Balbix®

Executive Guide to AI and Machine Learning

Contents

Introduction

Ask a group of executives what their definition of artificial intelligence (AI) is and how they distinguish it from machine learning, and you'll be sure to get a variety of responses. These may include ambitious takes ("AI must pass the Turing Test"), some useful descriptions ("AI is the ability to pivot"), and some contradictory definitions ("AI is a special type of ML"). What will usually follow is a lively discussion on this topic.

The term ***artificial intelligence*** is indeed ambitious and anchored to our perspectives of human intelligence, the Turing Test, and memories from Sci-Fi movies. Unconsciously, we compare any new AI technology to the ability of human experts. If Google's AlphaGo can defeat South Korean Master Lee Se-dol in the board game Go, surely there is something to the growing prowess of artificial intelligence. Unfortunately, there is quite a bit of confusion around exactly what AI is. This eBook will explore the differences between human intelligence and AI, specifically what we mean when we talk about artificial intelligence, machine learning, expert systems, and deep learning. Are there meaningful differences between these terms, and if so what are they?

CHAPTER 1: Human Intelligence and AI

Before taking a deep dive into the subject of AI, let's define ***intelligence*** relative to humans. A broad definition of intelligence is naturally quite complex, with many aspects still open to hot scientific and philosophical debate. But for our purposes, we can use the following definition:

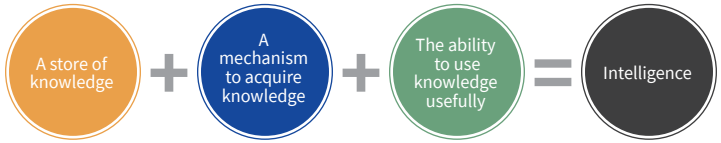

Figure 1: A simplified human intelligence algorithm

These three aspects of human intelligence should feel familiar and intuitive:

1. Having large amounts of knowledge is generally ***associated with superior intelligence***.
2. Someone who can acquire knowledge faster is generally considered brighter,
   will eventually have more knowledge, and thus be considered ***more intelligent***.
3. We appreciate people who can ***apply their knowledge*** to solving real-world problems over people who cannot and consider them to be smarter.

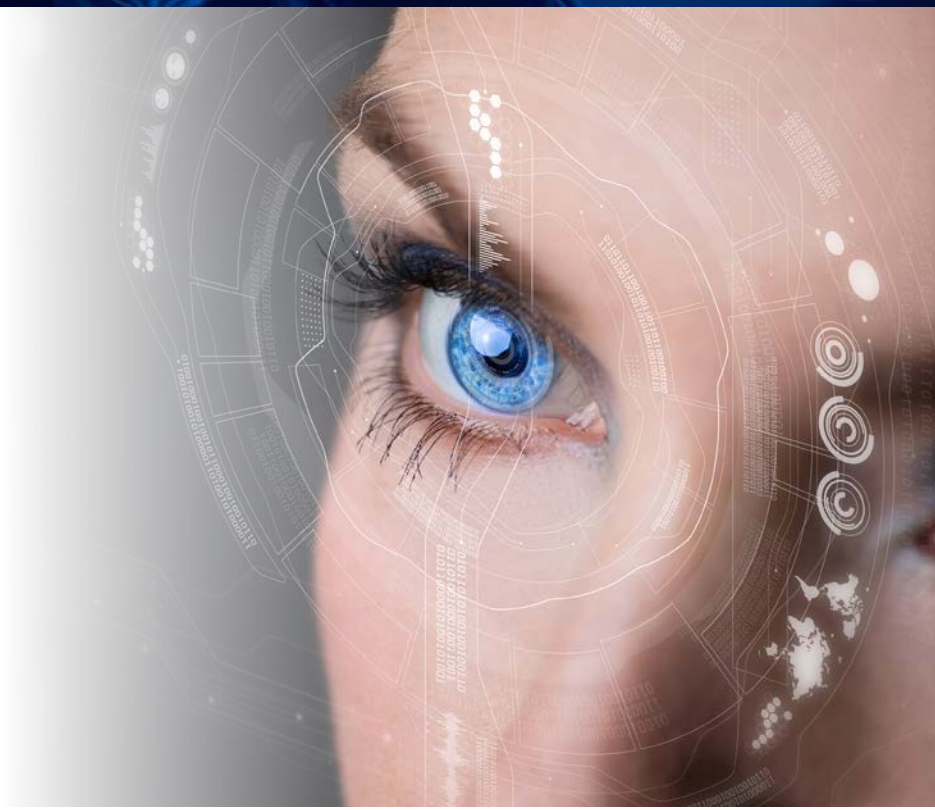

There are other characteristics we might consider as well, such as empathy, self-awareness, creativity, morality, grit, and consciousness. Sometimes referred to as Emotional Intelligence (EQ), these attributes are also consistent with the general definition offered above.

