Agent-based Cyber Threat Intelligence and Operation

New ACTION Institute Teams Humans and AI to Fight Cyberattacks

"These are some of the very best people in AI and security. They have been at the forefront of expanding the foundations of AI, machine learning, game theory, and computer security."

Imagine a modern city. Imagine the infrastructure that keeps it running smoothly: an aqueduct, a power plant, a smart-city system running on open-source software to monitor and distribute services. Now imagine a hostile nation-state that targets that software system, identifies vulnerabilities, gains administrative access to the main server, and shuts down the aqueduct and the power plant, causing paralysis and chaos.

The scenario is hypothetical—for now. It's also realistic enough that the National Science Foundation recently named UCSB as the lead institution on a five-year, $20 million grant to pursue new approaches to cybersecurity.

"The basic idea of the ACTION Institute is to combine two cultures," explains computer science professor **Giovanni Vigna**, the institute's director. "One looks for new ways to do AI and the other looks to use AI in new ways to improve security. We hope that by putting them in the same room, something amazing will result."

The problem is complex. Countering the evolving tactics, techniques, and procedures of bad actors as they attack massive, mind-bogglingly complicated, and rapidly changing systems is increasingly beyond the capabilities of human defenders alone. There's just too much happening, and it's happening too quickly.

"Solving the time-and-scale problem will require automation," Vigna says. "And it has to be smart automation, and that means AI."

Complicating matters even further is the idea that attackers will also be using automation and AI to overcome cyberdefenses. Responding to these threats will require reasoning and acting based on small amounts of data, while adapting to untrainable and unspecified scenarios.

"In the case of security," says UCSB CS professor **Ambuj Singh**, a co-PI on the project, "humans and AI may have different perspectives on the implications of actions. This is where the synthesis of humans and AI becomes useful."

The integrated approach will use AI agents to prevent or repel attacks in time and at scale, supported by extensive domain knowledge, logic-based reasoning, and human-agent and agent-agent interactions.

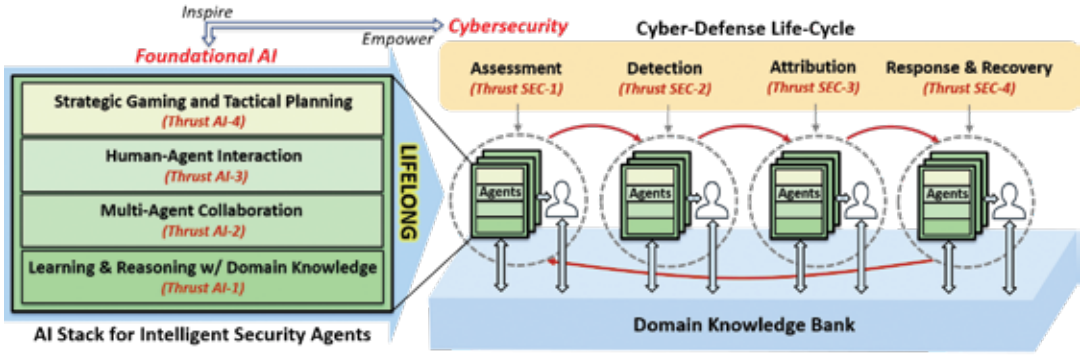

Research thrusts are highly integrated and interdependent—four each in foundational AI and cybersecurity.

According to the project proposal, over time these intelligent security agents will become increasingly robust and effective, capable of composing strategies and plans in the face of uncertainty. They will be more collaborative with each other and with humans for mutually complementary teaming, and for adapting to unfamiliar attacks.

Development work on the project is being shared among 11 academic institutions, with the goal of answering a series of research questions about the capabilities needed to build these intelligent security agents—questions covering such topics as learning and reasoning, human-agent and agent-agent interaction, and game theory.

"These are some of the very best people in AI and security," says Vigna. "They have been at the forefront of expanding the foundations of AI, machine learning, game theory, and computer security."

They will each be working primarily in one of eight highly integrated and interdependent research thrusts—four each in foundational AI and cybersecurity.

The ACTION Institute, in turn, is just one part of a much larger investment by the NSF to advance a cohesive national approach to AI-related opportunities and risks.

"If you use AI wrong, you can hurt entire classes of people," explains **Chris Kruegel**, another ACTION Institute co-PI and UCSB professor. "And that has occurred. So we have to be careful about what we encode in the agent's knowledge, what we learn from the data, how we learn it, and how we align it to conform to the highest ethical values. Developing AI that is ethical and trustworthy is not an option; it's the only thing you can do, and it's ingrained in this community."

"We want to have ethical AI," adds Vigna. "We don't want it to be making decisions that could cause harm … We want to be sure that decisions are made with a human in the loop, but in an efficient, targeted way that makes the best use of that person's capability."

World's Largest Raspberry Pi Cluster Finds New Home at UCSB Computer Science

What is it?

A supercomputer consisting of 1,050 Raspberry Pi 3B+. The next largest system that we know of is a cluster of 750 at Los Alamos National Lab.

Why is it?

The short answer, according to creator **Chris Bensen** in his series of videos detailing the project, is "because it kicks a**." The longer answer lies in Oracle's support for a series of outside-the-box projects designed to attract the attention of engineers and developers.

How did it get to UCSB?

It started with PhD student **Animesh Dangwal**'s internship at Oracle Labs. Professors **Rich Wolski** and **Chandra Krintz** explain that as a result of Dangwal's outstanding work, his Oracle supervisor came to visit UCSB.

"In the course of discussing Animesh's research, we mentioned that we do a lot with Raspberry Pis," says Wolski.

They showed off their own cluster of 18 devices, and the supervisor wondered if their lab would like some more that Oracle wasn't using.

"I'm thinking he had maybe 20," Wolski says. "I said yeah, absolutely!" He thought they could make room in the shoe rack that houses their current cluster. "The next thing that happens, here come a thousand of these things."

"This has been sitting at Oracle for two years not being used," adds Krintz. "They really wanted someone to take it forward. There's so much great technology here that they wanted to pass along for its next generation."

Bigger on the inside: Professors Rich Wolski and Chandra Krintz with engineer Chris Bensen and his creation.

How will we use it?

"Our students are going to go crazy on this," says Wolski. "We're going to use it in a couple of ways. We have large scale, high performance computing problems that we want to experiment with on a low power device. And from a high performance computing perspective, this is a very low power device. The other part is, we didn't have an instrument like this where we can train students on the systems end of it. This is a significant system."

"It's very different if I taught you how to manage 12 Raspberry Pis," says Krintz, "than it is to do something at this scale. This gives us the ability to train students in a way that nobody else on the planet can."

"You can make donuts in your kitchen," adds Wolski, "but if you want to make donuts for the Super Bowl, it's different—this is the Super Bowl."