COSMOS
Navigate the Compliance Universe

Report on Patient Privacy Volume 18, Number 6. June 30, 2018 Artificial Intelligence Can Help Monitor Workers, Reduce Threats

By HCCA Staff

Systems enabled by artificial intelligence (AI) are showing some promise in learning to detect insider threats, but the systems still need to be augmented with a robust organization-wide insider threat policy and human oversight, experts say.

Dozens of vendors now market solutions intended to detect insider threats, including many that claim they are "AI-enabled." AI and machine learning are relatively new but are making good progress in detecting anomalies, especially if they're given larger data sets with which to work. Eventually, they'll teach themselves to recognize threats—both from insiders and from outsiders—before much damage has been done.

Still, security personnel can't turn over the task of insider threat detection completely to computers just yet, says Michelle O'Neill, director of corporate compliance, Summit Health Management in New Jersey. "Human vigilance is still important, which is why we speak to watching for employee behaviors," O'Neill tells *RPP*. "In addition, although artificial intelligence may help point out 'red flags,' those red flags do have to be investigated in person."

The idea behind AI and machine learning is to have the system gather enough information about "normal" behavior so that AI-enabled solutions can compare peoples' behavior to themselves (to detect anomalies over time) and to others (to detect anomalies between similar workers). Then, the AI system would need to determine whether there's really a threat or if something nonthreatening is going on.

"AI surfaces the insider threats hidden by all the noise," says Stephan Jou, chief technology officer of Interset Software, Inc., which offers AI-enabled security software. "You can take advantage of all the data and all the systems you already have deployed."

'Fun' a Motive for Snooping

Insider threats are a huge problem for covered entities and business associates. According to Verizon Enterprise's 2018 Data Breach Investigations Report, the health care vertical is the only industry vertical that has more internal actors behind breaches than external. "This somewhat bleak finding is linked closely to the fact that there is a large amount of both errors and employee misuse in this vertical," the report says.

Misuse takes the form of privilege abuse—using logical access to assets, often databases, without having a legitimate medical or business need to do so—in 74% of cases, the report says.

"Interestingly, the motive (when known) is most often (47%) that of 'fun or curiosity.' Examples of this are when an employee sees that their date from last weekend just came in for a checkup, or a celebrity visits the hospital and curiosity gets the better of common sense," the Verizon report says. "Not to be forgotten, our faithful friend avarice is still alive and well, with financial gain being the motivation in 40% of internal misuse breaches."

The challenge in detecting insider threats "isn't that we don't have the data," Jou said in a recent webinar. "We

have plenty of data. The challenge is the fact that we don't have enough humans" to look at all the data. The current tools to detect insider threats have limitations, Jou said: they're reactive, scattered, and they have a high —60% to 80%—false-positive rate.

Machine learning takes two paths, Jou said: unsupervised learning, where the AI system is left to self-discover patterns, and supervised learning, where the system is provided with examples of certain patterns and tasked with finding more.

For AI to identify insider threats, it first has to understand the "unique normal" for each worker who uses a system, which includes how the worker uses the system plus when and where he or she uses the system. Then, the AI can flag situations that are outside that "unique normal"—for example, if a particular user logs in from outside the country, or at an unusual time, or accesses different files than normal.

"These models can be trained on historical data to learn and detect something that is out of normal," Jou said. "The idea here is you do not want to look at every one of these clues in isolation."