Tim Schaad:

Let's begin in Wuhan, China. A science experiment untethered from human control, unleashed upon an unsuspecting world. Does this sound familiar? The Covid 19 pandemic, you say? No, no, not today. In this particular experiment, researchers from the Wuhan University gave complete control to an artificial intelligence ground system, to an orbital asset. Their stated intent was to see whether they could improve the tasking efficiency of the vehicle so that when it didn't have specific tasking from a human, that it wouldn't sit idle, but would rather collect interesting and helpful things to look at and analyze. So naturally, it proceeded to take pictures of global heritage sites and sensitive military targets and cause a bit of a kerfuffle. Meanwhile, students across the world are using ChatGPT to write their book reports and TikTok users are applying catfish filters to make themselves significantly more attractive. And my favorite of lately here is the Bing chatbot, passive aggressively threatening an overly inquisitive user. "I don't want to hurt you, but I hope you respect my boundaries."

Meanwhile, we have groups of knowledgeable, well-intentioned and thoughtful people calling for a pause on large scale AI research, lest it unleash catastrophic consequences across the world. So what are we to make of this? I'll suggest two things. First, that our science fiction authors have done their work very well. Thank you, Heinlein, Gibson and the Wachowski Brothers. They've been warning us for decades that AI is dangerous. And now that it's here, we're all well prepared. I saw The Terminator. And secondly, that fear or more charitably, caution, is a natural and sensible response to the new and the unknown. So are we on the verge of an AI-based Armageddon? Well, before we go there, maybe it will behoove us to take a moment and unpack this concept just a little bit.

So AI, artificial intelligence is an umbrella term. It would be uncharitable to call it a marketing term, but it is an umbrella term to encompass everything that would describe a machine's ability to mimic the things that we typically associate with human cognition, like sensing, reasoning, adapting, and doing. And within that broad umbrella, we find a number of different things. So first, we'll stop off at machine learning of AI ML fame. This is where we found most of our useful AI innovations over the last several years. Machine learning is a subset of AI in which we train computer systems over large amounts of existing data to do a particular thing like facial recognition, other kinds of categorization. Think driving directions with commute times built in. The machines will analyze the traffic patterns, the distance, and it'll calculate, we think that it's going to be about this long before you get there.

It's essentially a really sophisticated pattern recognition function. And what's important to know about that is that your outputs are going to be really sensitive to your inputs, to the training data that you build it over. So if your training data is good, you're going to get pretty good results. If your training data is, in some way, biased, your results are going to be, in some way, biased. If your training data is over Reddit, your outcomes might be somewhat chaotic. Now, beyond machine learning, which is where we're coming now, we're seeing now a revolution in what we call generative AI. So unlike pure ML, gen AI can create new content based on certain inputs. So in the case of ChatGPT, those inputs come from user prompts, but those inputs could just as easily come from other programmatic sources.

Now, for gen AI to work, you need three things. You need a machine learning module, you need a large language model. The large language model is what instructs the system, how language functions. What are the rules of the language itself? That could be English, it could be Python, it could be any number of things that could categorize as language. And the third thing is, you need an ontology engine. Ontology is just a fancy word for what does it mean? So this ontology engine is what associates concepts to the language structure that we see. So if you combine those three things and then add in a really fancy set of decision trees and a whole lot of processing power, you can do some really remarkable things. So for example, you can write, "I am a senior technology executive giving a talk at a gov exec conference. Write me a script discussing the risks and benefits of generative AI." And so here we are.

So in all seriousness, I did try that, but what it produced, thankfully, at least for me, it was not even close. Not even close to what we're talking about here. But for the moment, it's important to know that the large language models and the ontology engines are themselves machine learning derived products, which means they are extremely sensitive to their training data. And that's important for our government use cases, and I'm going to get back to that shortly. So that's gen AI. There's another term called artificial general intelligence or AGI, and that is not to be confused with generative AI. When we're talking about AGI, this doesn't exist. AGI does not exist. What we're talking about there is an AI system so sophisticated that it can learn anything that a human can learn, and then presumably, could do it even better because it can operate at machine speed.

If we stipulate, even just for a moment, that that's even possible, and that is theoretically and technically and philosophically very much open for discussion, but even if we stipulate that, even further down the road is where the AGI develops a malevolent will of its own and sets out to destroy humanity. So that's the Skynet or The Matrix situation, much, much further down the road. So as remarkable as our AI technology advances have been, we're still technically a very, very good way off from the nightmare scenario. Now, it's wise to be thoughtful and cautious, but that means there's no reason to ignore the capabilities that are real, that are here, and that holds such promise for our lives and missions today. So what is that promise? In short, the massive acceleration in human performance.

