



# HensonNews

SPRING 2026

Produced by Salisbury University's Henson School of Science and Technology

## Henson & Artificial Intelligence SPOTLIGHTS

In this issue of *Henson News*, faculty and students share their perspectives and experiences with AI – the good, the bad, and the surprising.

### From the Henson Dean's Office



*It is likely that readers of this newsletter are inundated with doom-and-gloom news stories about the rise of generative artificial intelligence (AI). This is true in higher education as*

*well. While the negative downstream effects of AI are real and worth mitigating, there is a great deal of exciting innovation and creativity surrounding the adoption of these technologies as well. In fact, Henson School of Science and Technology faculty and students are on the cutting-edge of that work. Our collective belief is that while some of the functionality of AI will either eliminate or minimize our more mundane tasks, nothing will replace a.) the ability to probe questions that have not yet been asked and b.) learning how to use advanced technologies to solve important and complex problems. That is what our outstanding faculty teach students at Salisbury University, and that won't be replaced by AI any time soon!*

– Dr. Michael Scott, Dean

### OpenClaw, MCP and Agentic AI

Dr. Junyi Tu  
Computer Science Department

Artificial intelligence is undergoing a shift from passive assistants to autonomous, action-taking agents. A key driver of this transformation is OpenClaw, an open-source, local-first AI framework that enables agents to operate directly on a user's machine – executing shell commands, managing files, navigating browsers, and interacting across messaging platforms. Its rise illustrates the broader movement toward AI systems capable of understanding goals, decomposing tasks, and completing end-to-end workflows.

OpenClaw began in late 2025 as Clawdbot, created by developer Peter Steinberger, who initially built it as a simple relay connecting WhatsApp to a Large Language Model. Its unique “agent that actually does things” approach drove rapid viral adoption. After trademark challenges, it briefly adopted the name Moltbot before settling on OpenClaw in early 2026. Within months, it became the most starred non-aggregator GitHub project, surpassing React and gaining attention from OpenAI CEO Sam Altman, who later hired Steinberger to lead personal-agent development.

The Model Context Protocol (MCP) originated with Anthropic, the company that created Claude Code, as an open standard to solve integration fragmentation in 2024. MCP introduced a universal, self-describing interface for tools, data, and structured prompts – quickly adopted by OpenAI, Google DeepMind, and the broader ecosystem. By 2026, MCP servers provided standardized access to hundreds of external systems, becoming a foundational layer for agentic frameworks like OpenClaw.

OpenClaw represents a major advancement in artificial intelligence because it transforms AI from a passive, text-based assistant into an active system capable of performing meaningful operations in real environments. Its local-first “gateway” architecture allows AI models not only to interpret instructions but also to take direct action – managing files, executing commands, automating browsers, and interacting across communication platforms. This shift moves AI closer to functioning as a true digital collaborator rather than a tool confined to generating suggestions. MCP amplifies this progress by acting as a universal connector that lets OpenClaw interact with a wide range of external services – such as email platforms, messaging tools, databases, and productivity systems – without requiring custom integration code. Together, OpenClaw and MCP enable AI agents to operate autonomously across diverse digital ecosystems, marking a pivotal evolution toward more capable, flexible, and integrated AI systems that can handle complex workflows from end to end.

Because OpenClaw-style agents possess deep system access – including shell execution, filesystem control, and browser automation – securing them requires treating AI processes with the same rigor applied to high-privilege software. To mitigate risks highlighted by real incidents where autonomous agents accidentally modified or deleted user data, users should run agents in sandboxed or containerized environments that restrict operating system privileges, applying strict least-privilege controls when connecting through MCP.


## Using AI to Predict Proteins Structures

Riley Bozman

Biology – Biomedical Science Concentration Major

When people first think about biology, what often comes to mind are images of beakers, plates of bacteria, or maybe even a DNA double helix. While these images justifiably depict different fields within biology, they fail to capture its interdisciplinary nature in which chemists, physicists, and, most notably, computer scientists play invaluable roles. As I began my research with Dr. Kirsten Guckes, I too fell into this trap. Prior to my work in the Guckes lab, I had an interest in computer technology, particularly digital simulations, and sought ways to incorporate it into my research, which led me to discover computational biology.


