IBM Security

Security and Artificial Intelligence

FAQ

Jeff Crume
Doug Lhotka
Carma Austin

Security and Artificial Intelligence: FAQ

Introduction

In the world of cybersecurity, we can all agree on one thing: Change is constant. We must continuously review what we did yesterday and identify ways to improve. To keep up with our persistent adversaries, we must constantly try new technologies in an attempt to find better ways to defend or proactively prevent an attack. We must assess our policies and enhance our methodologies daily. In short, if we aren't improving, we will not be able to maintain the level of security needed to safeguard what is most important.

This foundational principle of change is well understood by every security vendor. It seems that every year the industry adopts a new buzzword geared towards selling a new technology – we're all guilty of looking for the silver bullet to our security challenges. Terms like "actionable" and "automated" have flooded our inboxes over the last few years, and now we are increasingly seeing trending buzzwords including artificial intelligence (AI), cognitive computing and machine learning.

Many vendors today claim that AI is their secret sauce, and this may, in fact, be true. This critical technological advancement seems to have blossomed overnight, but many industry professionals are skeptical of the promises being made. Sadly, the AI buzzword bandwagon has contributed to the confusion, thereby inhibiting the market from truly adopting the technology. There is no question as to whether or not AI exists today. What it is, and what it can do for us, are the real questions. Herein lies the purpose of this paper: To set aside claims and aspirations and attempt to define what AI and its components are capable of doing for us in the security space today.

Frequently Asked Questions

- **What is artificial intelligence?**
  Artificial intelligence (or AI for short) is defined by Webster's Dictionary as "a branch of computer science dealing with **the simulation of intelligent behavior in computers**."[1] The Oxford Dictionary elaborates on this basic definition, calling it "the theory and development of computer systems able to perform tasks normally requiring human intelligence, such as **visual perception, speech recognition, decision-making and translation between languages**."[2]

- **OK, but what does that mean in plain English?**
  AI essentially involves making computers more able to match or exceed human

[1] Merriam-Webster, https://www.merriam-webster.com/dictionary/artificial%20intelligence
[2] Oxford Dictionaries, https://en.oxforddictionaries.com/definition/artificial_intelligence

intelligence in its various forms, by mimicking the human ability to discover, infer and reason.

- **Isn't AI just another marketing buzzword? Is it real?**
  AI has been used to hype a lot of things that stretch the definition of the term beyond recognition. That said, some promises of AI have already been realized, while others remain in the realm of research. Instead of thinking of AI as a single, monolithic feature, it is better to envision it as a collection of related technologies that support a common goal. For example, the following is a list of some of those capabilities that, when taken together, make computers more intelligent:
  - Reasoning and problem solving
  - Knowledge representation
  - Planning
  - Learning
  - Natural language processing
  - Perception
  - Motion and manipulation
  - Social intelligence
  - Creativity
  - General intelligence[3]

- **What is machine learning?**
  Machine learning is a subfield of AI and computer science that has its roots in statistics and mathematical optimization. Machine learning covers techniques in supervised and unsupervised learning for applications in prediction, analytics and data mining.[4] Machine learning can be (and often is) used independently of other AI or cognitive technologies. In fact, this is the most prevalent type of AI we see today. Many machine learning algorithms and techniques are already in use within various solutions that look for patterns or anomalies in data.

- **What is deep learning?**
  Deep learning is a relatively new set of methods that is changing machine learning in fundamental ways. Deep learning isn't an algorithm, per se, but rather **a family of algorithms that implement deep networks with unsupervised learning.** These networks are so deep that new methods of computation, such as GPUs, are required to build them.[4]

---

[3] Wikipedia, https://en.wikipedia.org/wiki/Artificial_intelligence

[4] M. Tim Jones, "A Beginner's Guide to Artificial Intelligence, Machine Learning and Cognitive Computing," June 2017, https://www.ibm.com/developerworks/library/cc-beginner-guide-machine-learning-ai-cognitive/cc-beginner-guide-machine-learning-ai-cognitive-pdf.pdf

Security and Artificial Intelligence: FAQ

- **Deep learning sounds great. Are there any limits to what it can do today?**
  Yes. Despite the results of applying deep learning algorithms, problems exist that we have yet to solve. A recent application of deep learning to skin cancer detection found that the algorithm was more accurate than a board-certified dermatologist. But, where dermatologists could enumerate the factors that led to their diagnosis, there's no way to identify which factors a deep learning program used in its classification. **This is called deep learning's black box problem,** and presents challenges for model validation, particularly in gaining regulatory approval. The evidence shows that it's accurate and effective, yet if we cannot describe how the decisions are made, how can it be fully validated?

