The Impact of Artificial Intelligence on the Cybersecurity Industry

Bryant University Honors Program
Honors Thesis
Student's Name: Lindsey Shearstone
Faculty Advisors: Francis Varin and Alicia Lamere
Editorial Reviewer: Michele Varin
April 2023

Table of Contents

ABSTRACT

As our world becomes more digitalized, cyber criminals have an increasing landscape to launch their attacks. Developments in Artificial Intelligence are being used both to attack and defend networks, therefore, what is the next step for cybersecurity companies when it comes to beating these criminals? A study was conducted that utilizes previous literature sources written on the topic of Artificial Intelligence (AI) in the cybersecurity industry. In addition, the insights of professionals in the industry today are included through a survey and interviews to dive into the details of this battle and what lays in its future. The purpose of this study is to educate the public on the current role of AI in the industry, as it is a relatively new advancement that is slowly becoming mainstreamed. Preliminary research has shown that Artificial Intelligence may be the key to defeating these criminals, but there is much to discuss in terms of the use of this technology by cyber attackers, how it is being implemented into defense mechanisms, and the issues in the industry that may prevent this technology from growing as fast as it could be. Results from this research show that AI has a positive impact on the cybersecurity industry, but there are challenges with its implementation that have prevented it from becoming truly mainstreamed. However, as cyberattacks continue to evolve, AI will be key to winning the battle in this technology arms race.

INTRODUCTION

Today, our world is centered around technology. Whether it is due to the pandemic forcing almost all aspects of daily life to go online, or just the natural transition to more developed technologies, for better or for worse, our lives revolve around the internet and our mobile devices. With that comes an increasing area for cyber criminals to hack into our accounts and access our personal and business information. How can we protect ourselves and our assets from this severe threat? Cybersecurity is becoming more and more prevalent, as we need to protect the devices that our lives now revolve around. From Instagram accounts to business networks, there needs to be increased security to fend off these attacks that are seemingly more malicious than they have ever been, thanks to the evolving AI technology that can aid them in their attacks (ChatGPT being a recent example). Artificial Intelligence may very well be the answer to providing the security we need to keep safe what we treasure most.

The purpose of this thesis is to inform the general public about the current trends of the technological world. It can help make people aware of the advancements that are happening, so they can try to implement them for their own security in their businesses or in their homes. This research can make people conscious of the fact that even though security is advancing, so are the cyber-attacks, they should be well prepared for what is in store for the future.

This thesis aims to bring together most of the research previously done on the topic into one place, and to also add the insights of professionals currently working in this constantly changing industry.

RESEARCH QUESTIONS

Artificial Intelligence can be a game changer for the cybersecurity industry. The following questions are explored through this thesis:

- Where does AI stand in the industry today and where is it headed?
- What are the challenges associated with AI right now?
- How do bad actors utilize AI to their advantage, and what does that mean for cybersecurity?

LITERATURE REVIEW

In this review of Artificial Intelligence's use in the cybersecurity industry, there is a discussion about the current trends of the technology in the industry, and where it could head in the future. It also discusses the relevance of the skill resource gap, red team/blue team exercises, and how cyber criminals are using AI for their benefit.

Artificial Intelligence Evolution to its Current Trends

As the world becomes more computerized, cyber security is becoming increasingly crucial in businesses for the protection of their assets and information. (Hunter, 2020; Simonovich, 2021). Cyber-attacks pose a major threat to the well-being of businesses and personal privacy, so it is imperative to find the best solution for this protection, and AI seems to be the answer.

Cybersecurity is a set of technologies that are designed to protect computers, networks, programs, and data from attacks by cyber criminals (Xin et al., 2018). Most organizations employ a hybrid solution comprised of both software and hardware appliances in multiple layers to guard against attacks. These include firewalls, antivirus software, and intrusion detection systems (IDS) (Diogenes and Ozkaya, 2018; Xin et al., 2018). In the past, signature-based detection has been used as the most common antivirus software (Hall, 2021; Xin et al., 2018). These methods detect known attacks by comparing their signatures to those that have been stored in a malicious signature database. They are useful in detecting known attacks, but when it comes to identifying new ones, or zero-day attacks, it is not effective, as the software doesn't have the associated signature in its database to categorize it as a threat (Hall, 2021; Xin et al., 2018).

In Hall's article, he states some striking numbers: as of June 2020, there had been nearly 16 billion recorded breaches, and 1 in about 100 emails is a phishing attempt (Hall, 2021). This shows just how large the attack surface is, and it is only going to continue to grow. Attackers are changing their signatures, and deploying AI of their own, so the best way to conquer this is to step up the defensive game and incorporate AI into the defense mechanisms. The traditional signature-based detection, a form of reactive detection, has failed in defending networks from these evolving attacks. So, the implementation of a proactively preventative solution, that of AI and ML (machine learning), makes the necessary transition that will help with these advancing attacks (McClurg, 2020). However, until the AI and ML models have been trained, traditional analysis will provide the best bet for defense, versus an untrained AI/ML model. The network and endpoint protection that have been in place for years is becoming obsolete though, and needs to be advanced, as the hackers have learned to adapt. This is why these Artificial Intelligence models need to be trained so they can fill this role (Labs 2021).

Artificial Intelligence and Machine Learning algorithms can track different phishing sources and distinguish between valid or deceptive websites and platforms (Hall, 2021). It can be used to evaluate vulnerabilities in a network and can identify potential points of entry for hackers (Hunter, 2020). AI learns from previous human experiences, and its mathematical algorithm allows it to continuously learn from new input data (Addo et al., 2019; Addo et al., 2020). This ability of an automated system that continuously learns new threats on its own by analyzing the behavior of the malicious code through these algorithms simplifies the threat discovery process for these cyber-attacks (Diogenes and Ozkaya, 2018).

So how exactly does it work? Without getting too technical, Artificial Intelligence uses a programmed behavioral analysis, to continuously monitor and detect threats, and in turn provide an immediate defensive response alert to both known and new (zero-day) attacks (Simonovich, 2021). A massive amount of data is required for the algorithm to normalize the business environment and be able to detect any abnormalities that could be malicious activity. This use of AI and ML within anomaly detection (Halsey 2021; Xin et al., 2018) will increase

the detection rate of known attacks and reduce the false positive rate of unknown attacks (Xin et al., 2018).

AI aids cybersecurity at both a macro and micro level. At the macro level, for example, Artificial Intelligence provides companies with next generation firewalls (NGFW) that will help ensure better protection. As mentioned, the AI/ML algorithms will detect and block suspicious files, without needing to use a historical, signature-based database for comparison. If the files meet a specific threshold of specific behaviors, it is isolated and analyzed. Additionally, every time the algorithm is used, it learns from that experience and the previously analyzed behaviors and will become even more proficient in detecting those suspicious files in the future (Hall, 2021).

At the micro level, from the viewpoint of a singular device, heuristic analysis is at play. Heuristic analysis is a method of detecting viruses by inspecting the code for anything suspicious. Heuristic-based detection makes it much harder for zero-day attacks to break into a particular network. For example, the AI software for facial recognition provides another layer of protection to a mobile device, and to the network it is linked to (Hall, 2021).

Currently, there are many different developments being made to further this AI revolution. Simonovich discusses the new software of DeepArmor Industrial, which formed through a collaboration between Siemens Energy and the AI startup SparkCognition. This software uses Artificial Intelligence to flag cyber threats before the attack even occurs (Simonovich, 2021). Sophos, a global leader in next-generation cybersecurity released an announcement with some of their new developments in 2021. This included the SOREL-20M dataset for accelerating malware detection (a dataset containing 10 million disarmed malware samples that are available for download for research and training of models), the AI-powered Impersonation Protection Method (compares the name of emails sent to the user with those in a data base to flag suspicious messages), and YaraML Automatic Signature Generation Tool (a program that allows AI to “write” its own signatures) (“Sophos Announces 4 New…”, 2021). These developments only mark the beginning of what is to come for AI advancements.

As this technology advances, however, one thing to consider is the certifications and regulations that will come with it. Halsey emphasizes in her article the necessity of administering a standards-based approach, specifically like the ISA/IEC 62443 series, to be able to equip people and businesses with the knowledge of common procedures to be used in their security platform. A study showed that 81% of engineers and managers said that standards helped companies improve their compliance to regulations surrounding security (Halsey, 2021). Furthermore, Hunter's insights show that having a third-party certification system has proven to be preferred to that of a self-certification system that is currently in practice under the National Institute of Standards and Technology. The Cybersecurity Maturity Model Certification (CMMC) establishes this system of a third-party certification that may be soon required for eligibility to compete for defense contracts in the US. However, this may prove to be a barrier to entry in defense supply chains, and may end up being redundant, as it may be extremely difficult to keep up with. Hunter states that AI may be a better answer than this system (Hunter, 2020).

Future of AI in Cybersecurity and its Benefits

The future of AI is not known for sure, but many sources provide their insights and forecasts as to what can lie ahead. The technology is growing at a rapid pace and has many benefits, with some concerns within the cybersecurity industry (Hall, 2021; Maguire, 2022; Yampolskiy, 2017). According to a study explained in Hall's research, the market for AI in the cybersecurity industry is expected to reach $46.3 billion by 2027, showing the immense growth projected for this technology. He classifies four main benefits of the use of AI: the technology gets better over time, it can handle a large amount of data, has a faster detection and response time to attacks, and a better overall security for the business using it. He also addresses the concern that AI can render false positives and incorrect results if there is a lack of diversified data sets, but mentioned that Deep Learning, another branch of AI, could be the answer to reducing these false positives (Hall, 2021).