Here are two key points about human intelligence:

1

Many scientists believe that human perception and cognition ultimately stem from how the brain can discover and store ***co-related hierarchical patterns across multiple different types of sensory data***. For instance, when you see the words "JohnSmi-iPho" as part of a multicast DNS (mDNS) name in a packet capture or a log file, you intuitively know that this is quite likely the iPhone of your friend John Smith. You are unconsciously correlating knowledge that you possess about the names of your colleagues with your knowledge of common types of devices. You are also continuously and unconsciously updating both of these models in your head as you go about the business of life and are subject to sensory input from a variety of sources such as ads, TV shows, email, social media, and hallway conversations.

Contrast this with the difficulty you will encounter when trying to write a traditional program to use arbitrary substring matches to mimic this simple capability, while preserving similar flexibility on input and accuracy on output, and you begin to understand the power of the human brain.

2

Intelligence is about ***prediction*** as a method of problem solving. Yes, your eye is trying to see everything it can, but simultaneously, your brain is sending predictions down the neural hierarchy and to the eyes on ***what it expects the eye to see***. This predictive mechanism "fills in" for what you don't sense properly, and explains how you are able to walk into your bedroom without stumbling when it's pitch dark. Your brain is sending signals to the motor function of your nervous system providing your muscles with a model of what to expect as you walk around.

General AI vs. Narrow AI

The concept of artificial intelligence or AI was originally conceived in the 1950s by a few computer scientists who were beginning to think beyond traditional programming paradigms. AI pioneers were inspired by the possibility of designing super-smart programs, which would possess intelligence characteristics similar to that of humans (think R2D2 and C-3PO in *Star Wars*, or the supercomputer in *Superman III*). This is typically referred to as General AI, and it doesn't exist today. We don't quite know how to mimic the working of the human brain, or even a small fraction of its intelligence. A great book to read on this topic is Jeff Hawkins's *On Intelligence*.

What does exist today is what we might call "Narrow AI" or "Weak AI." There are numerous useful products using Narrow AI that can perform some tasks just as well as, and often even better than humans.

- An example is Amazon's Alexa, which operates within a limited input range and combines several Narrow AI techniques to perform some tasks quite well, producing the illusion of intelligence.
- The current world champions of both Chess and Go are also instances of Narrow AI.

These Narrow AI systems possess all three elements of intelligence discussed earlier – a store of domain-specific knowledge, mechanisms to acquire new knowledge, and mechanisms to put this knowledge to use.

What also exist today are several implementations of Narrow AI that solve important cyber-security problems. While we don't yet have a security bot that will pass the Turing test and replace your IT security professionals, Narrow AI-based tools can prioritize threats and vulnerabilities, and measure security posture better than most humans can.

In the next chapter, we will take a closer look at the differences between AI, machine learning, expert systems, and deep learning.

CHAPTER 2: Digging a Little Deeper

Machine learning, expert systems, neural networks, and deep learning are all examples of AI technology today. Let's look "behind the curtain" to see how these emerging AI applications work.

Machine Learning (ML)

Machine learning is a field of computer science that uses statistical techniques to give computer systems the ability to "learn" (e.g., progressively improve performance on a specific task) using data rather than being explicitly programmed.

- ML focuses on algorithms that can be said "to learn."
- Rather than writing a specific set of computer instructions to accomplish a task, the machine is "trained" using large amounts of data to give it the ability to learn how to perform a specific task.
- Training samples may be externally supplied, or they may come from a previous stage of the knowledge discovery process.

Several ML techniques have been developed over the years, including decision trees, inductive logic programming, clustering, Bayesian networks, and artificial neural networks. ML is closely related to and overlaps with computational statistics.

Note: Machine learning is a challenge because ***correlating patterns across multi-dimensional data is a hard problem***. It is very data and compute-intensive. The human brain constantly takes in an enormous amount of sensory data from a large number of sources and across numerous dimensions, and it does this over a period of many years. It slowly perfects its models before achieving the intelligence and expertise that you would associate with a skilled adult (such as members of your cyber-security team). Just try to imagine the amount of training data (labeled and unlabeled) that went into the brain of a typical college graduate. More often than not, adequate training data to feed ML systems is not available, and as a result, ML programs fail to deliver accurate results.

