29 May 2019
The Australian, Australia

Author: Wilson da Silva • Section: Special Report • Article Type: News Item
Audience : 94,448 • Page: 10 • Printed size: 1000.00cm² • Region: National
Market: Australia • ASR: AUD 22,161 • words: 1619 • Item ID: 1125761215

isentia.mediaportal

Licensed by Copyright Agency. You may only copy or communicate this work with a licence.

AI stands at the front line of the cyber war

WILSON DA SILVA

There's a war on in cyberspace. Every day, hundreds of attacks are mounted against computers and networks across Australia, with hackers searching for vulnerabilities, trying to access restricted systems, steal data or corrupt networks. And it's not just banks or credit card companies: our government and defence systems are also a target.

For example, the breach in February of the federal parliament's computer network, which quickly spread to those of the Liberal, Labor and National parties. Luckily, hackers were detected early, and the Australian Signals Directorate called in. Prime Minister Scott Morrison later told parliament that "a sophisticated state actor is responsible for this malicious activity" — diplomatic code for a foreign power. Experts pointed to China.

It's estimated 22 countries can launch offensive cyber operations, and there were more than 50 state-sponsored attacks in 2018, according to the Council on Foreign Relations in New York. Just this year, Chinese hackers tried to steal research on military maritime technology from 27 universities, US and European think-tank networks were breached, and Indonesia accused Chinese and Russian hackers of modifying voter databases to disrupt presidential elections.

"Ten years ago, this was almost science fiction," said Dr Gareth Parker, a research leader at the cyber and electronic warfare division of Defence Science and Technology (DST) in Adelaide. "A decade ago, people were still thinking in terms of antivirus protection."

Now, there's a phalanx of threats: data scraping, keystroke theft, distributed denial-of-service attacks, worms and trojan horses, remote port scanning, spoofing, ping floods, smurfing, phishing and even eavesdropping of optical fibre networks. As you'd expect, no-one in Australia's defence community will admit to breaches or even attempted attacks — but logic would dictate that they are as much a target as anyone else.

Which is why research into countering cyber-attacks has been booming. And why at the pointy end of the best is research being done by the world's defence agencies, Australia among them.

At the CSIRO's Data61 data research division, computer engineers developing highly sophisticated algorithms that hunt through a computer network, constantly monitoring behaviour and seeking to identify and contain an intruder.

But the engineers are not writing the software: they're letting the algorithms write themselves. It's called 'adversarial machine learning', and led by DST and Data61, with researchers at the universities of Melbourne, Swinburne and Monash, it is one example of how artificial intelligence, or AI, is playing a major part in defence research today.

"Let's say an adversary tries to infect, attack or poison a server which is located in your network," said Data61's Dr Richard Nock. "The goal of the machine learning algorithm is essentially to figure out what's happening and correct the behaviour of the network. But to do that, it can only learn by being exposed to examples and observations of intrusions."

It's the cyber equivalent of battlefield wargames: once the core machine learning software has been developed and deployed into a simulated defence computer network, scores of computer engineering students will try to hack the system, either sneaking in undetected or trying to wreak havoc and disable the network. In short, they'll try to emulate what an enemy might attempt against a defence installation or a computer system supporting military personnel in the field.

"So you would simulate the attacker, and essentially train the machine-learning algorithm against that kind of attacker," added Nock. "It's like training the immune system."

Machine learning is the AI technique that powered Google DeepMind's AlphaGo to become the first program to defeat a 9-dan world champion at the ancient Chinese game of Go in March 2016. It relies on 'reinforcement learning', which uses a neural network driven by a finely crafted algorithm that is exceptional at one thing: learning what it needs to in order to achieve its goals.

As the neural net is exposed to more and more attacks, it builds a database of incursions and teaches itself to recognise patterns that no human could possibly discern. You can't "look under the hood" and necessarily understand how it does it, just see that it does this very well.

Which is part of the problem with neural nets: how do you know if a sophisticated intruder, knowing that you have a machine learning algorithm defending your network, hasn't itself developed another machine learning algorithm that is adept at disguising itself from detection once it's in the system? Or even worse, subtly "trains" your defensive algorithm to ignore its presence over time?