Going off the second of Hall's four classified benefits, there have been many sources that support the fact that AI can handle huge amounts of data that would be unthinkable for a human to do. The amount of data generated from previous cyber-attacks need to be analyzed

to create proper solutions to those that may arise in the future, and this analysis is way beyond the ability of humans, which is where AI comes in (Addo et al., 2019; Addo et al., 2020; Hall, 2021; Hunter, 2020; Maguire, 2022).

In both of Addo et al.'s books, he stressed how the quality and quantity of these data sources is crucial. This is what is required to train the AI to the best of its abilities, and what will foster it to continue to grow, and learn to adapt to new variations of cyber-attacks (Addo et al., 2019; Addo et al., 2020; Hall, 2021).

A much-debated concern of AI is whether it will replace humans in the future. There have been mixed opinions on this, some sources saying it will (Hunter, 2020; Yampolskiy, 2017; "AI likely to replace humans…", 2021), others say humans will still be needed (Addo et al., 2020; Labs, 2021). Obviously, nothing is known for sure, but this is a reasonable debate that is arising with this growing technology.

In a survey of IT leaders, 41% believed that AI will replace their role in their business by 2030, while 9% said it would not ("AI likely to replace humans…"; 2021). Hunter provides his commentary on how AI algorithms can be utilized to assess vulnerabilities, identify points of entry, and other problems that require real-time correction, rather than employing humans to do this work (Hunter, 2020). Yampolskiy takes the extreme approach, prefacing that in the case of a failure of a super intelligent AI (SAI) system, it could lead to a global tragedy. He states that Stephen Hawking, Bill Gates, and Elon Musk have all expressed their concern about AI's potential to evolve to where humans can no longer control it (Yampolskiy, 2017).

On the flip side, many sources claim that even with this growing technology, humans will not be redundant, and will still be needed to create and monitor the code. Addo et al. explains in his book how AI will detect the threats, and analyze pure and raw data, while humans will continue to be needed for correcting actions and for defense against the attack (Addo et al., 2020). One of Labs' interview responses also stated that software will not detect everything without the help of humans, for an absolute minimum, humans will be needed to tune the AI or ML to the particular environment (Labs, 2021).

Red Team/Blue Team

The concept of the Red Team/Blue Team exercise plays an important role in the development of efficient security systems. It simulates the battle between the cyber attackers, the red team, and the defenders, the blue team. There have been many literature sources that dive into the benefits (Diogenes and Ozkaya, 2018; Hautamaki et al., 2019), as well as the current issues with these exercises (Yamin and Katt).

The red team will perform an attack on a particular environment, while the blue team tries to defend against their actions to protect the environment's assets (Diogenes and Ozkaya, 2018; Hautamaki et al., 2019; Yamin and Katt). The book by Diogenes and Ozkaya dives into the details of this exercise, thoroughly explaining the roles and purposes of each side. This is an ongoing cycle that will improve with best practices over time (Diogenes and Ozkaya, 2018). The literature review by Hautamaki et al. also reflects this exercise, and how it emphasizes interaction and collaboration, yet it states that there is a lack of research in the practice of teaching within the cyber security environment. It claims that there are massive simulation environments that are currently being developed to enable more education and research in this aspect of the industry (Hautamaki et al., 2019).

The article by Yamin and Katt's research show that the exercise is a good tool for skill development, but it also explores the inefficiencies that lie in the current exercises. These issues arise in its development as it can take up to months to prepare this simulation, making it very costly and time consuming, further delaying the process of closing the current skill gap. Their literary research has shown that automation (AI) can be a solution, as it can reduce the cost and time required for preparation and execution, provide better training that is always available, and make the exercises repeatable for systematic training (Yamin and Katt).

Increasing Resource and Awareness Gap

As mentioned, with the world becoming more digitalized, it provides a larger surface for cyber criminals and their attacks (Halsey 2021; Maguire 2022; Simonovich, 2021). This makes it necessary for an increase in cyber security professionals to tackle this growing number of cyber-attacks, which there is a lack there of (Halsey 2021; Erdogan et al., 2020; McClurg, 2020; Simonovich, 2021; Yamin and Katt). As the technology is advancing, there is

a gap between the new generation of learners, and those currently in the field, which puts the industry at a disadvantage to the evolving hackers. It is important to facilitate ways to close this gap through education of these new advancements (McClurg, 2020; Erdogan et al., 2020).

Erdogan et al. created a method to help fill this gap of the lack of skill and awareness seen in the industry. Their paper discusses the logistics of this model that would help train and evaluate the skills of its participants (Erdogan et al., 2020). Features of this paper correlates with McClurg's article that emphasizes the importance of integrating this sort of training into schooling at all levels, and to introduce more opportunities for education in AI, to bridge this resource gap (McClurg, 2020).

Malware Evolution and Criminal Advances

As the defensive mechanisms are advancing, so are the malware attacks, making it important to analyze both sides of this evolution. Many books and articles dive into these current cyber-attacks and how there is a clear competition between the hackers and the defenders, each trying to best the other (Diogenes and Ozkaya, 2018; Kaloudi and Li, 2020; Labs, 2021). One survey by Kaloudi and Li on Artificial Intelligence based cyber-attacks in existing literature concluded that these threats are constantly changing, and many are incorporating the use of AI-driven techniques in their process. It explores the existing studies of these AI-based attacks, and reviewed 11 case studies with this topic, allowing the authors to provide their insights into these new threats. The authors classified these attacks into five categories: next-generation malware, voice synthesis, password-based attacks, social bots, and adversarial training (Kaloudi and Li, 2020). These AI-based attacks from cyber criminals are beginning to become more prevalent in today's age, and will only continue, making it clear that the best hope to defend against these constantly changing, automated attacks, is with AI (Hall, 2021; Kaloudi and Li, 2021; Yampolskiy, 2017; Xin et al., 2018).

Another literature source references this constantly changing environment, with interviews of the CEO of Verve Industrial, the global director of Cisco, and the Vice President of Operational Technology Security of Tenable, three cyber-security companies currently dealing with these threats. The article asks the question: are criminals smarter or are hacking tools better? (Labs, 2021). It shares similar thoughts with the article by Kaloudi and Li in the

sense that it is the increasing quantity of attacks, assisted by AI, that happen every day, that are the real problem due to how easy it is for hackers to share their techniques, and how much harder it is becoming for companies not using AI security to control (Labs, 2021; Kaloudi and Li, 2020).

The book by Diogenes and Ozkaya also dives into these attacks, showing where they started, with well known viruses, malware, trojans, and human error, and how they advanced to what they call "targeted attacks", crypto and ransomware. Although it doesn't mention the use of AI in these attacks, it still shows how the cyber-threat landscape is expanding and becoming harder for businesses to prevent and defend against (Diogenes and Ozkaya, 2018).

METHODOLOGY

Research Design

After conducting the literature review to get preliminary knowledge, eight interviews were conducted with cybersecurity professionals. These professionals came from companies that both developed their own AI engine and products, to companies that buy these products and implement them in their environment. Titles of these professionals range from Chief Information Officers (CIOs) to Vice Presidents, company founders, sales managers, and research directors. The questions asked were as follows (1-6 refer to developing and implementing AI within an environment, 7-8 refer to the training and shortages of employees in the industry, 9 pertains to malware evolution, and 10 pertains to the future of AI/ML) :

1. How is Artificial Intelligence being implemented into different products at your company? Do you collaborate with companies focused on AI advancement and products or do you develop your own technology?
2. If your company is the one developing the AI software, how do you go about that process?
3. What are the considerations when licensing security software in terms of cost?
4. Do the executives in the company understand and value AI in terms of security software?
5. Will the purchased AI software replace a method already in place or will it be in addition to preexisting security measures?
6. AI needs a large amount of data to be trained, how are you going about getting that data and diversifying your data set?

7. Has your company experienced a skill resource gap as the technology advances?
8. Have you ever participated in a red team/blue team exercise? If so, how did it go, and do you have any suggestions on improvement?
9. What is your experience with the evolving malware, and what are you or your company's responses to that? Have you experienced any attacks that have used AI?
10. What do you think is in store for the future of AI and its impact on the cybersecurity industry? Do you think AI will replace humans in the future or will humans still be needed to develop the software?

After conducting these eight interviews, a survey of seven questions[1] was constructed that references these questions and their responses. It was created on the survey platform Qualtrics and was sent out to 60 cybersecurity professionals. Out of the 60 respondents, 45 surveys were entirely completed and used in the analysis and 3 were partially completed and used to analyze only the first two questions.

Data Analysis

After downloading the data from Qualtrics into Excel, PowerBI was utilized to create histograms to explore the data. A number of tables were constructed in Excel to examine the differences between the survey selections and also to segment the data based off the companies that buy an AI cybersecurity product versus the companies that sell the product. Question 3 of the survey asked respondents to select a statement describing the company. One of the options was buying and implementing a product in their environment, and the other three options were different variations of selling the product. 34 of the 45 respondents were from companies that buy these products, while the remaining 11 sold these products.

Two of the seven survey questions provided quantitative ranked data, ranking challenges of implementing the AI from 1 to 5 (1 being the most challenging) and ranking how cyber-criminals can utilize AI most effectively (1 being the most effective). Given the small sample size, t-tests were run to compare the differences between the mean of the ranks to try to find the true order of challenges and cyber-criminal use effectiveness. One thing to be wary of with this analysis is that the respondents were required to assign a rank to each option for both

[1] Questions in Appendix A

questions. Because of this, it is unknown how large of a difference there was between ranks. For example, respondents could have had two challenges they feel are most challenging, but they had to select just one to be on top, and the other three challenges could have been way behind, but there is no way of knowing that from the data. Future research may break this question up and allow people to scale how much of a challenge each one is on its own, rather than looking at them all together.