  Another application, called Deep Patient, was able to successfully predict disease given a patient's medical records. The application proved to be considerably better at forecasting disease than physicians – even for schizophrenia, which is notoriously difficult to predict. So, even though the models work well, no one can reach into the massive neural networks to identify why.[4]

  Deep learning has significant potential for a certain class of problem solving, but is not suitable for all situations, is difficult, costly and time consuming to implement, and works best when focused on a narrow target.

- **What is cognitive computing?**
  Cognitive computing is a subfield of AI which builds on neural networks and deep learning. It applies knowledge from cognitive science to build systems that simulate human thought processes. However, rather than focus on a singular set of technologies, **cognitive computing covers several disciplines, including machine learning, natural language processing, vision, and human-computer interaction**.[4] Of those, Cognitive computing focuses most heavily on natural language processing.

- **What is so special about cognitive computing?**
  Here's an excerpt from a white paper by IBM's Head of Research John Kelly,[5] that describes it well:

  Cognitive computing refers to systems that learn at scale, reason with purpose, and interact with humans naturally. Rather than being explicitly programmed, they learn and reason from their interactions with us and from their experiences with their environment. They are made possible by advances in a number of scientific fields over the past half-century, and are different in important ways from the information systems that preceded them. Those systems have been deterministic; cognitive systems are probabilistic. They

[5] Dr. John E. Kelly III, "Computing, Cognition and the Future of Knowing: How Humans and Machines are Forging a New Age of Understanding," October 2015, https://cra.org/crn/2016/09/computing-cognition-future-knowing-humans-machines-forging-new-age-understanding/

Security and Artificial Intelligence: FAQ

generate not just answers to numerical problems, but hypotheses, reasoned arguments and recommendations about more complex – and meaningful – bodies of data. What's more, cognitive systems can make sense of the 80 percent of the world's data that computer scientists call "unstructured." This enables them to keep pace with the volume, complexity and unpredictability of information and systems in the modern world.

- **Is cognitive computing the same as AI?**
  Cognitive computing is an AI technology, but there are other aspects of AI that don't involve cognitive.

- **How do all these AI technologies relate?**
  Deep learning is a special type of machine learning. Cognitive computing uses both. All three are examples of AI technologies. Here's a stylized Venn diagram to illustrate.

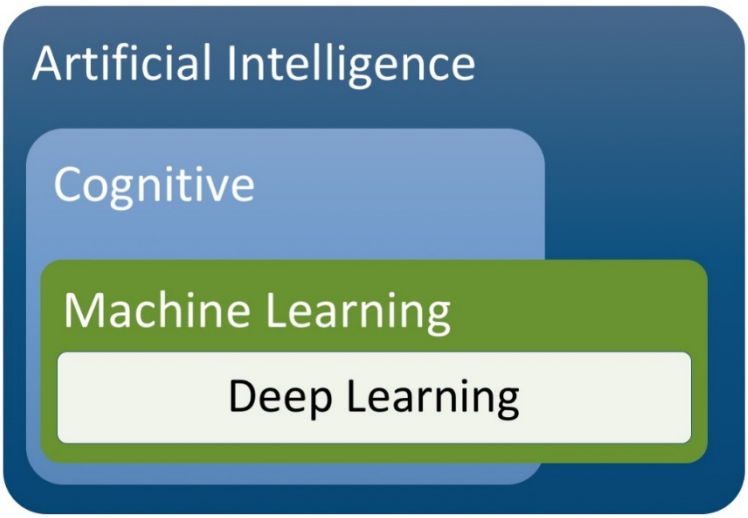

- **What is IBM Watson®?**
  Watson is IBM's AI, machine learning and cognitive computing platform. It provides a wide range of AI technologies to process both structured and unstructured information from a wide range of sources, understand what they mean, and add them to its body of knowledge (a.k.a. corpus) for subsequent use. Some components of the platform can then respond to complex questions using natural language based upon this corpus to augment human intelligence and improve insight and efficiency in critical tasks. Other components provide machine learning and deep learning capabilities for use on structured data.