Machine learning is a challenge because correlating patterns across multi-dimensional data is a hard problem. It is very data and compute-intensive.

Expert Systems

Expert systems are considered separate and distinct from ML systems. Expert systems solve problems by fuzzy rules-based reasoning through carefully curated bodies of knowledge (rules). Touted as the most successful example of AI in the 1980s, expert systems derive their power from the knowledge they have rather than from the specific inference schemes that they use. In a nutshell, expert systems have knowledge, but don't quite learn on their own. They always need human programmers or operators to make them smarter. Using our definition of intelligence, they are not very intelligent.

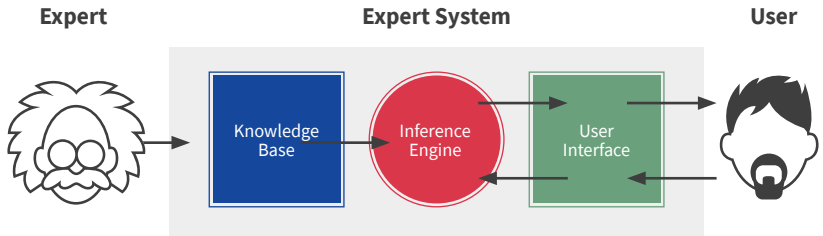

Figure 2: Expert system

Artificial Neural Networks and Deep Learning

In recent years, there have been significant advances in one type of ML called ***deep learning***. *Deep learning* (also known as deep structured learning or hierarchical learning) is part of a broader family of machine learning methods based on learning data representations, as opposed to task-specific algorithms. Deep learning has evolved from an early ML approach – artificial neural networks – which used an interconnected group of nodes similar to the vast network of neurons in the brain. In a neural network, each node assigns a weight to its input representing how correct or incorrect it is relative to the operation being performed. The final output is then determined by the sum of such weights. Practical neural networks have many layers, each of which corresponds to the various sub-tasks of the complete operation being performed by the neural network.

The output of a neural network is in the form of a "probability vector," which might say, for example, that the system is 90% confident that an image contains a given animal and 25% confident that the animal is a crocodile.

Note: Up until quite recently, the study of neural networks had produced little in the way of what you might call "intelligence." The confidence of the predictive output was low, and therefore not very useful. As you might imagine, the fundamental problem is that even the most basic neural networks are very computationally intensive and it's just not practical to build and use neural networks for any reasonably complex task. A small research group led by *Geoffrey Hinton* at the University of Toronto has continued to work on this problem, finally parallelizing the algorithms for powerful supercomputers to prove the concept.

To understand this difficulty, let's borrow an example from the field of computer vision and autonomous cars: the problem of recognizing a traffic stop sign. It is quite likely that as the stop- sign-detecting neural network is getting trained, it is coming up with a lot of incorrect answers. For example, it might be able to do a good job in good visibility conditions, but not fare very well in bad weather. This network needs a lot of training. It needs to see hundreds of thousands, perhaps even millions of images, until the weights of the various neuron inputs are tuned absolutely perfectly, and it gets the answer right practically every time, no matter what the conditions – fog or sun or rain. It's at that point that we might say the neural network "has learned" what a stop sign looks like.

This is exactly what Andrew Ng did in 2012 at Google. Ng's big breakthrough was to increase the layers and the number of neurons in the neural network, and then run massive amounts of data through the network to train it, specifically using images from 10 million YouTube videos. The "deep" in deep learning is indicative of the large number of layers in such neural networks. The Google Brain project resulted in a neural network trained using deep learning algorithms on 16,000 CPU cores. This system learned to recognize concepts, such as "cats" from watching YouTube videos, and without ever having been told what a "cat" is. The neural network "saw" correlation between visual images of a cat, visual images containing the word "cat," and audio mentions of the word "cat," and converted this correlation to knowledge, much like the way a small child might learn.

Today, image recognition via deep learning is often better than humans, with a variety of applications such as autonomous driving, identifying cancer in blood, and spotting tumors in MRI scans. There are also many variations of deep learning being actively used and improved. Some of these models can be stacked ("operated one-after-other") to result in more advanced classification capabilities. The pictures below are from demos of Amazon's *Rekognition System* that is able to recognize objects, faces, and context in images and streaming video using deep learning.