After running the t-tests, before analyzing the p-value, the Bonferroni Correction was applied. This is an adjustment made to the confidence level (alpha) when several tests have been run, which was the case for this analysis. For example, for the first set of t-tests of the challenges question, 10 tests were run to analyze differences, so the alpha for these tests was 0.05 divided by 10, or 0.005. This was done in order to control the overall Type I error rate, or the probability of false positives (rejecting a null hypothesis that is actually true).

RESULTS

Interview Responses

The eight interviews conducted gave great insight on where AI stands in the industry today, and where it has significant impact. The following is a summary of each question's responses, in the order presented above, and grouped by the topic of the question. Interviewees will remain anonymous and will be numbered to see connections across question responses. Please see Appendix J for the full information collected from the conducted interviews.

AI Development and Implementation within an Environment

AI is being implemented into a variety of different products in the cybersecurity industry, with some companies developing their own products, and some companies buying those products and implementing them into their environment. Companies of interviewees 1, 2, 3, 5, and 7 utilize vendor-based relationships to get these products, and interviewee 3's company also does some in house developing. Companies of interviewees 4, 6, and 8 develop their own AI/ML cybersecurity software and sell it. Interviewee 8's company also develops AI/ML software for use cases outside of the cybersecurity space. Interviewee 7 says that they don't see many companies building out their own security technologies because it is cost prohibitive.

Cybersecurity professionals need to cover a lot of ground in terms of securing their network. They need protection at the desktop level (end point protection), the network level (like a SIEM (Security Information and Event Management) solution), and for e-mail (dissipating filters), which is one of the biggest threats today. Security starts by knowing your assets (Interviewee 3). Once you know what you have, you know what you need to protect. All interviewees mentioned how traditional cybersecurity was signature based, and because of the evolving attacks, has now transitioned to behavioral based detection. However, there is still no bullet proof security, or silver bullet, according to interviewees 3 and 4 respectively. Different tools stop different attacks. For example, firewalls stop external threats and endpoint management or EDR (endpoint detection response) alerts on user-based detections (Interviewee 4). All of these tools have been enhanced by AI in one way or another. For example, Interviewee 3 talked about AI enhancing e-mail protection. Traditionally, spam filters would look at who sent an email, which brand it is attached to, what the subject was, and maybe keywords that could hint at a phishing email (mortgage, loan, lottery, etc.). Now, there are AI/ML algorithms that scan the attachments of the e-mail and look at the associated behaviors. When the engine identifies a phishing email, they add it to the ML knowledge and flag that e-mail as a phishing email for everyone using the solution, not just those on that individual network. Interviewee 3 says the number one attack today is cryptologic, which attacks a network straight from the e-mail attachment, and encrypts everything on the network, further proving the necessity to have protection on all doors of the network. Another example from Interviewee 3 was firewalls. Traditionally, analysts would manually set network rules for the firewall to follow. Now, there is an AI component. Also referred to as FMC (Fire Management Control) systems, firewalls still have network rules, but they are now built on top of intrusion detection provided by the AI.

Other areas that AI has been used are vulnerability management (Interviewee 7), penetration testing (Interviewee 3) and UBA (User Behavior Analytics) (Interviewee 7). In vulnerability management, AI is used to quantify vulnerabilities. It does this by looking at different attributes of the weak points and classifies them in terms of how high of a priority it should be. Penetration testing is hacking your own system to look for vulnerabilities. Interviewee 3 said how he hires companies to do this and to find vulnerabilities in their network, and what's

more is that some companies are using AI to automate this hacking process, which will be discussed more when discussing red team/blue team exercises. In UBA, AI creates a baseline for the employees, or users, normally do in their day-to-day job, and then creates alarms for any deviations from that normal.

However, the area that seems to be of the most focus is anomaly detection (Interviewee 6) and incident response (Interviewee 7). Most of the time, analysts don't know what they are looking for in advance, so they can't have a simple search filter, again they need to look at the behavior. The best way to do that is to use AI/ML and statistical analysis because there are loads of data to be analyzed to identify deviations from the baseline, rather than just looking for a specific action (Interviewee 6). AI keeps track of what is considered "normal" and identifies abnormal behavior faster and better than a human can, as humans cannot look at that amount of data and quickly make conclusions on it (Interviewee 7). This is a good concept because it increases detection rates (Interviewee 6), however, there is a risk of getting a lot of false positives as well (Interviewee 1) which can be frustrating for analysts. Another issue that Interviewee 6 presented was the divide that this technology creates between the vendors and customers. Some customers lack transparency with the algorithms, and therefore cannot understand how and why they work. This dampens the collaborative nature of the industry, however Interviewee 6 assures that this will change with time, as it will be mandated across the industry to have more openness about the algorithms, allowing people to better understand how AI is enhancing their solutions. Within this context, using AI is all about trying to stay ahead of the bad actors who are also leveraging AI to hack systems. It is a consistent battle between the good and bad, each side fighting to stay ahead of the other (Interviewee 7).

There are so many different AI/ML and non-AI/ML products circulating the industry, so how do you choose which one to buy? Interviewee 5 goes as far as saying the products themselves don't really matter, as long as the security team knows which control is being satisfied, and how well it is being satisfied. The respondent stated, "You're better off buying the cheapest thing off the shelf at Walmart and putting some consistent and capable resources behind it than buying the most expensive product you could find and understaffing it." However, the

respondent also said that it is best practice to choose from the vendors on the Gartner or Forrester top 5 list, which are two companies that do research and reviews on all the cybersecurity products circulating at a given time, and this was also alluded to by Interviewee 7. What was interesting was that Interviewee 5 said that the AI component of products isn't a huge driver for their purchase choice, except for the fact that the whole industry is basically mandating that threats need to be detected based on behavior, however that happens, and AI/ML is the best way to do that.

Interviewees 6 and 8 gave insights into how they develop the AI/ML products. At a very high level, Interviewee 6's company's data science team reads up on an attack, works out an AI approach that could be used to prevent that attack, tests this approach to see if it is a viable solution, and if it is, engineering components come in to try to build it into a product. Interviewee 8 explained how the AI/ML engine is already there, it is just iterations of the engine that are being built. Their company has 300-400 engineers that are developing the software, and 300 data scientists. The cost of developing these products is in hours of labor for building the use cases, which can range anywhere from 10 to thousands of hours depending on said use case. Because the only cost for these products is labor, most of the employees are in Eastern Europe as the labor costs are lower than in other countries, like the United States.

When licensing security software, cost is usually the biggest consideration, according to Interviewee 2. Security analysts want to know that they are getting the features they want, how many seats (licenses) they get, how many computers it can be installed on, how long the contract is, who to talk to if there are issues with the software, etc. These are all considerations when deciding on which solution to purchase, and at what cost to the company. Costs of AI/ML cybersecurity products depend on the vendor, product, and size of the company, as stated by Interviewee 1, and it can get really expensive. A lot of vendors bill by seat, or the actual individual licenses they use for the product. Others will bill by throughput, or the amount of data that is being pumped through the solution. Interviewee 1 worked at a large company that didn't blink an eye when they saw a $6 million bill, and also worked at a small company who could see a $1 million bill and reconsider if the software is really needed,

so it truly depends on the size of the firm buying the product, and if they deem it worth it or not. The respondent went on to say that any cutting-edge technology isn't going to be cheap, but it is a must have for the company. If the product is going to help them have more actionable alerts in the future, companies won't be afraid to open their wallets.

Looking at a few more numbers, Interviewee 2 reported that their company paid around $80,000 for 300 users for a couple years (need to renew after period is up). Interviewee 1 stated how AI driven Splunk (tool that collects logs, runs queries on the logs, and alerts on them, with embedded AI capabilities of anomaly detection) can be up to millions of dollars per year depending on how much data you are sending through the solution. Interviewee 5 said that a product could cost anywhere from $50-100 per user per year, and that a basic Symantec product is as low as $25-30 per user per year, and that more next generation, industry-leading products aren't much more expensive as prices have come down over the past couple of years. If a company has 100 computers, it buys 100 licenses of the antivirus solution. Interviewee 3 spends somewhere between 12-15% of their IT budget on cybersecurity and also discussed how costs of these software have decreased over time and become more affordable. The first SIEM solution 7 to 8 years ago cost close to half a million dollars and up each year, and Carbon Black (cybersecurity leader) was also really expensive a few years ago. However, because more players are entering the market, there are more solutions available, so products have become more affordable to compete with other vendors, but it is still an expense.

One of the big advancements that led to AI being more affordable was the development of the cloud. Interviewee 7 said the "Cloud flipped the paradigm of AI being cost prohibitive on its head." Before the cloud was available, it would cost companies a lot of money to build out the infrastructure needed to house an AI product. However, with the cloud, vendors can build out a solution and overlay the AI across the cloud, carving out individual virtual environments for their customers. Without needing to build the infrastructure, much of the cost is avoided, as customers only need internet connection to access the solution and can now just pay per user. As Interviewee 7 said, "Today, there is not a cost issue anymore, everyone can leverage AI and that is the beauty of the cloud."

Even though solutions are becoming more affordable, you still "get what you pay for", said Interviewee 4. There are some great technologies out there, but there are also some not-so-great technologies. There are tools available that are a lot cheaper, but with that, there will not be the level of incident detection and functionalities that could be achieved with another product. For some people, that is okay, and they just want the bare minimum, but they have to consider how much it would cost them to recover if they were to get hacked. It is a common mistake for security people to think they are getting, as Interviewee 4 put it, "apples to apples" when comparing different products. Therefore, the best practice as stated earlier is to use third-party analysts that do reviews on the products. If a security team wants a really high-end, top-quality tool, they can see who the leader in the space is by looking at the magic quadrants created by Gartner.