"It is a bit of an artificial intelligence arms race," agreed Nock. But their project also aims to defeat such an attack strategy — using superior mathematics. "It's essentially a matter of the underlying mathematics of your training algorithm being the most robust it can be ... so it actually sees possible tampering," he added.

Also of growing concern is detecting suspect behaviour across the internet, an internal data network, or even over Wi-Fi — especially as a growing amount of data is being encrypted. That's where Project Deep Bypass comes in: developed by Data61, the Uni-

Licensed by Copyright Agency. You may only copy or communicate this work with a licence.

versity of Technology Sydney, the University of Sydney and DST, it “sniffs” high-speed network data traffic and characterises encrypted traffic.

Again, machine learning is key. Deep Bypass uses three different deep learning models and corresponding neural net architectures to filter data rapidly and search for key statistical characteristics. In effect, it “fingerprints” data and even recognises content with some accuracy. If the content has been previously flagged as being of interest to defence or intelligence agencies — such as a terrorist propaganda video or hate speech video — Deep Bypass will recognise this, and even identify which video was played, with 97 per cent accuracy.

“That was fairly impressive,” said DST’s Parker. “But that’s what you want to be able to do, to have a high-level look at the network traffic and understand what people are using it for. And, of course, identify malicious behaviour.”

DST is managing the Department of Defence’s Next Generation Technology Fund, which is spending $730 million over a decade to develop novel solutions to defence challenges, is replete with work on AI. Another project shepherded by DST and Data61, this time involving Monash and Deakin universities, is designed to use AI to find flaws in network or other software that can be exploited by an attacker.

All software has bugs — errors or flaws that produce incorrect results or makes programs behave in unintended ways. Such errors arise because software is designed and written by humans; it’s estimated that every 1000 lines of software code has at least one error or flaw. Finding and fixing bugs is difficult, and even the best programmers, using the most sophisticated debugging tools, cannot find all flaws, or anticipate how multiple software and hardware layers will interact with each other.

That’s what hackers rely on: they study tranches of software minutely, looking for flaws that can give them a backdoor into a network, or bypass access controls and obtain unauthorised privileges and manipulate a computer or server. While this is an ongoing headache for every IT manager, in a defence context it can mean the difference between life and death in combat.

“It’s quite a complex process finding vulnerabilities, and then recognising those bugs that pose a security risk,” said Parker. So the researchers are developing an automated analysis process for software using a technique known as symbolic execution, which has “game-changing potential for computer security,” he said.

The researchers are also interested in combining symbolic execution with machine learning to create a new cyber weapon: rapid threat analysis, which would allow defence engineers fighting off a cyber-attack to not just identify malicious code, but understand its ultimate mission or target. In the past, such approaches have taken a lot of time and computer grunt; combining the methods suggests this could now be done in minutes. The goal is to develop techniques that allow portable tools which can respond immediately, and autonomously — even to unknown cyber-attacks — as they occur.

Cryptography is another area. Long on the front line of protecting both military and civilian networks, it relies on scrambling communications with devilishly complex mathematical formulae that would take decades of computational time to crack. But even better is quantum cryptography, which depends on the spooky properties of quantum mechanics to make it entirely unbreakable.

Problem is, quantum cryptography works best over short distances and on secure fibre networks. At the Australian National University, however, physicists are working on techniques to allow quantum cryptography to be used on secure defence communications via satellite with a quantum-encrypted laser communications system. It is paired with “quantum memory” which can capture, and store, information encoded in laser beams sent to a satellite without reading or tampering with the data, thereby keeping its quantum cryptography state intact.

“The complexity stays on the ground, where you generate the quantum state, send that to a spacecraft which then retransmits it,” said Dr Francis Bennet, an instrument scientist at ANU’s Mount Stromlo Observatory. “You don’t actually have to trust the hardware on the spacecraft because it is unable to make a measurement of those quantum states without completely destroying them.”

Bennet is working on the laser communications component, while the quantum memory is under development by several groups, including at the ANU node of the Centre for Quantum Computation & Communication Technology, where Professor Ping Koy Lam is using an atomic spin-wave approach, while colleague Dr Matthew Sellars uses a rare earth element, called erbium, embedded in a crystal.

The goal of the machine learning algorithm is essentially to figure out what’s happening and correct the behaviour of the network