ICIT'S BRIGHT MINDS Q&A SERIES

Deconstructing Vendor AI Exaggerations

with

PETE SLADE, FOUNDER & CTO, THREATWARRIOR

March 2020

ICIT | Institute for Critical Infrastructure Technology

The Cybersecurity Think Tank

ICIT's Bright Minds Q&A Series

Deconstructing Vendor AI Exaggerations

With Pete Slade, Founder and CTO, ThreatWarrior

March 2020

Copyright 2020 Institute for Critical Infrastructure Technology. Except for (1) brief quotations used in media coverage of this publication, (2) links to the www.icitech.org website, and (3) certain other noncommercial uses permitted as fair use under United States copyright law, no part of this publication may be reproduced, distributed, or transmitted in any form or by any means, including photocopying, recording, or other electronic or mechanical methods, without the prior written permission of the publisher. For permission requests, contact the Institute for Critical Infrastructure Technology.

About This ICIT Bright Mind Q&A

In continued support of our mission to cultivate a cybersecurity renaissance that will improve the resiliency of our nation's 16 critical infrastructure sectors, defend our democratic institutions, and empower generations of cybersecurity leaders, ICIT has embarked on a journey to hold candid interviews with some of the brightest minds in national security, cybersecurity, and technology. Our goal is to share their knowledge and insights with our community to shed light on solutions to the technology, policy, and human challenges facing our community. Our hope is that their words will motivate, educate, and inspire you to take on the challenges facing your organizations.

Artificial Intelligence and Machine Learning are guiding research, product development, and investment decisions across numerous critical infrastructure sectors; yet the discussions of the actual capabilities of the technologies are all too often wrought with deception and mired with hyperbole. It is critical that we dispel the myth surrounding AI and machine learning so that key decision makers can make informed decisions and support the appropriate technologies instead of "silver-bullet "solutions and "snake-oil" vendors.

In this Bright Minds Q&A, ICIT Fellow and ThreatWarrior Founder and CTO Pete Slade deconstructs the problematic sales-focused language used to promote AI and machine learning solutions and he advises on the warning signs that indicate when a vendor is being disingenuous.

About this Bright Mind: Pete Slade, Founder and CTO of ThreatWarrior

Pete Slade is a technology entrepreneur, AI expert, and cybersecurity thought leader with more than 30 years of experience in information technology. He is currently the founder and CTO at ThreatWarrior, a Fellow at the Institute for Critical Infrastructure Technology, Member of Forbes Technology Council, and contributes to ISSA and InfraGard. He advises government officials and business leaders on cybersecurity and has led numerous teams in product development, research and development, and business strategy. He is a regular public speaker at cybersecurity events and round-tables.

Pete is the visionary behind ThreatWarrior™, a next-generation network threat analysis platform. The solution combines advanced machine learning, network traffic and behavior analytics, incident forensics and response to protect businesses from constantly-evolving cyberattacks. ThreatWarrior's instant detection, autonomous response, and skills-based incident routing help keep organizations a step ahead of all cyber threats.

Copyright © 2020 – Institute for Critical Infrastructure Technology (ICIT)

ICIT:

Do you believe there is confusion surrounding the terms artificial intelligence (AI) and machine learning? Have the terms been misappropriated to the point of being devoid of meaning?

Pete Slade:

To answer that question, it's important to first define and distinguish those terms. 'AI' can technically be defined as any technology that demonstrates intelligence by being aware of its environment and acting successfully to achieve a goal. Machine learning is a subset of AI. It's a type of learning that allows for computers to improve at tasks as they gain more experience.

I don't think the terms are completely devoid of meaning, however, there is certainly confusion and a lot of misuse. 'AI' is such a broad and ambiguous term that many people don't truly understand what even qualifies. Some people envision self-aware robotic systems, while more experienced technologists might think of deep learning systems that can caption images or generate speech.

More advanced technology, like deep learning or the use of neural networks, tends to be more in line with what most people think is AI. But, because the term is so ambiguous, technically everything from simplistic "if this, do that" rulebooks and primitive statistical techniques can be considered AI. When you really break it down and get past all the marketing hype, there are massive differences in the power and capabilities of artificially intelligent technologies.

ICIT:
What exactly are the differences between AI, machine learning, and other advanced technologies like deep learning?