As for company executives, cost is obviously an important aspect when deciding to purchase a new technology. As mentioned earlier, these technologies can be really expensive, but there are also ulterior costs to consider. If companies can use AI to cut through the amount of noise in their environment, and get real actionable alerts to their analysts, they may actually be saving money in the long run. This is because it cuts resource hours, eliminating the hours that people spend looking at "garbage" alerts that they won't have to look at in the future (Interviewee 1). Another way it could cut costs in the long run is to either collapse tool sets or to buy a tool with AI capabilities that once it is fully running in the environment, allows analysts to potentially turn off other products and stop using them, which saves more money (Interview 1).

Executives also care about security posture. They would have to spend a lot of money to try to clean up after a data breach. If some malicious event hits the news, companies may need to spend billions of dollars' worth in remediation to try to fix their reputation with their customers (Interview 1). Because of this, as Interviewee 2 said, "the minute you say ransomware and that you can stop it fairly quickly, they get on board."

Executives care about the company's reputation and its expenses, but do they understand AI? Interviewee 3 is an executive and stated that their peers understand AI but not from a security perspective. There is constant education to get executives, and employees in general, to

understand what security is and why you have to have it. Because of this, 4-5 years ago, Interviewee 3 spent $5,000 for a 250-person subscription to an education program with thousands of online courses. With some opposition, and some gratitude, employees were enrolled in these courses that educated them on topics like phishing emails and social attacks.

The next question was a clean sweep amongst interviewees. All respondents said that AI would be in addition to preexisting measures, or it would replace a method, but still be layered with other tools. There is no "plug and play" AI product, as Interviewee 1 said, or no "magic bullet solution", as Interviewee 3 said, where you just buy a product, test it, and then it is ready to go. Companies may claim this in advertisements, but it is not the case. This is because AI is still a relatively new concept and there is usually a lot of tuning needed for the model to work effectively in each unique business environment. Therefore, AI products will always be adding context to other tool sets or used as a separate addition (depending on what the company already has in place), with the hopes of improving accuracy and increasing speed to remediation and visibility in the environment (Interviewee 1).

Interviewees 2 and 3 discussed this layering concept. Interview 2's company uses AI products in addition to their normal antivirus, Windows Defender, which isn't as expensive as a product from a company like CrowdStrike (leading cybersecurity company). They explained that it depends on the company, as some like to run both to have the extra protection, like Interviewee 2's company, and some just run one product to cut the cost of running the other, however it isn't a huge cost to the company to run both. Interviewee 3 discussed how they layer solutions and keep adding on every year. They find it necessary to layer the security with a variety of tools and hope that one of those layers will plug a hole, or a vulnerability in the network. The company either comes up with another layer or replaces a layer (Interviewee 3 says they replace their endpoint protection every 3-4 years, and reevaluate every solution to either keep it, remove it, or replace it). Across the board, it has been shown in the industry that tools that incorporate AI/ML are getting ahead of the traditional legacy tools (Interviewee 4) and are enhancing existing technologies (Interviewee 6). Security monitoring has been around for a long time, and they are just now starting to be able to automate a lot of the actions through the use of AI/ML (Interviewee 6).

All of these AI/ML tools are great once they are up and running in the environment, but in order to get there, they require a large amount of data to train the algorithms on what a "normal" environment looks like. There is so much data circulating in the industry (Interviewee 4), it is just a matter of capturing it correctly and efficiently. The amount and type of data required for tools depends on the type of learning associated with the tool, as discussed with Interviewee 6. Deep learning is a subset of AI that requires large amounts of data to train the models. Supervised and unsupervised learning also requires a lot of data, but not as much as deep learning. For supervised learning, the data needs to be labeled, or in other words, tagged as good or bad. In terms of cybersecurity, this would mean malicious or not. Unsupervised learning tools do not require the data to be labeled, so it depends on which type of learning is behind the tool, and in this industry, it is usually supervised learning. Interviewee 6 explained that the problem that arises from this need to label the data is that some people don't like to share their true positives or provide the data from when they got attacked or breached. Because of this unwillingness, there is not a huge representation of "bad" labeled data in the datasets. Instead, there is an overabundance of benign data, or data for business-as-usual. Another challenge discussed with Interviewee 6 was the cost to maintain the data. Large amounts of data cost a lot of money to store, the more data a company has, the more they have to pay to store it. Some companies spend millions of dollars on data storage, but typically it can be anywhere from $10,000-20,000 each year. Cold storage (compressing the data) is a cheaper way to store data, but it is much harder to access quickly. Warm storage is more expensive but allows analysts to search and query the data directly, without having to decompress it first.

Regardless of how the data is stored, it is necessary to have the data handy to put into AI/ML algorithms. These technologies have feeds that can consume data from multiple different sources, as talked about by Interviewee 7. AI/ML technologies can pull data from computers, servers, cloud technologies, building management technologies (temperatures and card access systems), and compile it all into one database. From there, it automatically learns the kind of patterns that employees follow within the business environment, like how they interact with their computer, what sites they use, rooms they access, etc. The more data that is fed into the machine, the more it will learn this behavior, so the goal is to give it as much data as possible.

These sources range from the company who is using the specific technology, to a whole community of data. Traditionally, data would be pulled only from the network that the tool is running on, however it can now also pull data from the newly developed cloud (Interviewee 3). Obviously, for a company to train AI/ML on their specific business environment, they still need to pull data from their own install base (Interviewee 2) to feed into the model, but now they can get more data than that. AI/ML solutions can gather data from any feedback received from a new virus that presents itself within the cloud community and add it to their database. This means that the whole world is feeding information to the AI through various cloud technologies and vendors. Interviewee 3 stated that AI/ML cannot be isolated for one client, it has to be an aggregation of the whole community. An example of this provided by Interviewee 2 is the company CrowdStrike, who is a major player in the next gen AI field. They "crowd source", hence their company name, all the data that they need to train their models from their customers. Their sensors collect data from each customer they supply their product to, including their unique experiences and the attacks that are happening to them, and aggregate it all together. This allows companies to train their models with good quality data, that even includes the unrepresented "bad" tags as talked about above since generally, products and vendors tie into their contracts that they will be taking data from their customers and sending it back to their machines (Interviewee 5). Companies don't have the option to conceal their data on breaches with these contracts, it is companies who don't sign these contracts or create their products in house that hide that data. The more customers a vendor brings on, the larger their sample size gets. If they have their product installed on thousands of clients across tens of thousands of computers, running 5-6 times a day, it provides them with a pretty great data set (Interviewee 2).

One topic that was particularly interesting was the concept of a "honeynet" as explained by Interviewee 6. This respondent gave the same response of getting data from customers, but they also added that they capture samples of malicious activity to add to their dataset. They have researchers who create fake systems on the Internet and wait for those systems to be attacked. Once it is, they can see how it was done and what technologies were used, and they take that information back to their models to better train them for upcoming threats. It is a

really smart way of filling the gaps of the "bad" tagged data that is sometimes majorly underrepresented in the dataset.

Employee Trainings and Shortages in the Industry

The cybersecurity industry, as well as every other industry, fluctuates in terms of skill and resources available over time. Interviewee 1 mentioned how there would be a news article that mentions there is a 0%, or very low percentage, unemployment rate in cybersecurity that would lead to a flood of people into the market to try to secure those jobs. This is great for hiring managers who need people, but with time, it ends up hurting the security people who have been there the whole time. Interviewee 8 discussed this quite a bit. They said that in the past, there were no developers, so when there was one, they would get paid double (or significantly more) because the company is desperate for that employee. However, once the market floods, it ends up leading to pay cuts. One of the respondent's friends was the highest paid person at their company, but the company said he either had to renegotiate his salary (no one wants a lower salary) or they would have to let him go. Interviewee 8 compared this phenomenon to a pendulum swinging – everyone goes towards the market, flooding it, and then it eventually goes back to a shortage, and this cycle continues.

According to all interview respondents, there is a skill resource gap present in the industry today. Interviewee 3 discusses how there are gaps on many levels. Because the industry changed how to build and deploy an application by using cloud-based technologies, there is a gap in understanding how to build a secure application that uses those cloud services. Interviewee 4 explains how the gap in the industry has brought the need for these AI/ML technologies that are consistently evolving and changing to make up for the limited human resources. With the skill shortage, companies are much more reliant on solutions, specifically ones that prioritize different alarms and alerts, so that the human resources can be used efficiently and without waste. This prioritization of alarms gives different tier levels of incident response. A low-level alert, like someone talking to a competitor, should not be an incident that a high-level security analyst should handle, it should be someone in a more entry level position. This frees up the high-level analysts to do other important tasks in the

company, or look at other alerts, rather than wasting time on something that someone with much less experience can handle.

Interviewees 2 and 5 suggest that it is not so much the skills that create the most gap, but the shortage of human resources that is the bigger issue. They claim that most teams are pretty well-trained, the problem is the number of people with that training. Interviewee 2 had only three people on their security team.

