Rotem Guttman: Welcome to season two of the SEI Cyber Talks. I'm talking today with Tom Longstaff. Thank you for joining us.

Tom Longstaff: Thank you very much. I love being here.

Rotem Guttman: So, we have you here today to talk about SEI and CMU's vision for the future of AI engineering.

Tom Longstaff: Absolutely.

Rotem Guttman: So, what are the-- what does it mean, AI engineering? What is the--

Tom Longstaff: So, the whole idea behind AI engineering is that if you're ever going to take advantage of the world of artificial intelligence or machine learning, or ever get any progress in that, then you have to actually have an engineering discipline in creating these tools so you can build them, you can trust them, and you can defend them. So, that's the whole idea. You can't just throw anything out there in AI that you might want to use and expect to have the war fighter support it in the way that they need to be or have critical operations that really then can take advantage of the latest things that AI can provide. And that is the exciting part of the vision that I'm really happy to share with you today.

Rotem Guttman: And so, what is the goal? What is the need for executing that vision? What's-- why do we need to use engineering practices in order to develop this?

Tom Longstaff: Well, look, the world is going in this whole direction of artificial intelligence/machine learning in everything that they're doing. We see it in the private sector, but it's happening in nation states. It's happening in criminal elements. Everywhere in the world, people are adopting these techniques and rocketing forward to create these higher capabilities that we will have to face. We'll have to face them operationally in theaters. We'll have to face them in criminal environments. We're going to face them everywhere.

We have the same ability to use those same tools, but if we do the same thing, if we only create ad hoc capabilities, then the best we can hope for is parity. You never want to go into an engagement with parity. You want to go into an engagement with dominance where you're really going to be far ahead of the people that you're engaging. The way that you do that is you create a discipline of engineering around the creation of these artificial intelligence tools.

Now, maybe that sounds a little bit like, "Oh, I'm creating bureaucracy. I'm going to slow everything down," but, in fact, no. What you do is you create a repeatable, rapid development, fielding, testing, and trusted operational context, where you can take a brand new capability fresh out of Carnegie Mellon University and move that directly into the field much faster than your adversaries can. That’s how you use AI to win.

Rotem Guttman: Absolutely. And so, what is the need specific to DoD that we can address with these capabilities because you mentioned assisting the war fighter?

Tom Longstaff: Absolutely. So, if you think about it, all the things that we do in the department of defense, that we do within the U.S. government, involves a whole lot of human activity. So, we do a lot of very detailed decision-making and planning and understanding how we’re going to do our engagements long before an engagement starts. When it’s a human manual process, it means a lot of people are cooperating to create and test out and understand what we’re going to do. Machine learning and artificial intelligence can help in the decision support systems to enable the human beings to much more rapidly come to closure on what is the most effective way to move ahead. So, decision support is one the most important areas we’re going to help.

And then, once you’re beginning an engagement, there’s the whole world of how do I manage the information, the situational awareness, of what’s happening before, during, and after an engagement. A lot of the ways you might think about this is just standard imagery, cameras, satellite imagery, overhead imagery from flying aircraft and drones, imagery taken from the ground, tons and tons of information coming back to somebody that has to assess what is going on and what they should do next. Well, if you use object detection and you use various machine learning aspects to label and track and understand exactly where things are evolving in that situation, you can provide a full, comprehensive, and understandable picture to those people that need to make a decision fast and then move out from that area.

So, the war fighters absolutely have to have those capabilities. They’ve got to move it in. And not just the people sitting back in Washington or sitting back at headquarters, but soldiers in the field need that information. They need it rapidly. They need it labeled.

Rotem Guttman: So, they can be more effective.

Tom Longstaff: Absolutely, so they can be more effective. So, that’s another good example of how we might actually apply it. So, basically, that’s why you want to those style systems. There are many others that we can also support in addition to vision, things like understanding audio and stress, things like really trying to understand better how to combine multi domain operations and make them really effective using patterns, which honestly this is what machine-learning tools do. All they really do is identify correlations and patterns in large amounts of disparate data. And that’s where we need to create those capabilities and getting them out there fast.

