

Computer science grad earns IBM fellowship for AI research

Naman Ahuja's work focuses on turning unstructured text into accurate, usable data for real-world applications

By Kelly deVos, ASU News
April 30, 2026

Editor's note: This story is part of a series of profiles of notable [spring 2026 graduates](#).

Modern medicine runs on data organized into tables and charts that help doctors make fast, informed decisions. Survival rates and prescription side effects are often distilled into formats that can be scanned in seconds. But that clarity doesn't come easily. The data must be extracted and organized, often through a time-consuming process that can leave critical information buried deep in research reports.

Naman Ahuja wants to change that.

This May, Ahuja will graduate from the [School of Computing and Augmented Intelligence](#), part of the [Ira A. Fulton Schools of Engineering](#) at Arizona State University, earning a master's degree in computer science.

Over the past two years, his research has focused on how to get artificial intelligence, or AI, systems to convert long, unstructured text into accurate, usable tables. This spring, that work earned him an IBM Infrastructure Master's Fellowship Award, a prestigious honor that recognizes research with strong real-world and industry impact.

Looks right, isn't right

The problem exposes a key limitation of modern AI. Large language models can read and summarize documents with ease. But when asked to extract precise information and organize it into something structured — like a table a doctor or analyst could rely on — they often struggle. Important details can be missed, information can become inconsistent, and models may generate claims that aren't supported by the original text.

The result can look polished, but it isn't always reliable. Ahuja's research focuses on closing that gap.

“In the real world, a lot of data exists in complex and semi-structured formats, like PDF documents or Wikipedia pages,” he says. “These documents have some structure, but they’re still complex and contain a lot of information.”

His solution, developed through his master’s thesis, rethinks how AI should approach the problem.

Instead of asking a model to generate a table in one pass, Ahuja breaks the task into steps. First, the system extracts atomic facts from the text. Then it builds a plan for how the table should be organized. Finally, it fills in the table incrementally, updating entries as new information appears.

The approach mirrors how a human might do the same work: read carefully, decide what categories matter, and then populate the table piece by piece. His thesis argues that reliable structured generation is about breaking complex tasks into smaller, verifiable steps that reduce errors and improve traceability.

That traceability matters most in high-stakes environments like health care, where clinicians often conduct systematic reviews, reading large volumes of research and extracting key findings into tables for decision-making. It’s a time-consuming, manual process where mistakes can carry real consequences.

“We can use these heterogeneous documents and convert them into a structured data form so that we can access that information more readily,” Ahuja says. “That can help reduce repetitive data extraction and allow clinicians to focus more on interpreting results.”

Ahuja’s approach is also designed for what researchers call living data, or information that evolves over time. As new studies are published, systems like the one he is developing can update existing tables instead of rebuilding them from scratch, maintaining consistency while incorporating new evidence.

That emphasis on reliability and real-world usability is part of what drew attention from IBM’s fellowship program.

[Vivek Gupta](#), a Fulton Schools assistant professor of computer science and engineering and head of ASU’s [Complex Data Analysis and Reasoning Lab](#), or CoRAL, where Ahuja conducted his research, sees the work as part of a broader shift in AI.

“Naman’s work really captures what we’re trying to do in CoRAL. We’re focused on complex structured data, especially how to generate it and evaluate it correctly, so we can build AI systems people can trust in real-world settings,” Gupta says. “He’s been very thoughtful about both the methods and how to evaluate them, and the IBM fellowship is well-deserved.”

Turning the tables on the future

For Ahuja, the path to that work began in Hyderabad, India, where he completed his undergraduate degree in computer science before coming to ASU in 2024. His focus has consistently been on turning research into something usable beyond the lab.

At ASU, he served as a teaching assistant for a graduate-level natural language processing course, delivered a guest lecture on neural networks, and presented his research at an

international conference in Vienna. Along the way, he credits Gupta's mentorship with helping him navigate the inevitable setbacks of research.

"Dr. Gupta has helped guide me through all the different tough aspects, which is important when experiments don't pan out the way you think," Ahuja says.

Outside the lab, Ahuja tries to stay balanced. He plays basketball regularly, explores new music, and unwinds by watching stand-up comedy in both Hindi and English.

As he prepares to graduate, Ahuja is already stepping into the next phase of his career, having accepted a full-time role at Amazon in Seattle, where he will continue working on large-scale systems.

"I'm interested in working on the core systems, how these models are actually built and how they can be used to solve real-world problems," Ahuja says.

The move aligns closely with his long-standing interest in applying research to industry challenges. And while his immediate future is now set, the broader problem he's focused on isn't going away. The world is producing more text than ever, and the need to turn that text into usable knowledge is only growing.

For now, that still often means someone, somewhere, building a table by hand. Ahuja's work suggests they don't have to.

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

Main image



Naman Ahuja poses in front of the Tempe, Arizona, skyline. This May, Ahuja will earn a master's degree in computer science and has received an IBM Infrastructure Master's Fellowship Award for his innovative artificial intelligence research. Photo courtesy of Naman Ahuja