Interviewees 6, 7, and 8 said almost the opposite. They all claim that it is hard to find good security people (Interviewee 6) or people who actually understand security (Interviewee 7). Interviewee 6 talked about how the market is starved of high-level analysts. Amazon, Google, and Microsoft have over 200 thousand security people between them, not leaving many for other companies to hire. Additionally, vendors in general are starving the market because they can pay their analysts more than the cybersecurity department at some other corporation or a smaller vendor, which drives the high-level analysts to leading vendor companies. Interviewee 7 discussed this shortage of people at the higher end of security. The respondent's vendor said that there was a shortage in finding a good security person with 10 years of experience and with cloud experience. COVID forced the world to embrace the cloud through remote work and learning, but there still are not a lot of people who understand how it works. However, in time the cloud might actually help close the gap because companies can leverage resources in the cloud, enhancing collaboration and the ability to outsource some tasks, and not have to do everything in house. Interviewee 8 mentioned how their company used to have 4 of the top 20 data scientists in the world, but now they only have one. Smaller companies are finding it hard to compete with other companies that can pay their employees more. What's more is that this shortage of security people affects the customers of these AI/ML products as well. Customer companies are struggling to understand some of the AI/ML products they buy or want to buy but cannot because they can't hire any good data scientists to understand them. The result, according to Interviewee 6, is that there are a lot of people in the middle, not necessarily entry level, but not with a lot of experience. The staffing shortage has allowed newcomers to work an entry level job for only a short time before getting a new

position. This just worsens the gap, as these people are now missing some key skills needed to be considered high-level analysts.

There are a variety of paths into cybersecurity that affect how a person is perceived by a company. The traditional route, according to Interviewee1, is to first work in IT (Information Technology) and build good foundational knowledge and expertise within that space, and then head into learning the security side of things. The other route, which is becoming more common, is people who are majoring in cybersecurity for their higher education. These security students have a risk analysis impact mindset, but they don't have the technical knowledge that comes from working in IT. It is harder to manage when these students are only looking through a security or risk level lens, and not understanding how it works on the IT side.

This presents challenges when hiring cybersecurity analysts. Interviewee 1 stated, "If you have somebody that has the IT background, and they have security knowledge, they are a unicorn, and should be hired immediately." The term "unicorn" is used to show just how rare that can be to find. Usually, it is picking one or the other, IT or security, and then the company needs to train that employee on the other side. Interviewee 6 had a similar response saying how they tend to recruit people who come from IT, development, or tech, and then they train the security part. The respondent claims that yes, one can complete a degree in cybersecurity, but one can't really only study cybersecurity. Security is predicated on understanding something about what you are securing, whether that is infrastructure, web services, or code, and that factor is what they referred to as a "bottleneck" in terms of training people. Building on that, Interviewee 5 also agreed, stating they found that people that have years of background in systems or network administration and then transition into security tend to pick up the concepts of the newer capabilities much quicker than someone who is only focused on security. Interviewee 3 had a slightly different approach, saying that they try to build all of their analysts in house, and they complement that with an outside contractor. This respondent only has one person working on cyber and not a whole team, so being able to utilize that external team that shares resources is becoming more common across small and mid-size organizations to close those resource gaps.

Once a company has a team of employees, it is important to train them to keep up to date with new technologies and advancements in the industry. Red team/blue team exercises are one way for analysts to gain hands on experience, though there are mixed opinions about them from the interview respondents, as well as mixed answers on if the respondent has completed one themselves. Interviewees 3, 5, 6, 7, and 8 have all participated in red team/blue team exercises in their careers, ranging from once to multiple times, and on the red team, blue team, or both. Interviewee 5 has a full red team in their company, and their SOC (Security Operations Center) is their blue team. The red team presents their results to the blue team who is responding and monitoring the network. Both teams compare notes to see how far the one team was able to get without the other team detecting them. Their company calls this aftermath purple teaming, which will be discussed more shortly in suggestions for improvement on these exercises. Interviewee 7 also had an internal red team, or a full-time pen tester (penetration tester), but mainly discussed how companies can make a profit by being the red team for other companies. For example, Interviewee 7 knows a man who can do everything from cybersecurity to breaking into buildings (using technology to pick locks and bypass system). Companies will hire him to try to break into their network. From the red team perspective, it doesn't matter how they get in, they just have to get in. Interviewee 6 differed as the company has an automated red team. It is sometimes hard for companies, especially smaller ones, to have a red team, or pen tester in house because their budget is already stretched thin. As a result, companies either just do a small test, which usually isn't enough to actually validate that what they are doing is enough, or they could use an automated tool (AI capabilities). This is seeming to become more common in the industry, which will be further discussed shortly. On the flip side, Interviewees 2 and 4 have never themselves participated in this exercise (and neither have any customers of Interviewee 4), although Interviewee 4 knows that his company has done them in the past, and some regions had their solution deployed in these exercises at other companies. Interviewee 2 said that they have done tabletop exercises, but never a full-blown red team/blue team. The respondent has always been on small teams in their career, so unfortunately never had the chance to participate in one due to the company not having the proper resources to put one together. However, the respondent is hopeful and wants to experience one in the future.

Even though these exercises are great learning experiences (Interviewee 3), there is room for improvement. The biggest one being the introduction of the purple team. Interviewee 1 discussed how the red team does their task, then sends a massive documentation of everything they did to the blue team, and then they do their task. With this traditional process, there is no cooperative factor between teams, leading Interviewee 1 to say this method is outdated. However, the purple team does this process side by side together, red team and blue team, or in tandem with each other (Interviewee 5). The red team is constantly working while in communication with the blue team, allowing the exercise to be more cooperative and agile, with no resource waste, and is becoming more common in the industry today (Interviewee 1). Interviewee 5 explained this transition in terms of the delay of results of the exercise. Traditionally, after the red team finishes their side and sends over the documentation, it could be weeks or months after the simulated attacks were performed. Having to go back weeks from a blue team perspective is quite challenging. Depending on what the company data retention is, the data might not even be there anymore (some systems only have data for 3-4 days). Doing the exercises side by side makes it a more efficient experience, and also avoids any problems with data loss. Concluding the purple team discussion, Interviewee 1 stated, "If you are in an interview and they ask you about red vs blue, you better start talking about purple because that is where it's going."

Another area for improvement within these exercises is transparency. This kind of goes hand in hand with purple team, but more importantly transparency with communication on when these exercises are taking place. Interviewee 1 recalled times as an incident response analyst where they would get multiple events that look really bad, and they would spend time and resources checking out those events, only to find out after the fact that it was a red team penetration test. In this case, analysts spent real time looking at fake cases, which isn't beneficial for anyone involved.

Another suggestion for improvement from Interviewee 5 was that participants need to be able to understand what all the possible permutations would be for the particular scenario that was run in the exercise, and not just focus solely on what happened. This allows participants to

increase their knowledge of what could possibly attack them in the future, and how it can change and adapt, rather than just looking at the singular scenario.

An interesting development in these exercises due to advancements in AI is automatic penetration testing, which was discussed with Interviewee 6. One of the respondent's friends works for HackerOne, which is a company that sells automated pen testing. It is important to find a balance between what is automated and what is not. If the pen testing is fully automated, it isn't as smart as a hacker. Automation is brutal, you can see it immediately from a detection point of view because it wasn't designed to be stealthy. However, if pen testing was done entirely manually, it takes a long time, and a lot of companies don't have the resources for that. Yet, humans can be stealthy as they have a thinking adversary and can adapt quickly. Each side has their advantages, so a combination between the two is the way to go in order to increase how much a company can test in a certain time as well as the quality of the testing, and that is where most of the market is going.

Malware Evolution

As Interviewee 8 put it, AI is a "double-edged sword", it can be used for both good and bad. As a result, there is a constant battle between cybersecurity professionals and hackers to see who can stay ahead of the other in terms of technological capabilities. Signature based detection was good, but it wasn't good enough to beat zero-day attacks and new methodologies that can come from hackers using AI to their advantage (Interviewee 1). Interviewee 6 stated that there was a 20-25% increase in terms of malware complexity according to one of their annual Director Reports. Companies are getting better at catching hackers, thanks to the help of AI/ML technology, and Interviewee 7 also said that the end user, or the weakest link in security, is more educated now than they have ever been in terms of knowing what not to click on. So as a consequence, the malware is starting to focus more on evading detection and is becoming more complex to do so. Interviewee 3's response was slightly different, saying that hackers don't necessarily need to be that sophisticated anymore. The respondent said that the volume of non-sophisticated attacks is increasing and briefly mentioned that there can be some increase in sophistication. Interviewee 3 used to learn to hack back in the day, since attacks then required certain skills, none of which are allowed to

be taught anymore. Today, you still need some skills, but you can also perform an attack from sending emails or conducting social attacks (phone calls). Interviewees 2, 3, and 6 all agree on the fact that there is still a large volume of phishing attacks, and that those are still effective ways to conduct attacks because the human element is still the weakest point. Interviewee 3 recalled a time where one of his employees got an email that looked like it came from their CEO asking the employee to do an urgent bank transfer. The employee almost did it, but luckily it was caught in time before any damage was done. The respondent does mention that their desktop users have been educated and they do catch a lot of phishing emails with the tools they have in place, but some do slip through the cracks. Interviewee 6 builds on this saying that attackers don't need to develop an AI to hack companies. All they need to do is send a well-crafted email, the AI would be overkill. "It's like pulling a cannon to shoot a fly" (Interviewee 6).

As for determining if a certain attack used AI/ML capabilities, it is still hard to tell, according to Interviewees 2, 3, and 5, but all think it is very possible. Interviewee 5 explained that usually, they try to tie the attack back to the particular group that was generating it. If that is possible, which sometimes it isn't (room for improvement), they can then dig into identifying the variant, how it gets used, what technology was used, and if it is AI or not. There isn't going to be a piece of malware that has AI embedded in it because the engine requires a large size that wouldn't be easily hidden on the system. Instead, hackers can use AI/ML to design the malware (Interviewee 6). Having AI/ML create attacks aids the hackers in detection avoidance, which according to Interviewee 5 (and stated above by Interviewee 1) is prevalent in the industry today. While none of the interviewees could get into specific attacks on their companies due to competition in the industry and confidentiality, a wealth of information was said about potential uses of AI/ML by cybercriminals.