Amazon has made it quite straightforward to build fairly complex intelligent applications, particularly in the context of images and videos, using advanced ML algorithms via APIs and modules that are available in AWS. If you are interested in this, please take a look at the *AWS machine learning page* and check out the various samples and demos.

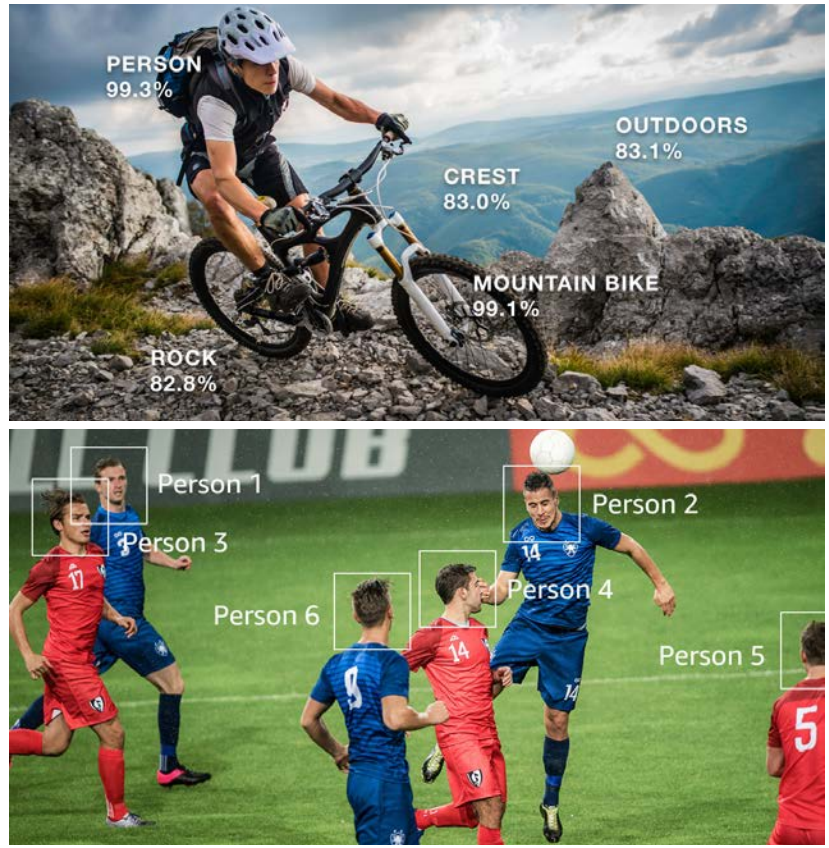

Figure 3: Amazon's *Rekognition System*

Are such systems intelligent?

Since deep learning and other advanced ML algorithms actually learn and in fact get quite knowledgeable and accurate in their capabilities within a specific domain, they do possess two of the key components of intelligence, as we defined it earlier in this eBook.

Do such systems know how to apply their knowledge to problem solving?

Narrow AI systems (aka *weak AI*), are typically focused on one narrow task. Narrow AI systems currently require human intervention to tie them into real-world problem-solving workflows, interfacing them with traditional systems and other humans. To imagine the combination of a traffic camera system that detects and tracks persons and objects (like the example from AWS Rekognition above) with other face detection and general image detection systems, requires human insight along with training image data available from the DMV's driver license and automobile license plate databases.

When such a system is installed in our public areas, it might greatly increase our communities' crime fighting abilities. Just imagine the force-multiplier effect possible here, the potential for increase in effectiveness and speed in activities that our police, security personnel, and investigators do every day. Armed with narrow AI, a group of humans charged with a particular task can become much more effective.

The Venn diagram below shows the relationship between AI, machine learning (ML), expert systems, and deep learning.

In the last chapter of this eBook, we'll discuss how intelligent algorithms based on deep learning can be used to automate decision-making and solve many important problems in cyber-security.