Artificial intelligence within biology can be a very powerful tool, granted you know how to use it. In my work, I was particularly interested in modeling how proteins are structured and folded in 3D space, and how this influences their interactions with other biological compounds. Before the advent of AI, to gain insight into a protein's shape, you had to physically crystallize a pure sample, direct X-ray beams at it, evaluate how those X-ray beams scattered, and use that information to identify the protein structure. Like many things in science, this is easier said than done. For reference, crystallizing a novel protein can cost over \$100,000 for a simple protein and can quickly skyrocket to over \$1 million for membrane-bound proteins. Even after spending this exorbitant amount of money, there is still



AlphaFold generated model of the VasX toxin.

less than a 50% chance that the imaging will work, depending on the protein's cooperation. Fortunately, a tool called AlphaFold allows us to make predictions about a protein's structure without expensive techniques and expertise that are impractical for an institution like Salisbury University that focuses on developing research skills in budding scientists.

AlphaFold is an AI system developed by Google that can generate predictions of a protein's structure based solely on its sequence. Although these predictions have limitations, both computational and experimental methods have shown these models to be highly accurate. Using this tool, I generated a putative structure of a bacterial toxin, VasX, which is the focus



Dynamic model of VasX interacting with a bacterial membrane.

of my research. With this data, I can predict what this toxin may target within a cell to inflict damage. For example, I was able to simulate how VasX may interact with a bacterial cell's membrane, informing new hypotheses about how VasX kills competitor bacteria. Having this information allows us to generate focused projects that are more feasible for novice researchers entering the field. It is important to note that these models are not perfect, so it is vital to confirm our hypotheses in the lab setting. Nevertheless, my work in the Guckes lab using AI has opened my eyes to its importance in biology, in allowing undergraduates to perform high-level research without the barrier of high-cost laboratory equipment.

## Using AI to Impact Teaching

Dr. David Keifer

Chemistry Department

I wanted to give ChatGPT a real test to see if it could aid my teaching, so I asked it to give me a high-level critique of my course materials for CHEM 122, General Chemistry II, based on the scholarly literature on teaching and learning. I dumped my entire folder of course materials into it: the syllabus, in-class work, lecture slides, everything. ChatGPT responded with four pages of what sounded like sensible comments and suggestions, starting with things I was doing that are backed up by the literature but also pointing out some things that it deemed to go against the literature. I actually ended up implementing some of its suggestions!

Let me give you an example. I started flipping my courses a few years ago, and I have students submit their notes on my pre-class recordings. ChatGPT gave what seemed to be an evidence-based argument for why requiring students to submit notes was a terrible idea: It enforced compliance with course policy but did not lead to learning. Its argument was compelling, but

then I realized I ought to be a bit more skeptical. The real answer, according to my read of the literature, was more nuanced. Requiring notes to be submitted can have educational benefits, but only if the instructor invests significant time in teaching students how to take good notes and providing feedback on notes. Otherwise, students just figure out what I am looking for in their notes and then write that down, rather than actually using their notes as an educational tool. I ended up deciding that I would be better off investing my time and my students' time in other aspects of the course, so I will no longer require students to submit their notes. That was a tough decision for me because I worried that students would no longer watch the pre-class recordings if they did not have to provide any proof. After some further conversations with ChatGPT, I decided that I could offset that risk by increasing the weight of my daily mini-quizzes that I give at the beginning of class. If students do not actively pay attention to the recordings and take good

notes, then they will almost certainly do very poorly on the quizzes. Hopefully, with a little bit of coaching, I can help students learn how to take educationally beneficial notes without actually having to grade all of them. It could be a win-win! I plan to implement this change next semester to see how it goes.

Here's what I took away from this experience: AI can impact teaching just as much as it can impact learning, but there are good and bad ways of using it. It pointed out several things in my course structure that I had never even thought about, and I believe there can be real benefit in considering its comments. But it is absolutely vital to use its responses as a jumping off point rather than uncritically believing what it tells you. It turns out that everything I tell students about the risks and benefits of AI applies to instructors, too.