Rotem Guttman: Well, that sounds like it would be fantastic to have in the field. How do we get there? How do we transition from this view of an AI engineering discipline, from where we are right now which, in a lot of cases, really is the Wild West?

Tom Longstaff: That's the core of our vision, the absolute core of our vision, Rotem. We-- today, we have a lot of individual specific tools in machine-learning/artificial intelligence, and more being created every day. At Carnegie Mellon, as you know, we are the premier AI institute in the world, which means that there are ideas and things being generated on a daily basis.

Rotem Guttman: It's sometimes hard to keep up even.

Tom Longstaff: It is really hard to keep up. But what we don't really have is we don't have a repeatable way of taking those capabilities and translating the vision of how they might be used into something that can effectively go into the value chain of DoD, that can actually seamlessly work within the standard operating procedure of the operator, that's actually going to function in there. That's the key to AI engineering. We want to be able to both use artificial intelligence capabilities in the creation of software and create capabilities in the software that ultimately deliver this machine learning AI capability directly to the war fighter rapidly, repeatably, with feedback loops that come back into to SecDevOps, and understand exactly where we're making the changes that need to be changed, and then keep this cycle, Rotem, faster and faster and faster.

Rotem Guttman: So, what's the-- where do we stand today?

Tom Longstaff: Today, we have a lot of really good standard software engineering practices. We have DevSecOps, as I've talked about before, where we can rapidly develop software iteratively with security characteristics in place so that we can trust the results. That's been ten years of labor to get that involved so that we can actually do that within the department of defense. And now, we have great examples of that taking off within DoD.

To incorporate AI, that means incorporating a couple of things in addition to what we already know about software engineering. We need to incorporate the use of operational data all the way back in the design and development phases of Agile and DevOps so that that data and the data curation, the data manipulation, can be a part of how we develop finished product. There's a whole world of how do you create validation and verification of a tool that doesn't have necessarily a specifically defined outcome. The outcome is dependent on the data you run up against.

Rotem Guttman: And if you feed in bad data, you're going to get your bad results.

Tom Longstaff: You're going to get interesting results, not necessarily even bad results. I feed in bad data, if I have the right kind of system that's been V and V'd, I might still get good results that actually come out the other end. This is part of the nature of what a machine learning system can do is adapt even to data that's being manipulated if you understand how to do V and V against these systems properly. Nobody in the world has the answer on this today. We, at the SEI, are really trying to help-- our vision is to help develop the V and V so that we can trust the systems that are actually going to be fielded.

Rotem Guttman: And that can be especially challenging, at least in my experience, when the data that you're feeding in-- you might not even have full control over that.

Tom Longstaff: In fact, your adversary has control over it much-- many times. Think about it this way. You're going out into the field. You're using imagery. You're taking pictures. What are you taking pictures of? You're taking pictures of things the adversary's controlling, their movement. They're putting up whatever they want to put up. They might put up labels that cause your AI to think it's something else. They might do all kinds of things to try and manipulate the input that's going to be coming into our ML enabled systems, but with proper trust, with proper validation, with the idea of understanding what all these different ways of manipulating the system might be, we can create robust, survivable systems that will basically help us no matter what the adversary does.

Rotem Guttman: It seems like that is the path that we kind of have to follow--

Tom Longstaff: It is.

Rotem Guttman: If we're going to be successful in supporting the war fighter. And so, thank you for taking the time to walk us through that. Tom, thank you for joining us. For more information, please click on the links in the description below.

Related Resources

AI at the SEI: https://resources.sei.cmu.edu/asset_files/FactSheet/2019_010_001_539538.pdf
AI Engineering: https://www.sei.cmu.edu/research-capabilities/artificial-intelligence/index.cfm