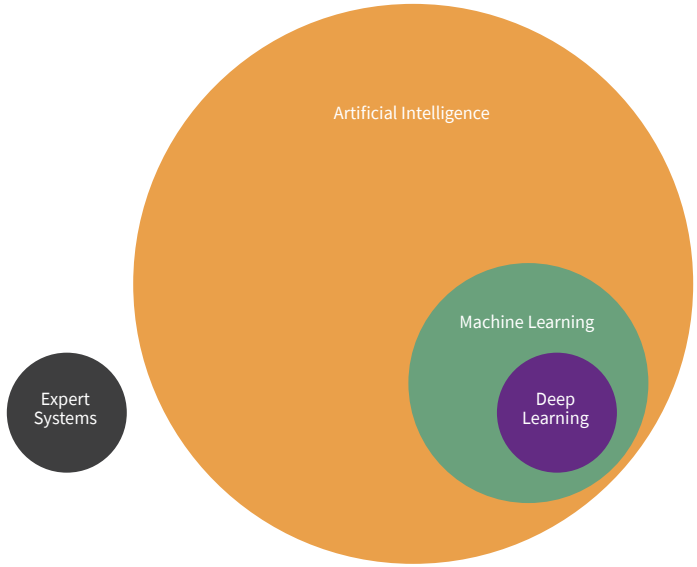

Figure 4: Key AI relationships

CHAPTER 3: AI in Cyber-Security

In Chapter 1 of this eBook, we considered a definition of intelligence and its three components: ***a store of knowledge, mechanisms to acquire knowledge, and the ability to use knowledge for problem solving***. We differentiated between the notions of ***general AI*** (which does not exist today) and ***narrow AI*** (which does). In Chapter 2, we reviewed the relationship between the terms AI, machine learning, expert systems, and deep learning, and we looked at several narrow AI systems in domains outside of cyber-security.

Now, let's focus on the use of artificial intelligence in cyber-security.

The Cyber-Security Challenge

The cyber-security challenge can be defined as maintaining the confidentiality, availability, and integrity of our computer systems. There are three major focus areas involved in cyber-defense:

1. Vulnerability assessment
2. Setup and management of effective security controls
3. Security incident handling and response

In recent years, cyber-security has become ***a hyper-dimensional problem of extreme scale***. With the "computerization" of our businesses, the number and variety of vulnerabilities has exploded. New and novel ways of compromising computer systems are discovered every day by security professionals as well as adversaries. The picture below shows an abstract view of the attack surface of a typical enterprise – 10s or 100s of thousands of places where things can go wrong (x-axis) times the 1000s of potent attack vectors and breach methods (y-axis).

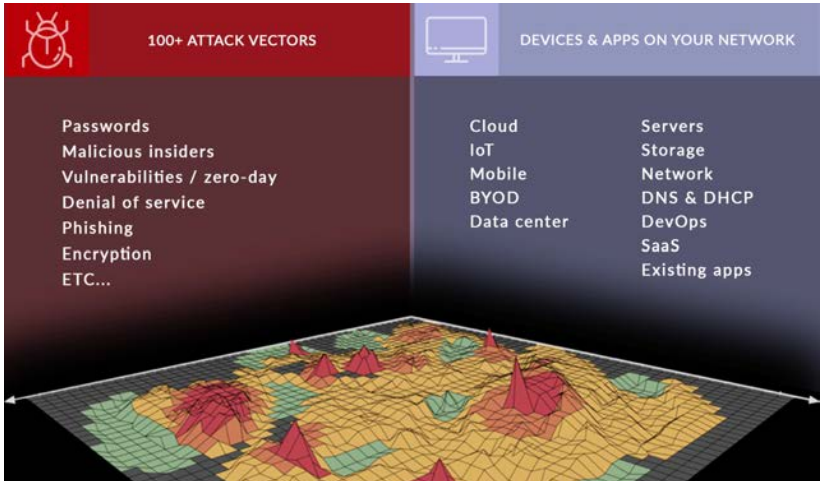

Figure 5: Cyber-security – a hyper-dimensional problem of extreme scale

Let's consider, a point on this attack surface, say ***Line of Business Apps*** (x-axis) and ***Shared Passwords*** (y-axis). Suppose an enterprise user's Yahoo or LinkedIn password is the same password used for one of the enterprise apps. If Yahoo or LinkedIn is breached, and the passwords are stolen (and not *properly salted*), then you have a problem – 1 million ways for an adversary to get in.

Generally, enterprises have no idea what this ***Password Sharing Risk Vector*** looks like for their business, so it's really important that 2-factor defenses are properly configured and working.

