

AI: The perfect partner for the future of cyber defence

By [Alex Healing](#)

18 Oct 2019

Building a safe future for business' critical data requires a collaborative effort between machine learning capabilities and traditional human decision-making.



Alex Healing is senior researcher of future cyber defence at BT

In a climate where Internet of Things (IoT) and cloud computing are pushing global enterprises' data flows into vaster and faster environments than we've ever seen before, security teams need to find ways to concentrate their efforts on protecting the data that matters most.

Using the concerted efforts of man and machine, security teams are exploring ways to cut through the noise of these increasingly busy environments. This is a problem that is only going to become more challenging as technology moves towards the cloud and networks require more context, more logs and a more flexible environment. The Chief Information Security Officer (CISO) needs to be investing in technology that aids analysts to reduce these challenges now and readies their business for a complex future.

By training both analysts and machines to define and visualise what normal network activity looks like, businesses stand the best chance of knowing when something isn't right and taking fast, decisive action.

But the question that continues to drive the future of cyber defence is: how can we spot near-invisible anomalies amongst masses of 'normal' data, and mitigate against both the insider and outsider threats that may be manifested in them?

Man and machine: finding the balance

The majority of our networks and systems have sophisticated automation capabilities to deal with high volume, low sophistication attacks. But developing Artificial Intelligence (AI) beyond automation is about accelerating what the analyst can do in the event of a more complex threat.

Businesses, therefore, should focus on using an intelligence augmentation (IA) strategy: the practice of using machine learning and automation to complement the intelligence of the human, so they're free to make meaningful contributions at a higher, more sophisticated level.

Security teams are facing an average of 174,000 alerts per week - and are only able to review around 12,000 of them, with approximately four days to resolution. When you consider the fact that today's average global business' broad data logs show roughly ten million events per second, it's no surprise that analysts are struggling to keep pace.

Couple the sheer volume of data with the subtle and, often unknown, nature of the patterns that arise within these logs, and you have a mass of 'noise' that humans just can't cut through by themselves. Analysts aren't just looking for known threats, but also the threats they haven't seen before. How can they be expected to spot an anomaly without any preconception of its form?

Enter AI. For security teams to stand the best chance of preventing threats, they need technology capable of focusing their search parameters by condensing billions of daily data points into hundreds. More than this, they need these data points, from across several data sets, classified and correlated into meaningful patterns and significant events to be presented to the user in a visual, accessible form.

Using the time and energy saved by automating this initial process, the analyst can really excel by focusing on events of interest and making meaningful contributions to the defensive process. They can study the unusual patterns and anomalies flagged by AI and determine their value and importance. Do we investigate this further? Is it a threat? How do we mitigate it and how will the action we take affect our live business operations?

Once an analyst answers these questions and determines the threat level - if any - they can feed this new information back into the AI-enabled machine; a process that enables the machine to learn alongside the analyst. This way similar threats are picked up faster and with a more specific classification.

The loop between analyst and AI is closed - the human and the machine work together to learn, develop and improve our cyber defences for the future.

A holistic view of security: the current state of play

Interactive visualisation is one way in which security teams can work towards incorporating this holistic loop into their security strategy. It's an effective means of increasing the bandwidth between human and machine by mutual interaction;

the machine is able to suggest its findings, whilst the human can explore these suggestions, interpret, validate and feedback new knowledge for the machine to learn from.

For example, Nexus, an AI-driven tool powered by artificial neural networks, is currently in development at BT as one of the first warning signs for analysts facing a potential attack.

Using graph analytics, Nexus consults its learned perception of a network's normal environment to flag up anomalies and plot them on an interactive graph — contextualising their behaviour by allowing direct comparisons to the 'normal' clusters of data on the same graph. Not only is this useful in the initial stages of an attack, but analysts can also use the visual data to discover and study behavioural patterns of advanced persistent threats (APTs): intruders that have remained undetected for an extended period of time.

In the case of initial attack discovery, if an analyst decides the anomaly Nexus has flagged may pose a real threat to their business' security, they can turn to another of BT's self-developed tools, Saturn. This is a powerful visual analytics environment that allows the user to visualise several diverse forms of data grouping. The security team can pull up the geographical positioning of the potential breach, pinpoint its exact locations across the network and study its behaviour within specific time parameters.

It's critical that the analyst is presented with these contextual and plotted representations of data. In cutting out the early, time-consuming stages of data analysis, security teams can focus their brainpower on answering the really valuable questions: why is this device acting strangely? How does it compare with other similar devices in the environment? How might it be impacting critical data stores?

With the AI-based cybersecurity market expected to rise from £9Bn to £26Bn by 2025, we can expect to see a whole new level of sophistication in mixed-initiative analytical tools, but our own Saturn and Nexus are some of the most promising options currently in deployment.

The dark side of transformation

As investment rises in AI for businesses' cyber defences, we should be wary of our adversaries progressing in the same vein. An increase in AI-driven attacks will require the need to fight fire with fire. This means working on AI that will hunt down AI-driven threats and deploy automated responses to mitigate them.

Just as we're seeing deep learning capabilities being used to help understand large volumes of network behaviour; we can expect to see attackers launching large-scale phishing campaigns. Rather than a group of attackers hand-crafting emails impersonating a friend or family member, AI can be used to mine individuals' or business' private data online.

We could see masses of emails being delivered - all tailored to a specific reader, and all completely automated. As attacks grow even more sophisticated, security teams may encounter advanced botnets and AI-enabled malware similar to those we've already witnessed in the media - intelligent threats that learns to avoid detection within a victim's IT estate.

Although there's certainly cause for businesses to prepare for this level of criminal-sophistication, it's important to remember that threat actors are subject to the same commercial challenges as businesses - their effort must match the reward.

Deploying AI technologies at this level still brings a mixture of reward and risk for investors on both sides of the fence; although the potential is huge, AI systems can increase the attack surface and produce new and unforeseen vulnerabilities.

The future will never be fully automated

Although the capabilities of AI are huge and yet to be fully realised, the value of the human will never truly be rivalled.

The challenge remains in being able to quantify uncertainty with a prediction. We must make sure that users of AI tools are aware that they will never be 100% accurate; they are indicators and aids but rarely definitive answers.

We must strive to automate our security processes wherever we can, but there will always be a collaboration between people and technology. The human is the innovator and the decision-maker, integral to deciphering cyber threats. AI systems provide high volume, high-quality data analysis - an equally necessary perspective. Only by creating technology that maximises the bandwidth between the two, can we hope to build a powerful team for the future.

ABOUT THE AUTHOR

Alex Healing is senior researcher of future cyber defence at BT

For more, visit: <https://www.bizcommunity.com>