  Watson Platform Solutions are not a single product, but rather a suite of integrated components that have specialized analytical capabilities that can be leveraged as individual components to build AI-enabled solutions. While all involve AI, some provide cognitive capabilities, many use machine learning or deep learning, and some use more traditional analytics.

Security and Artificial Intelligence: FAQ

- **Sounds good, but can AI and cognitive computing be used to actually solve real world problems?**
  Yes. For example, IBM Watson is being used at the University of North Carolina at Chapel Hill Cancer Center in order to identify and recommend treatment options for patients who have not responded to standard therapies. With roughly 8,000 new medical research papers published every day, it is impossible for a doctor or even a team of doctors to keep up with all the latest developments. Watson, however, can consume all this new research and then quickly apply it to improve patient outcomes. In a test of 1,000 patients, Watson's recommendations matched those of the experts 99 percent of the time. More importantly, Watson found other treatment options in about 30 percent of cases which had not been identified by the experts.[6]

  This is just one of many examples. Another is how IBM is now applying AI and cognitive technologies to the cybersecurity space, in order to allow organizations to identify threats and respond more quickly. Watson for Cyber Security has ingested over 2 billion documents in the corpus and is adding thousands more every day. It's reduced the time to analyze an incident from hours to minutes, greatly accelerating mitigation and reducing the impact to the organization.

- **Will AI make humans obsolete?**
  No. **The goal of AI is to augment human intelligence** – not replace it. There are still significant limits to what cognitive technologies can do, especially in the area of decision making, where humans are able to weigh factors that can't easily be expressed in algorithmic terms. AI is a tool. Like other tools, AI amplifies the work that people do. AI processes vast amounts of unstructured information into a coherent whole, resulting in greater efficiency and insights.

- **Will AI become sentient and take over the planet?**
  You've been watching too many sci-fi movies.

- **Is AI Big Brother?**
  Like any technology, AI can be misused. For example, Watson for Cyber Security needs to be trained with data to provide insights for identifying threat and attack scenarios. You have complete control to configure what type of data is sent to Watson for Cyber Security for analysis. We know that the bad guys are interested in abusing AI for their own ends, but we're far ahead of them at the moment. AI is just a tool that can be leveraged for good or malicious work.

---

[6] "Artificial Intelligence Positioned to Be a Game-Changer," *60 Minutes,* CBS News, October 9, 2016, https://www.cbsnews.com/amp/news/60-minutes-artificial-intelligence-charlie-rose-robot-sophia/

Security and Artificial Intelligence: FAQ

- **What are predictive analytics?**
  As its name implies, predictive analytics is a branch of advanced analytics which are used to make predictions about unknown future events. Predictive analytics extracts information from data using many techniques from data mining, statistics, modeling, machine learning and artificial intelligence to analyze current data to make predictions about the future.[7]

- **What can predictive analytics do? What are the limitations?**
  The answer here is "it depends." The predictive power of analytics derives from the quality and relevance of the data set being used, as well as the accuracy of the algorithm processing that data, and may even extend to the capabilities of the individual interpreting the insights gleaned from the data. That said, some events are easier to predict than others, and some first-of-a-kind events may never be predicted (this type of event is sometimes referred to as a "black swan").

  For instance, in a security context, a zero-day attack would be very difficult to predict, whereas a widespread malware campaign would be much easier. In other words, even though the name is "predictive," these techniques actually provide ***probabilities*** of a future occurrence, and by definition, black swan events cannot be predicted. Unfortunately, security black swans are not that uncommon. Combine this with the fact that predictive analytics for security is still quite new, so results vary, and expectations should be conservative.

- **What is threat hunting?**
  In the context of cybersecurity, threat hunting involves searching through large volumes of data to identify bad actors and threats to an organization's IT infrastructure. The goal is to prevent attacks before they happen and eliminate or minimize their effects. Threat hunting tools may ingest threat intelligence feeds, vulnerability analysis reports, risk assessments, malware analysis, HR employee records, security event data, system logs, social media feeds and more.