Interviewee 5 explained that there can be variants that use algorithms to identify if that variant is trying to be run in a sandbox (virtual safe space to test malware), and if it is, it can refuse to detonate itself so that it can't actually be analyzed in the sandbox. They also went on to talk about "timeouts" which have been seen more recently. This is when a piece of malware may wait a period of time before detonating code to avoid detection. For example, it might wait

until 2am (or a period of inactivity) to perform a second step. This is dangerous because even if it gets detected, unless there is someone monitoring the environment 24/7, chances are someone isn't going to look at that alert until 9-10am the next day, or even later depending on how many alerts they had that day. At that point, the damage is done. Interviewee 6 builds off this by saying that "malware is modular". This means that one piece of malware can be installed on a computer, but that malware can also download other modules to do various things depending on the hacker's wants. This malware is adaptable and can change when needed to either avoid detection or to complete a certain task.

Another dangerous form of AI was also explained by Interviewee 6, and it is called Adversarial AI. This attempts to fool models with deceptive data. There are companies out there that develop this to test AI/ML solutions (pen testing), so if these companies have access to this technology, hackers do too. This type of AI can fake an image to fool an image recognition software, which could be used to break into someone's bank account for example. Respondent 6 also heard that someone was using ML to create fake companies, including fake people on LinkedIn with fake generated faces. This algorithm was trained on LinkedIn data, so it looked authentic, as all these individuals were connected to the company. Building off of that, Interviewee 7 discussed how malicious actors are now trying to leverage corporations that they have already attacks compromised, using that as a pivot point to attack other organizations. The respondent believes that these are the biggest threats right now, leveraging compromised company emails to launch more attacks. The way this works is that once they get hold of a corporation, usually a vertical that doesn't have a lot of money to spend on security (school systems, police forces, municipal entities, fire departments, etc.), and compromise that corporation, they can then use the company e-mail to send e-mails to other corporations. This looks completely legit and authentic because it is coming from a real company. Another advancement here is that instead of sending the full-blown malware right out the gate, hackers will send bits and pieces of it over time, and continuously do analyses to evade the system and find vulnerabilities in the network. This concept of leveraging compromised corporations is leading to zero trust in the industry, and corporations always second-guessing communications that may even come from close friends or business partners.

Future of AI/ML in Cybersecurity

The future of AI is unknown, but it is certain that the technology is going to continue to develop over the next several years. As mentioned earlier in the paper, signature detection is now obsolete. It can't be used anymore because the attacks are more sophisticated and know how to get around those signature-based algorithms. Because of this, every modern security team uses AI in one way or another (Interviewee 8). If you aren't using AI, you can't help an organization much with their security. It has taken about 10 years to get to this point, and a lot of people have left IT because of it. With security, it is necessary to stay ahead of trends and constantly be learning. Some people in the world don't want to continuously learn. They want to learn one thing and never evolve (Interviewee 7). Interviews 2 and 5 both said that more tools are going to be AI driven. There are so many use cases for the technology that may not have even been thought of yet. Interviewee 6 mentioned this, saying that no one has quite understood what can be done with AI in its fullest potential. The respondent states that this is normal due to the concept of recursive innovation. This means that until you have a certain foundation in place, new innovation doesn't really happen. In cybersecurity, since AI is a relatively new advancement, professionals are focused on laying down that foundation and getting the AI mainstreamed across the industry, directing their focus away from how AI can advance other aspects of the industry. Interviewee 5 reinforces this by saying that AI is going to continue to grow and be used in the industry due to its current success, so it will be continuously developed and used in a lot of other areas. Interviewee 6 said with respect to AI "Something so powerful like this takes on power that is hard to predict… We are just at the cusp of something that's going to change everything."

Interviewee 1 said they haven't seen a solution yet where AI is a "plug-and-play" thing. It has a long learning curve, not for the users but for the AI/ML to learn the business environment and baseline all of the data. Until it is fully learned and developed, the AI will produce a lot of false positives, and even when it is fully developed, it still may generate these false alerts. So, in the future, cybersecurity companies are going to try to make their algorithms more accurate and will also incorporate more threat intelligence driven alerts (threat prevention and mitigation). Interview 6's response builds of this. There are certain areas like NLP (Natural Language Processing) and image recognition where AI is really efficient. However, there is

still plenty of data out there to prove otherwise, the biggest example being self-driving cars. When people talk about AI in security, it is all about detection because it is the obvious thing in the industry and what most people focus on. They claim that the industry is getting pretty good at detecting threats, it is acting on those detections that aren't getting a lot of focus across the industry (automating incident response).

One advancement that is almost definitely going to happen in the industry with regards to AI, as mentioned before, is the transparency of the algorithms. People are already demanding to understand how these algorithms are being trained, what those data sets are, and how they work. So, in the future, the industry is going to start seeing organizations that are more open about their algorithms and how they work (Interviewee 5).

Interviewee 2 discusses a few other potential advancements within AI. This includes code that writes itself (automation), or systems that build themselves. Predicting analytics behind what the attackers are going to do, and how they can be stopped could also be developed. Algorithms in the future may be able to predict when an attack is going to hit and put preventative measures in to stop it from occurring in the first place. With all of this potential AI advancement projected for the future, it makes some wonder about where humans will stand when all is said and done.

Interviewees 1-7 all said that AI won't replace humans in security, at least not any time soon. Even with advanced AI, humans will still be needed to make the phone call, pull the trigger, respond to something, etc. – there are some analytical things that only humans can do (Interviewee 2). Interviewee 3 stated, "As long as the enemy of the AI and machine learning is a human, the human will always win." Engines are as good as what they learn and the data they read, but on the other hand you have humans that excel in creativity, and they will come up with something that the machine didn't think of. Humans are still needed to challenge the AI/ML and will be needed for a long time. Interviewee 4 says that the technology has come a very long way, as it can shut things down automatically when they are happening, but there always needs to be a level of analyst involvement to confirm that it is shutting down the right things, for example. The respondent also discussed how some of their customers say they are trying to get to a place where they have basically no analysts, yet opinions vary on this across

customers. SOAR (Security Orchestration Automation Response) is a software that will either disable a user (put it into a quarantine type place) so they can't access the network, or it will block certain items (like a compromised email or even a user). Another version of this tool, mentioned by Interviewee 5, is IPS (Intrusion Prevention Systems), which also removes the human element by automatically blocking potentially malicious activity. It's tools like these that are convincing some people that they will not need any analysts, because the software does everything the analyst would be doing. Some believe that these tools will be able to solve everything that humans can, but Interviewee 4 disagrees. If the CEO of a company is suddenly being targeted as a malicious user, SOAR may block the CEO from the network. The CEO should not be cut out of its own network, it could be detrimental to the business. If that was to happen and there were no humans, how would the model know that the CEO can be an exception? Interviewee 5 mentions how the IPS tool makes it easy to see what was blocked, and if it was wrongly blocked, it can be whitelisted (telling the machine to ignore it next time). But this action has to be taken by a human, the machine won't do this on its own. It is important to think of every possible scenario that could happen at any time, and without a human to correct the machine, there could be major problems.

Academic studies show that AI excels at certain problems, while humans excel at certain problems. If they were to work together, they are far more efficient than each one individually (Interviewee 6). With that in mind, Interviewee 6 also said that if AI does happen to replace humans in the far future, it wouldn't happen in the security industry because the industry doesn't invest enough in AI compared to other industries like finance and healthcare that invest a lot more. The respondent claims that we are still 20 years away at least from full AI (if full AI was the focus), but at the same time, it is unsure if the world is even trying to solve that problem. AI doesn't come up with anything new on its own. It is intended to free up humans to do tasks that the machines can't accomplish. Interviewee 7 builds on that concept saying that AI isn't replacing people, it is doing the job that no one wants to do, better than a human. In turn, this allows the corporation to leverage the intelligence of a human to innovate or do things they weren't able to do prior to AI enhancement. Suppose an employee spends 50% of their time on those repetitive tasks, and 50% of their time thinking, helping the business, and driving innovation. Once you eliminate the repetitive task they need to do with

the AI, that employee can now spend 100% of their time being innovative. So, if anything, it will replace the people who like to do repetitive tasks.

Interviewee 8 had this perspective, saying that AI has already in a way replaced humans, just not entirely. Humans will still be needed to develop the code, but a lot of code has been written in the past 5-10 years, so it is a matter of putting jigsaw pieces together to build more code. Automation in code building is available and will only improve with time, but it isn't at the point, nor anywhere close to it, where humans will no longer be needed. Interviewee 5 builds on this saying that the industry is already past the point where some aspects are fully automated, and it is replacing humans for those certain actions. The industry is starting to see a lot of companies promoting the concept of being able to do more with fewer people. It is tools like SOARS (previously discussed) and IPS (intrusion prevention systems)

It is the automation aspect that if anything, would replace some humans. Interviewee 7 said that a lot of things can be automated, but people do fear it. Yet they proceeded to say, "I don't have any fear of that, I'll embrace it" showing the variations of opinions on the topic. People have to fight a lot of uncertainty and doubt when it comes to these AI technologies and automation.

Survey Responses

The following sections are information found from each question of the survey. The first three questions have 48 responses (3 partially completed surveys) and the last 4 will have 45 responses.