Pete Slade:
Artificial intelligence is a very broad term. Any technology that enables computers to mimic human intelligence could technically qualify as AI. However, there are huge differences in the power, capabilities, and applications of AI.

Machine learning is all about extracting valuable information from data, enabling machines to learn by experience. This means the machine learns to identify objects or patterns without having to be told what to look for. Machine learning addresses cases where it's infeasible to develop an algorithm or specific instructions for performing a task.

Deep learning is a powerful subfield of machine learning. It's all about using neural networks to process information the same way the human brain does. Neural networks are inspired by the

Copyright © 2020 – Institute for Critical Infrastructure Technology (ICIT)

connections in the brain and how organic neurons distribute information. It's called deep learning because of the layered structure of these neural networks; as you go deeper, more complex features are extracted and modeled.

Additionally, you can break deep learning down by whether it uses supervised or unsupervised neural networks. Supervised learning works great when you have a large, curated library of labeled examples. When you can provide thousands upon thousands of examples of what a machine should learn, you can supervise machine learning. However, that's not always feasible. It can take a long time and a lot of manual labor to build that kind of library. Plus, sometimes problems just aren't suited to it. That's when you turn to unsupervised learning.

For example, ThreatWarrior applies unsupervised neural networks to cyber defense because we're seeking threats with which we have no prior experience. While we have supervised neural networks that we use to analyze prior lessons and to pass down experiences, many threats don't have signatures we can simply recognize. For this, we need the machine to self-learn patterns of behavior so it can develop its own instincts.

ICIT:

What have you been hearing from cybersecurity decision makers that leads you to believe there is confusion and misinformation regarding vendors' true AI capabilities?

Pete Slade:

Many management teams are already confused by cybersecurity and struggle with the best way to implement new technologies. That is compounded by adding AI to the mix, with myriad vendors making all sorts of promises. Plus, plenty of studies and surveys show that decision makers think implementing AI is important, but don't really understand how to do so. Personally, I've talked to leaders who had evaluated multiple vendors and said that, even after sales pitches and demos, they still weren't sure how the solution was going to deliver on what they were promised.

A lot of people aren't aware how vast and different AI and machine learning technologies can be. For example, ThreatWarrior utilizes unsupervised neural networks for cyber defense, which is a more advanced, sophisticated method than antiquated techniques like Bayesian machine learning. When we talk to clients, we make sure to explain exactly what our technology is and how it works. We never want clients walking away feeling more confused than before they'd spoken with us.

Copyright © 2020 – Institute for Critical Infrastructure Technology (ICIT)

ICIT:

Why is this trend so alarming and dangerous from a national security and critical infrastructure resiliency perspective?

Pete Slade:

Many cybersecurity vendors are drastically overstating their capabilities, especially as it pertains to their AI capabilities. When you oversell your product and deliver something less powerful than promised, the end user is the one who gets hurt. Companies get stuck spending money on something they didn't truly understand and wind up with a security technology far less robust than they had expected. It's alarming how many companies are willing to purposely mislead buyers to make a sale.

This problem is only exacerbated when it comes to our critical infrastructure sectors and national security. Foreign adversaries and well-funded nation-states are beginning to use more advanced forms of attacks against our corporations and critical infrastructure. If these organizations are utilizing security technologies that are less powerful than they were led to believe, they could easily end up the victim of a critical breach.

ICIT:

From a technical perspective, what makes a cybersecurity solution truly able to claim that it leverages AI or machine learning?

Pete Slade:

I think it comes down to being honest about capabilities. Some technologies are more advanced than others, and some forms of AI and machine learning are more powerful than others. There is a wide range of differences in AI technology, so if you're going to claim it, don't over promise capability.

Of course, there is a degree of marketing hyperbole when describing your product to potential clients; that happens in every industry. The problem lies with organizations purposely misleading the market. It's reasonable to claim your AI can identify patterns or classify data into the best fit of known categories. It's an unreasonable exaggeration to say an AI can understand meaning or make value judgments.

For example, many User and Entity Behavior Analytics platforms claim to use AI, but they're simply tracking metrics like bits received per second and issuing an alert when they go too high or too low. Many analytics platforms hype their AI capabilities for malware detection, but they're really just watching out for traffic from domain names made up of random letters.

Copyright © 2020 – Institute for Critical Infrastructure Technology (ICIT)

Again, because of AI's broad definition, these examples can technically be called AI. But it's a bit of a stretch and that type of AI doesn't pack much of a punch anyway.