  While threat hunting leverages tools, much of the work is manually driven by an investigator who researches the answers to questions they develop proactively. For instance, they may choose to look for people in a particular organization with access to highly sensitive resources, who have recently expressed a negative sentiment toward the organization, as they could be potential insider threats. There are, of course, limitations to this sort of activity, since it involves a great deal of factors to try to detect anomalous activity and then predict future events. Because such activity is human

---

[7] "What Is Predictive Analytics?" Predictive Analytics Today, https://www.predictiveanalyticstoday.com/what-is-predictive-analytics/

intensive, expectations should be set appropriately. There is no magic AI technology that can find a rock in a large ocean (or swamp) of data on a consistent and reliable basis.

IBM Security and AI

- **What AI offerings does IBM have?**
  Watson for Cyber Security is an instance of IBM's cognitive computing capability focused specifically on the cybersecurity space. It consumes structured security information from threat intelligence feeds, security events and related data, as well as unstructured sources such as research papers, security blogs, websites and advisories. It then stores this as a massive corpus (i.e., knowledge base) consisting of more than 10 billion elements and refreshes its understanding at the rate of 4 million more elements each hour. In a sense, Watson for Cyber Security is like a security expert who reads the web 24x7, never forgets, formulates hypotheses about attacks based upon this highly dynamic knowledge base and gets smarter over time. Today, **Watson for Cyber Security can be accessed via QRadar Advisor with Watson.**

  QRadar Advisor is an application which extends IBM's **QRadar Security Information and Event Management** (SIEM) platform by sending indicators of compromise to Watson for Cyber Security for processing, which then returns insights into the nature and extent of cyber attacks. This app merges the security analytics in QRadar (which itself uses AI technologies like machine learning) to discover anomalies, along with the reasoning skills and knowledge base of Watson for Cyber Security (which is cognitive focused), to provide rapid analysis of the full scope and context of an offense.

  i2 Enterprise Insight Analysis can be used to perform threat hunting and investigations. **i2 QRadar Offense Investigator** integrates the QRadar SIEM capabilities with i2 in order to improve the efficiency of the investigations.

- **Where else in its security portfolio does IBM leverage AI technologies?**
  We are actively infusing AI capabilities across the portfolio. For example, **IBM QRadar** has used machine learning to identify security anomalies in the environment for many years. More recently, the **QRadar User Behavior Analytics** app leverages advanced machine learning to detect changes in user behavior and assign risk scores.

  IBM MaaS360 uses other AI technologies in the Watson family to provide insight into mobile endpoint vulnerabilities from unstructured data.

  IBM AppScan Source leverages machine learning to reduce false positives when looking for potential security vulnerabilities in source code.

Security and Artificial Intelligence: FAQ

- **Is this stuff real? Is anyone actually using cognitive capabilities to improve their security posture?**
  Yes. IBM customers using QRadar Advisor with Watson typically see an improvement in their security posture because they are able to complete investigations faster, more thoroughly and more consistently. They are also able to get through their daily backlog of events – even with a smaller, less experienced team of analysts.

  Here are some real quotes from real customers.

