI grows more powerful, it remains a tool that enhances human judgment, rather than replacing it.

The long game

ASU's researchers have kept working in the field when it was unfashionable and there were no textbooks. For these faculty, the explosive popularity of AI today is the latest chapter in a long pursuit of discovery that predates hype and sets the stage for what AI can do today and tomorrow.

For Kambhampati, who has been at ASU the longest, the work was never about technology.

"There are only really three fundamental questions facing humanity: the origin of the universe, the origin of life and the nature of intelligence. AI tries to have something to do with the last part, the nature of intelligence. So, how can you not be interested?"

About the author

A writer and editor with a focus on business and finance, Ellen Sheng's work has appeared in publications including The Wall Street Journal, CNBC, Forbes, Fast Company, Real Simple and Marie Claire.

This story originally appeared on [ASU News](#).

Main image




ASU researchers used IBM personal computer systems in 1984, when desktop computing first supported the development of AI. Photos courtesy ASU Archives

Text image(s)



A discussion unfolds beside ASU's Univac system in 1976, when large-scale machines and magnetic tape defined computing. Photo courtesy ASU Archives



In 1992, now-Professor Chitta Baral assembles an early AI-enabled robotic prototype. Photo courtesy ASU Archives



Professor Huan Liu (center) works with students in his lab in 2019. Liu is a global leader in social media data mining research. Photo by Erika Gronek



Professor Subbarao Kambhampati tests a robotic arm in his lab in 2017 where he explores applications of AI for autonomous systems. Photo by Marco-Alexis Chaira



Professor Ross Maciejewski developing AI applications for public safety in his lab in 2018. Photo by Jessica Hochreiter