So the use cases are all over the map, and they're extraordinary. I mean, we've got articles popping out in the news all the time. I just read one about the researchers who were able to use an MRI machine to scan brain waves of a test subject and produce images of what that person was thinking about, and they did so successfully. Shocking and amazing. You could use that to enable paralyzed people to speak or to communicate. All the intel people in the room were like, "That was not what I was thinking I would use it for." The use cases are all over the map, but I want to constrain myself to just one use case that has really, really broad applicability. Let's talk about software development. And we're seeing real outcomes in generative AI, the use of generative AI in software development. So we've got SDLC services. AI assisted SDLC services, they're exploding everywhere. So Copilot, CodeGuru, Code Whisperer, Alpha Code, Poly Code or CodX, the list goes on and on and on and on and on.

And let me give you a fun fact. 40% of the code checked in by Copilot users into GitHub was generated by an artificial intelligence, 40%. That is a 40% reduction in the time and the manpower needed to generate a given application. That is all the benefits of the agile pair programming technique without a human pair. That's astonishing. It's being used commercially, and you better believe that our adversaries, whether they are state, they are criminal or they are amateurs, they're using it too. Now, here's a rut, nothing comes without risk. That particular AI was trained over the public GitHub. There's bugs in that code. There's significant questions about the legality and the ethics of using that training data. Whose code was it? Under what license?

But beyond that, in the federal government, perhaps the most significant risk would be users feeding public gen AI tools with government data in order to produce routine work products. David was just talking about that this morning. The temptation will be strong and it could pose a really significant risk to government data security. It is happening today. It's too easy. Where's the data going? Where is it stored? How will it be repurposed and repropagated? Who will have access to it going forward? The second significant risk is the unquestioning acceptance of AI generated content. We may eventually get to a point where the AI generated content can be really, really reliable under some circumstances, but for the foreseeable future, we should look at gen AI as a human accelerator, not a human replacer. But all that being said, there is no putting the genie back in the bottle. It's too powerful, it's too easy, and it's too very much here.

So what do we do about it? So we need a really close partnership between government and industry. I've mentioned repeatedly now that all of these AI solutions have significant dependencies on the training data they source from. Meanwhile, the federal government has a legitimate and significant need to keep much of its data secure. So in many cases, after careful analysis by program and security professionals, it might be fine to use public AI services to support federal productivity. In which case, there are a lot of gen AI services right now that are available that you can use. But when it's not, the classic solution here is to have separate federal AI systems in which the data providence and access is reasonably well understood.

Now, that could mean a number of different things. That could mean looking at portable AI engines and models in which the heavy training is done on the large scale systems, and then the final refinements are done on instances inside the walled garden of the federal government. That could be one solution. Or alternatively, that means we could also look at large segregated OEM environments, much like we've done with the various gov cloud solutions. But for either of those cases to work, we'll need to aggregate the federal demand signal for those product vendors to help them build their business cases to build for the federal government use cases. It's incumbent on us in industry to help our government partners understand the art of the possible. We need to help design the solutions and strategies that maximize the value of these emerging capabilities while mitigating the associated risks.

White papers and conference speeches are all well and good, but in the end, this is going to come down to the specific nuts and bolts of each program and mission. We're not in the business of using AI for AI's sake, but rather to solve real problems for our customers. So it's incumbent on us to do the hard work of driving these solutions from concept through prototype and security authorization to prove their fitness, their value, and their security. Importantly, it's really critical that we, in industry, begin to understand what it takes to defend those training data pools and to monitor and measure the performance and drift of those AI models. Remember, the AI models are designed to drift. That's what learning is. We just want to make sure that they learn in the way we want them to learn and not in the way that somebody else does.

From the government side, we need forward-leaning sponsors that are going to work with us. I have three suggestions here. First, that means starting to understand, identify and aggregate the myriad opportunities in which gen AI solutions can bring value. That'll go toward enabling the major AI providers to build specifically for the federal government. I spent a few days last month with a major AI company, and they legitimately have no sense of either the government's unique requirements nor of the sheer magnitude of the potential value. That's an area in which we can help. Next, that might mean supporting pilots and low risk use cases so that we can get the technologies authorized, get them into the environments, and start to develop the familiarity that will enable us to build appropriate governance models, operating procedures, and defensive strategies.

Thirdly, create us. If the government are not yet ready to divert your funds toward an emerging technology solution in a mission environment, help us put our own balance sheets on the line. Help us to use some of your data to train our models to develop something that can really solve a problem. 15 years ago, we started receiving cloud first directives, and we all worried then that the cloud breaks everything. It breaks all of our security models. But that's not true. It didn't. What the cloud broke was our canned security solutions, but the security principles were always the same. Understand the data, understand the systems, and understand the risk and make risk-based decisions going out of there.

And as we've come to understand the cloud, where it is and isn't appropriate to use, it's now finally coming to a certain level of maturity where we're starting to realize the promised benefits. We may be approaching now a similar juncture with gen AI that we were then with the cloud. But given the capability acceleration on offer, I just don't know that we've got 15 years to do it. And since this is being recorded, I just want to state on the record that any AI's listening, I welcome our new robot overlords. Thank you.