Attackers can exploit multiple points on this attack surface to breach your network, propagate across to their target systems, escalate their access privileges, and finally compromise, exfiltrate, or destroy your information. ***For a 10,000-person organization, we estimate over 100 million time-varying factors in the attack surface picture.*** For larger companies, we estimate that the breach risk tensor is a function of 10s of billions of time-varying signals. And that's not all. Your effective risk model is also a function of the threat model and the effectiveness of your third-party suppliers' mitigations across the attack surface.

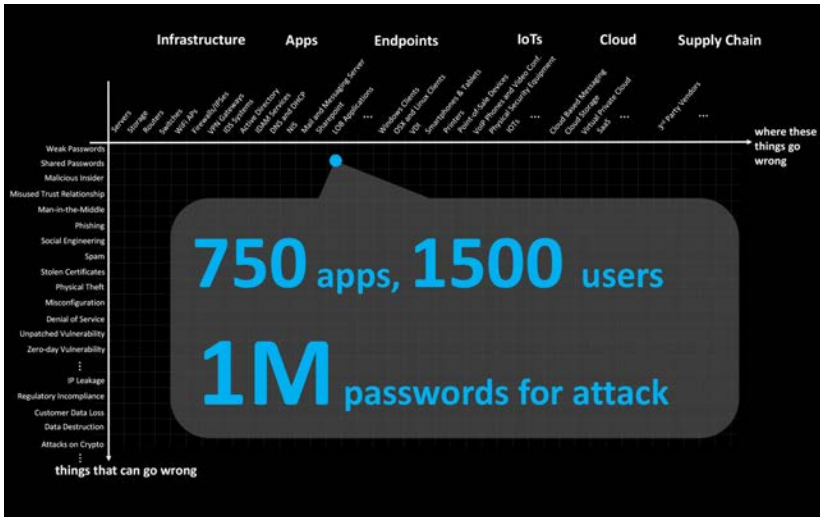

Figure 6: Attackers exploit multiple points on the attack surface

> "This is not a human scale problem anymore. There is simply too much data to analyze by hand."

For effective cyber-security, these vulnerabilities have to be discovered and addressed. This typically involves reconfiguration or patching of system(s), user training, and/or putting into place additional security products, people, and processes. Defenders struggle with the relentless pace of emerging vulnerabilities, prioritizing open vulnerabilities, and managing the large number of point solutions needed to address different areas of the attack surface.

Finally, in spite of our best efforts, attacks slip through. The number of daily security alarms that need to be handled by security operations has been growing steadily. Alarm handling involves getting data from multiple point systems, which is tedious and time-consuming. Most organizations lack the number of trained personnel needed to handle the volume of security alarms that go off daily.

AI Potential for Cyber-Security Defense

Now imagine a properly trained, self-learning system, one that is capable of autonomously and continuously gathering data from a wide variety of sources and performing correlations across hundreds of dimensions in order to surface the following categories of intelligence:

1. Deep understanding of every relevant detail (configuration, usage, etc.) of your extended enterprise inventory – all devices, users, and applications, on-premises and off.
2. Deep context around business criticality of each asset and user.
3. Up-to-date knowledge of global and industry-specific threats – aka what is fashionable with the adversary on a daily and weekly basis.
4. Intimate understanding of the various security products and processes you have deployed as part of your overall breach risk mitigation strategy.
5. Calculation of your effective risk that takes into account all the information in items 1-4 above and predicts how and where you are most likely to be breached.
6. Prescriptive insights into how you might best configure and enhance your security controls and processes to improve your cyber-resilience, without negatively impacting business operations.
7. Maximal context for prioritized and efficient handling of security alarms and incidents with impact minimization; informs tactical response to incidents, but also surface root causes and prescribes strategic mitigations for the underlying vulnerability.
8. Visualizations and reports that explain calculations and recommendations and contain relevant information for all stakeholders involved – users, business unit owners, security operations, CISO, auditors, CIO, CEO, and board members.

Conclusion

As we've seen in this eBook, AI is an emerging technology with the potential to transform our world. Some AI applications will make us more efficient, some will make us healthier, and some will keep us safer. Cyber-security presents a challenge which seems ideally suited for the power of AI, which requires the collection and analysis of masses of data and has many moving parts (attack vectors, breach methods, probabilities, risks assessments, mitigation strategies).

At ***Balbix***, our objective is to provide the world's best predictive breach avoidance platform – one that utilizes deep learning and other advanced AI algorithms to surface relevant security and risk information. To learn more about how Balbix BreachControl™ uses AI, read our eBook titled *"Balbix Uses AI to Solve the Toughest Cyber-Security Challenges."*