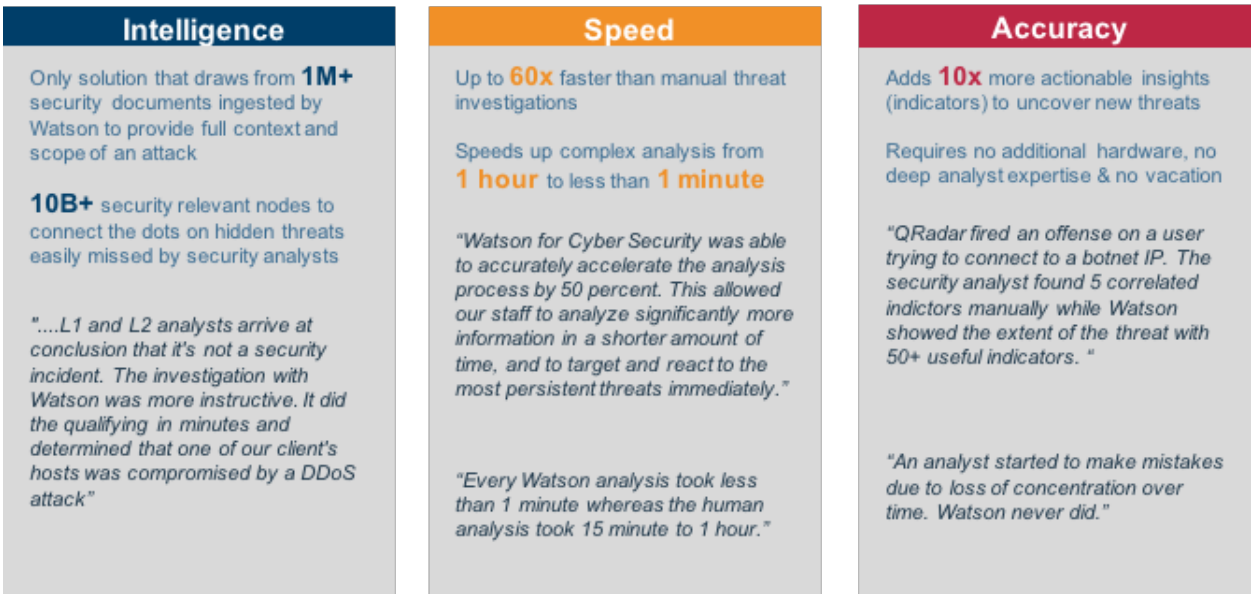

- **Can security response be automated?**
  Yes and no. For example, malware attack responses have been automated for years, with generally good results. Typically, this involves moving infected files into a special quarantine container to prevent further dissemination. Even then, there are regular reports that a signature identifies an operating system or other application component as malicious, which results in system outages.

  Well-known security expert Bruce Schneier put it this way:

  > ***Incident response is fundamentally uncertain and that makes it hard to automate.*** *For a whole lot of reasons. All attacks are different. All networks are different. Security environments are different. Organizations are different. Regulatory environments are different. In a lot of cases the political and economic considerations matter a lot more than the tech.*[8]

  Fundamentally automation requires certainty, and the one thing certain in security is variability. For specific cases where the downside risk is small (e.g., isolating one user's

[8] Bruce Schneier, "Security and Privacy in a Hyper-Connected World," InterConnect 2017.

machine), automation may be appropriate. For cases where the risk of disruption is high, such as in electrical utilities or healthcare delivery, automation should be used sparingly and with caution. Best practice at the moment is to have a human in the loop for all critical situations. For these reasons, we refer to the broader concept as orchestration, rather than automation – humans are in the decision loop, with machines rapidly executing the approved tasks.

- **Where can I get more info on all of this?**
  Here are some good references for further reading:

  IBM Cognitive Security White paper

  A Beginner's Guide to Artificial Intelligence, Machine Learning and Cognitive Computing

  5 Things You Need to Know About AI Buzzwords: Cognitive, Neural and Deep, Oh My!

Conclusion

The concept of artificial intelligence has been around since the 1950s. Its intent has always been to mimic how the human brain identifies and interprets complex patterns with the hopes of simulating neural networks that could someday begin thinking for themselves. Since the '50s we have progressed by leaps and bounds, and while we are nowhere near the classic science fiction image of a machine that thinks like a human, we are now entering the era of cognitive computing, where the interaction between computers and humans does take place.

Cognitive diverges from core machine learning and machine intelligence by having the ability to understand, reason, learn and interact with humans using natural language. Cognitive systems represent the first time a computer can provide insight into an unstructured data set that we, as humans, might not have been able to identify.

IBM Watson for Cyber Security is an AI system made up of many different subfields of AI, including natural language recognition and processing, predictive analytics, data mining, machine learning, deep learning and knowledge graphs by which to display the relationships within the data sets. This architecture works together to create a massive corpus of knowledge containing billions of data elements that were previously inaccessible to both humans and machines at scale. From every experience, the corpus gets larger and smarter, and delivers that intelligence to our finger tips.

QRadar Advisor with Watson has only begun to scratch the surface of the Watson corpus's capabilities, by creating a bridge between the world of unstructured data (Watson), deep analytics and machine learning on structured data (QRadar), and the human security analyst to accelerate the breadth, visibility and speed of an

investigation. We don't know exactly what the future holds, yet it's undeniable that the AI era is here and it promises to be a game changer for cyber security.

If you are interested in learning more about the IBM perspective on artificial intelligence, or the IBM offerings that utilize the technology, contact any one of the authors.

Jeff Crume, Distinguished Engineer, IT Security Architect, IBM Master Inventor, IBM Security Solutions Tech Sales
crume@us.ibm.com

Doug Lhotka, Executive CyberSecurity Architect, CISSP-ISSAP
dlhotka@us.ibm.com

Carma Austin, Executive Security Advisor
caaustin@us.ibm.com