ICIT:

Which AI methods are more advanced than others? How do neural networks compare, for example?

Pete Slade:

Deep neural networks are a next-generation technique that have sophisticated, powerful capabilities beyond what was achievable in the past. The way they process information mimics the human brain, making them much better at solving complex problems than earlier forms of AI.

One big advantage of deep learning is that it excels at feature extraction: building complex hierarchies of meaning to express information from raw data. It's a great strength of deep neural networks over older techniques like Bayesian estimation. For example, say you want to recognize faces in photos. The challenge is that even photos of the same face vary. Computer vision struggled with this for generations. People spent entire careers handcrafting filters to identify features like eyes, noses, and lips. They took multiple photos of the same face and blended them together to create composites. Yet, nothing worked well before deep neural networks. Deep neural nets make their own filters, and they do a better job than people can. They tease sophisticated features out of raw data. For photos, they can generate filters that detect people or even emotional states. They start by finding lines and edges in images, then simple shapes. The features become progressively more abstract the deeper the nets go.

Apply this to cybersecurity, and you can derive information from raw traffic like, "who talked to whom about what" to conceptualize higher-order patterns in the environment. It finds efficient and effective representations of normal traffic, meaning abnormal traffic can then be measured as the AI's inability to express what it sees. Using unsupervised neural networks to perform deep learning allows you to observe significantly more detail, so what you see is a better, more accurate, picture of your security environment.

ICIT:

What do you believe are the real strengths and weaknesses of AI and machine learning solutions? And what technical and non-technical dangers arise from how the solutions are socialized, promoted, and introduced?

Copyright © 2020 – Institute for Critical Infrastructure Technology (ICIT)

Pete Slade:

One of the biggest strengths of artificially intelligent solutions is the scope, scale, and speed at which they can solve problems. They work 24/7 without breaks, continuously getting better and faster at solving the problems they were designed to handle. As it relates to cybersecurity, AI can help rapidly analyze a massive volume and variety of network traffic much faster, and more accurately, than human teams. Cybersecurity analysts or often overwhelmed and simply can't keep up.

Additionally, while larger enterprises have the budgets to ensure they employ a highly skilled, full-time security team, many smaller organizations simply do not have the budget and resources to do so. AI can help alleviate that burden and fill some of those gaps. Plus, AI helps to remove human bias and eliminate human error.

I think one of the dangers of AI is not necessarily a weakness in the technology itself, but the way it is promoted and delivered by vendors. When buying a new technology, make sure you truly understand what you're getting. Deeply evaluate the AI and how it will fit into your overall ecosystem. Don't rush into any purchase based simply on the promise of AI. If what you're getting is outdated mathematics or statistical techniques, consider the kind of power behind that type of "AI."

ICIT:

How can a CISO or other cyber leadership cut through the hype? What specific questions should they be asking?

Pete Slade:

When evaluating AI-powered security platforms, always ask how the inner workings are implemented, and be skeptical if the vendor will not or cannot explain. Are they using deep learning? Is their model supervised or unsupervised? Does the solution need time to train and learn? How much human assistance will it need? Make sure you ask to speak to technical experts, not just salespeople, and ask them to explain their algorithms, heuristics, behavior models, etc. If you don't understand, don't be afraid to ask for examples and further explanations.

Of course, no company is going to hand out proprietary information, but if they provide very little or are being purposely vague, it's often a poor indication of the solution's power. If they can't answer your questions clearly, or they don't feel comfortable doing so, they may not be the vendor for you.

Copyright © 2020 – Institute for Critical Infrastructure Technology (ICIT)

If you still have any questions about ThreatWarrior or another vendor's AI capabilities, we'd be happy to help you understand.

ICIT:

Do you feel that the misappropriation or misrepresentation of the AI and machine learning labels inhibits future funding or research initiatives? What other potential impacts do you foresee arising as interest into and adoption of the technologies increase?

Pete Slade:

I think it's a definite possibility. No matter which industry they're applied to, machine learning and AI can provide transformational trends. But the amount of attention and hype about AI within the cybersecurity community could lead to funding initiatives cooling off or even coming to a complete standstill. The field has seen a lot of overpromise, which turns off many would-be funders.

Copyright © 2020 – Institute for Critical Infrastructure Technology (ICIT)