# When Answers Are Abundant, What Do We Want Our Students To Become?

Dr. Sophie Wang

Computer Science Department

“Professor, I used ChatGPT to study for my midterm. Why didn’t I still earn an A?”

This question is enough to make every educator feel unease because it precisely captures the challenge we are facing today. According to Stanford University’s *2025 AI Index Report*, artificial intelligence (AI) systems are rapidly approaching or even surpassing human baselines in reading comprehension, Ph.D.-level scientific problem solving, English language understanding, and mathematics competitions. This rapid advancement of technology presents unlimited opportunities, and young people, particularly our students, are among the earliest adopters. Surveys from around the world show that an increasing number of K-12 and college students are integrating AI systems into their learning.

In students’ mind, AI can draft essays, generate code, solve complex problems, summarize key points, and polish expression. It is always available, patient, meticulous, and rarely makes mistakes. Given such a seemingly great tool, why wouldn’t everyone use it?

Yet, university faculty responses are more complex. In traditional education, the teacher’s role has been stable and well-defined: interpreter of knowledge, evaluator of performance, and guardian of academic standards. AI challenges that stability. A recent nationwide survey of U.S. college instructors by the American Association of Colleges and Universities found widespread concern that generative AI might undermine higher education’s core mission. Many fear it will

lead to overreliance, weaken critical thinking, compromise academic integrity, and devalue degrees. Some institutions have attempted to revert to pen and paper exams, oral questioning, or in-class writing. Others have imposed outright bans. But as generative AI grows more sophisticated and pervasive, resistance alone is no longer sufficient; educators must confront a deeper question: What is the true purpose of education?

Thinking has long been considered a human intelligence. In 1950, British mathematician Alan Turing asked a question that would shape the field of AI: “Can machines think?” He proposed the imitation game, later known as the Turing Test, asking whether a machine that cannot be distinguished from a human in conversation can be said to possess intelligence. Over seven decades, AI research has evolved through cycles of hope and setback, from rule-based systems to machine learning, and now to generative models, continually seeking to answer Turing’s question.

The key breakthroughs of human civilization have never begun with answers, they have always begun with questions. When Albert Einstein was a teenager, he wondered: “What would I see if I rode alongside a beam of light?” This question planted the seeds of the theory of relativity. Today, as AI can produce answers within a split of second, the goal of learning should extend beyond obtaining answers to interrogating them: Are these answers reliable? What assumptions underlie them? What factors are overlooked? Formulating profound questions

is in itself a form of learning, compelling us to analyze, critique, connect, and create beyond existing knowledge.

Children enter the world brimming with curiosity. To them, nothing is self-evident. An ant, a cloud, a flower, anything can spark a cascade of “whys.” Yet, as children grow, the number of questions often dwindles. Modern education prioritizes correct answers and standard solutions during exams. Students learn to find answers but gradually lose the habit of asking questions. Up until this moment, this was less problematic: Correct answers indeed lead to good grades, and grades have long been a cornerstone of education. But in the era of generative AI, when answers are abundant and instantaneous, their inherent value is also diminishing. The ability to pose meaningful questions has become more scarce and more precious.

This is a challenge but could also be an opportunity for education. When machines excel at answering questions, the value of education and learning can no longer lie in providing and obtaining answers alone. Rather, education must cultivate individuals who can ask questions, think independently, and engage critically with the world.

In the face of this “intelligence tsunami,” education is not competing with machines. It is testing itself. It must determine whether it can safeguard the unique human capacity – the curiosity to explore, the courage to challenge authority, and the insight to ask questions for which answers may not exist yet.



HENSON SCHOOL  
OF SCIENCE AND TECHNOLOGY  
1101 Camden Avenue  
Salisbury, MD 21801-6860

ADDRESS SERVICE REQUESTED

SJ is an Equal Opportunity/AA/Title IX university and provides reasonable accommodation given sufficient notice to the University office or staff sponsoring the event or program. For more information regarding SJ's policies and procedures, please visit [salisbury.edu/equity](http://salisbury.edu/equity).