Question 1: In What Area Does AI have the Most Benefit

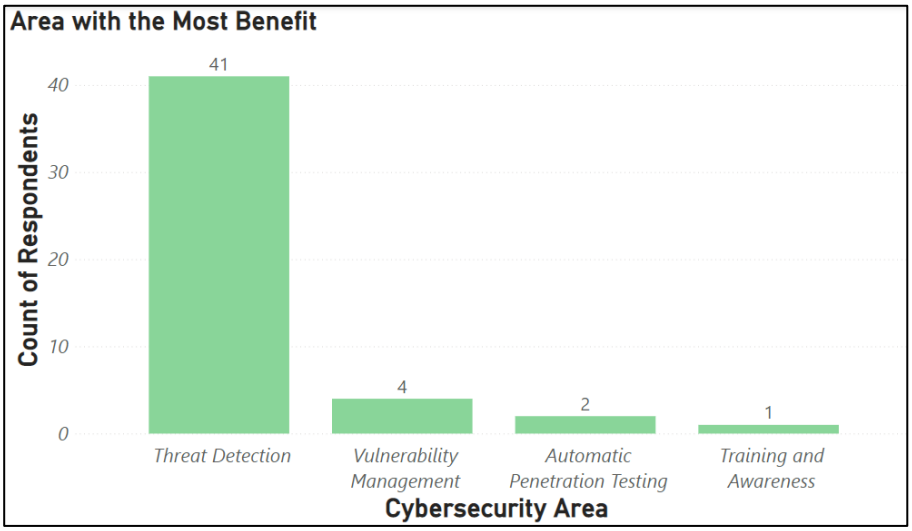

Figure 1- Information security area that AI/ML has the most benefit

According to the survey responses (as shown in Figure 1), threat detection is by far the area that AI/machine learning has the most benefit in. 41 out of the 45 respondents (85%) selected this option, which is clearly the majority. When segmenting the data between the companies that buy versus sell the product, the same trend is shown.[2] This is understandable because threat detection was where AI initially entered the industry, and it has remained a strong focus ever since. Threat detection through a behavioral-based approach is all about finding a normal baseline in one's environment, and then identifying any deviations or anomalies from that normal. Since this is what AI is best at, it makes sense that threat detection was AI's entryway into the cybersecurity industry and why cybersecurity professionals think that this area has the most benefit with an AI application.

[2] Figures and Table in Appendix B

Question 2: The Defense Mechanism Most Enhanced with AI Applications

Intrusion detection & prevention and User Behavior Analytics (UBA) seem to be the defense mechanisms that can be most enhanced with AI applications, as shown in Figure 2. These two mechanisms were selected by 31% and 33% of respondents, respectively, which shows that all of these areas can be enhanced as there are responses across the board, with not one having the vast majority like threat detection did in the previous question.

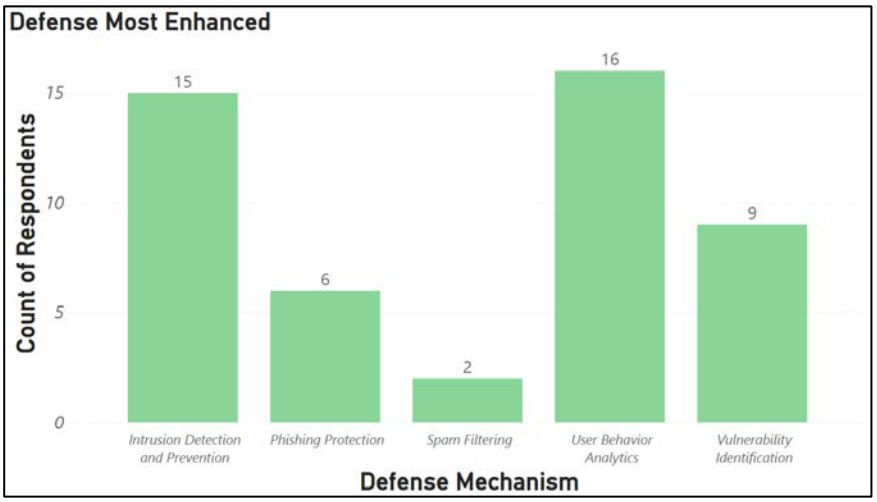

Figure 2 - Cyber defense mechanism that can be most enhanced through the application of AI/ML according to all survey respondents

Looking at Figure 3, what is interesting is that for the companies that buy products, the same trend is shown (UBA and intrusion detection & prevention being unmatched for top two), however for the companies that sell products, vulnerability management is tied with UBA (29% each), and intrusion and detection fell slightly behind with 21% of respondents.[3] This could be due to the lack of respondents from companies that sell these products (sample of only 14), or it may hint that the producing companies believe AI can enhance vulnerability management more than the consumer companies.

[3] Figures and Table in Appendix C

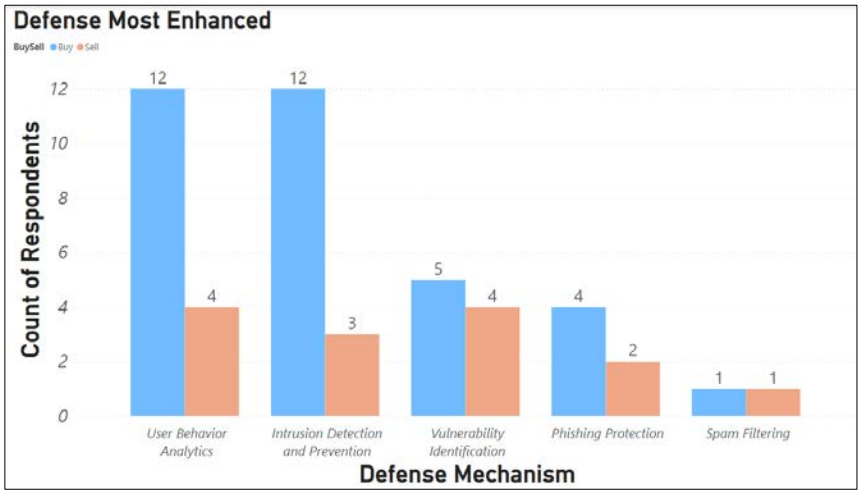

Figure 3 - Cyber defense mechanism that can be most enhanced through the application of AI/ML according to all survey respondents split between buy (blue) and sell (orange) data

Question 3: Company Description (Demographic – Buy or Sell)

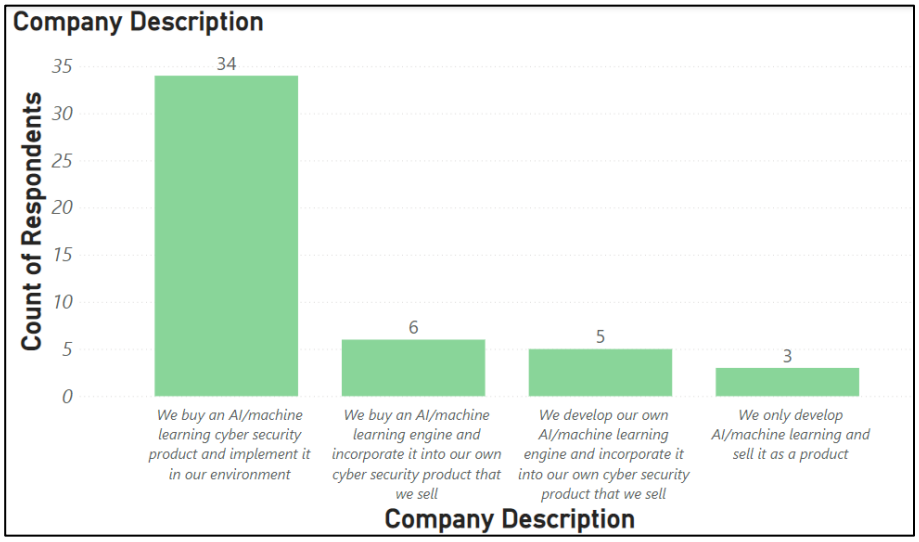

Figure 4 - Survey responses to which statement best describes the company the respondent currently works for

This survey question was used to gauge which companies were answering the survey without asking them to break anonymity. In Figure 4 above, the first column was declared as companies that buy these products, and the next 3 were grouped into companies that sell these products. A majority of the data comes from these product buying companies, so it would be worth trying to get more responses from companies in the other three categories. These two

groups were used to see if there is any difference in preference or opinion between the companies who make the product versus those that use it.

Question 4: Challenges of Implementing AI

When analyzing this data as a whole, strictly looking at the means of each ranked challenge, the order is as follows[4], with the first being the most challenging:

1. Expertise
2. False Positives
3. Effectiveness
4. Cost
5. Malware Complexity

This order remains the same for companies that buy the products, however, cost and effectiveness are switched when ordering the ranks of the companies that sell the products. To see if there was any true difference between these mean rankings, multiple t-tests were run for this question, looking at it as a whole and splitting it up between buy and sell data.[5] While most of the tests came back inconclusive with high p-values, a few conclusions could be made.

[4] Figures and Table in Appendix D
[5] All T-tests Calculations in Appendix E

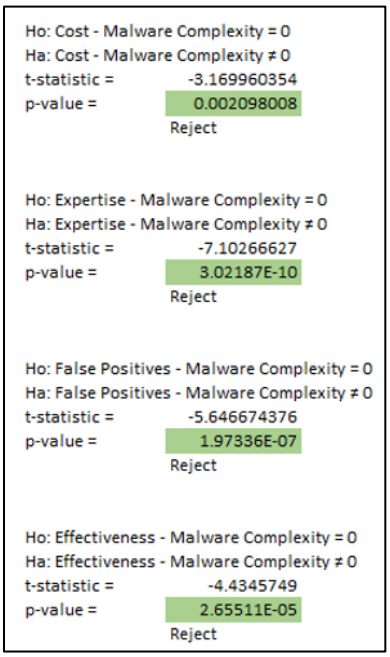

Figure 5 - T-tests comparing the means between challenges of implementing AI/ML based on survey responses

The clearest conclusion that can be made across all respondents, is that malware complexity is the least challenging obstacle out of the five options when implementing AI. This can be seen mathematically for the entire survey dataset in Figure 5. When comparing the four other challenges to malware complexity, each one returned a negative t-statistic, meaning that the malware complexity mean was significantly higher than the means or averages of the other challenges. Since a high average means a low ranked challenged, it can be concluded that malware complexity is the least of companies' worries when it comes to implementing AI. These conclusions remained true for the most part when segmenting the data based off of the buy and sell companies, with the exception of the cost challenge. There was no statistical proof that there was any difference between the means of cost and malware complexity when looking at only the companies that buy products or only the companies that sell products. The only statistical significance with the cost challenge was when looking at all the respondents as a whole. It can only be concluded that cost is less of a challenge than expertise, and more of a challenge than malware complexity, according to all respondents as stated earlier.

An interesting result was that of the sell dataset. For the industry experts that were surveyed, there was strong statistical evidence that expertise is a greater challenge when implementing AI than effectiveness. Since effectiveness was also ranked one slot lower than the overall data and the buy data, this makes sense. There must be some reason as to why producers of these products believe that effectiveness isn't as much of a challenge, and it may be because they are the ones creating the "effective" product. More research would need to be done to look into this, as mentioned earlier, it is unknown how big of a gap there is between the ranks when the respondents filled out the survey.

Question 5: Cyber Criminals Utilizing AI

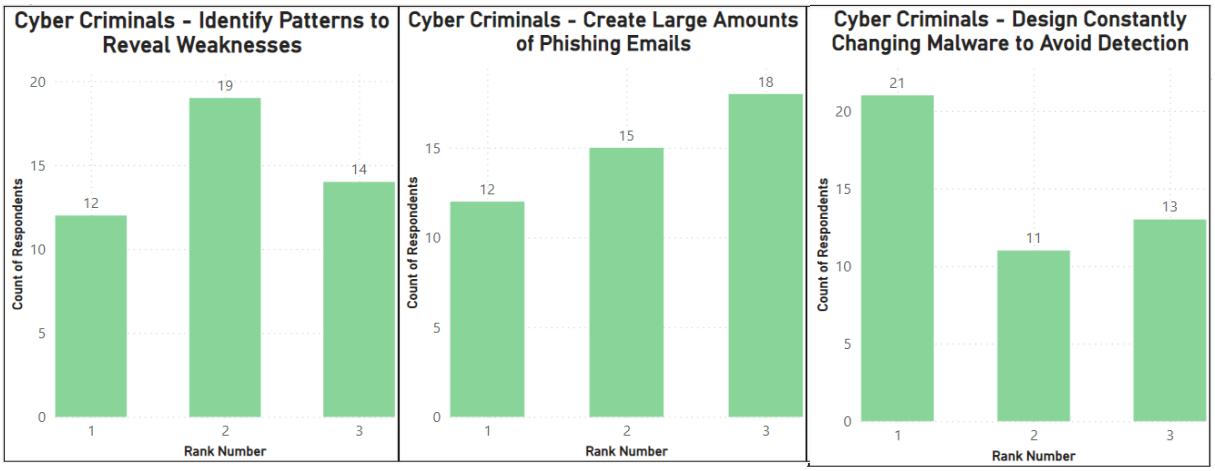

Figure 6 - Histograms that show how many respondents ranked each option of how Cybercriminals can use AI for malicious reasons, 1 being the most effective application

This question aimed to gauge how survey respondents think hackers use AI for malicious reasons. Because this was also ranked data, like the challenge question, t-tests were run to see if there was a difference between the means of each option. However, none of the tests came back with any statistical significance[6], so no conclusions can be made that there is a difference between the ways hackers can use AI. Whether it being due to there only being three options to rank, which confines the means and standard deviations, or not, it can be concluded that there are mixed opinions on how hackers are most using the technology. This is logical, as the respondents are not hackers themselves, and would only be able to form opinions based off attacks that have either happened to them, or ones that they have read about on the news or heard from some other source. It has been seen through plentiful research and new advancements in IOAs (indicators of attack) that cybercriminals have used AI to leverage more stealthy techniques or mass scheming. It is a matter of what the respondents felt was the most effective for them, which can only be an educated guess and not known for certain.

Even though the t-tests came back insignificant, it can be said that the survey respondents believe that cybercriminals can use AI to identify patterns in computer systems that reveal weaknesses, create a large number of phishing emails to spread malware or collect

[6] T-test calculations in Appendix F

information, and design malware that is constantly changing to avoid detection by automated defensive tools. This is because the ranks across all three options were relatively even, as seen in Figure 6. However, the order as to which is most effective through the use of AI is unknown.

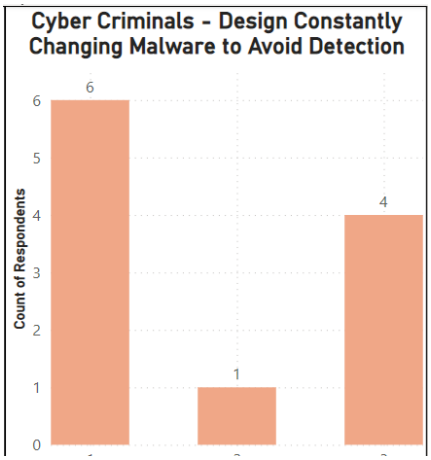

Figure 7 - Histogram showing how respondents who sell AI/ML products ranked how Cybercriminals can use AI to design constantly changing malware, 1 being the most effective use

This can be seen in the segmented data as well[7], with the exception in the sell data segment of the "design constantly changing malware to avoid detection" option. It seems that these companies either think this is the most effective, or the least effective, as only one respondent ranked it as number two (see Figure 7). This could be due to the fact that there are only 11 respondents in this segment, so more respondents may be needed to fully make this conclusion.

Question 6: Where AI is Currently

This question aimed to gauge current trends of AI as it stands in the industry today. Some of these results were shocking (seen in Figure 8). It was expected that more people would have selected the option of AI being able to accomplish mundane, non-complex tasks, as it is the least complicated of the three options, but only 62% of all respondents selected it. What was even more interesting was that 91% of all respondents selected the option of AI being able to analyze and correlate events beyond a human's capability. This was expected as that is the main reason why AI is being used – to figure out what a "normal" environment looks like and call out any anomalies. However, this task seems much more complicated than that of mundane ones, so it begs the question, why don't more people think it can accomplish the mundane, repetitive tasks if they think it can analyze events better than a human?

It was also surprising to see that 38% of respondents selected the option that AI takes the place of advanced human analysis. It was expected that much fewer respondents would select

[7] Figures and table in Appendix G

this option, as many literature sources have claimed that humans are definitely still needed right now to supervise the AI as it learns. This is an action of AI that was assumed to be of the future, which makes sense why 62% of respondents did not select this option.

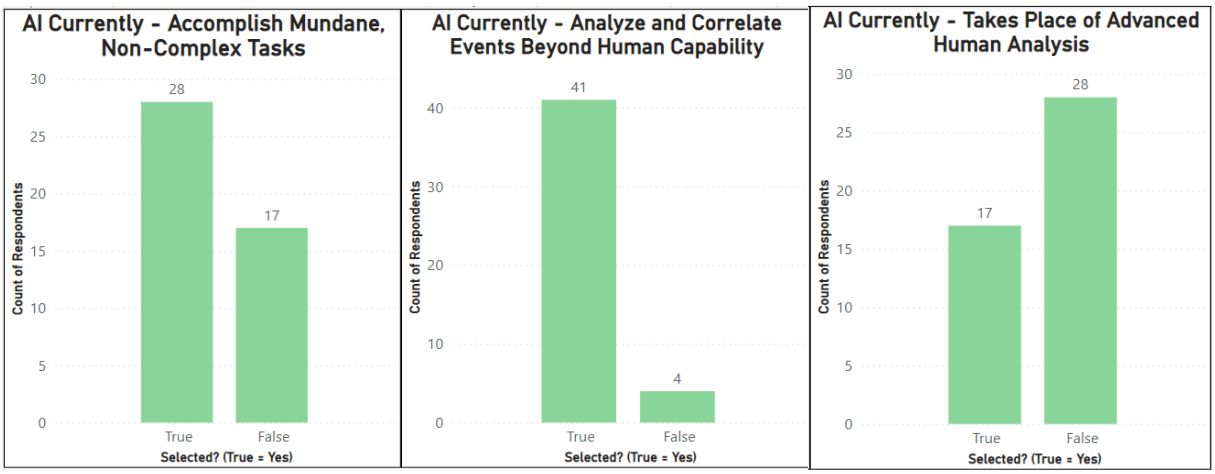

Figure 8 - Histograms that show the number of survey respondents that selected (True) each task that states where AI currently stands in the industry.

When segmenting this data into the companies that buy or sell products, the same trends were seen. Most people selected the option of AI being able to analyze and correlate events beyond that of a human, with 94% of buying respondents and 82% of selling respondents. Because the sell data is so small, more data would be useful to try to conclude if there are differing opinions between the two groups.[8]

Question 7: Where AI could be Headed in the Future

The future of Artificial Intelligence is unknown, but this question aims to gauge what tasks respondents think AI will be capable of in the future. As seen in Table 1, the vast majority of respondents believe that AI will be able to predict threats and actively defend against attacks (84% and 87% respectively). If these fruitions come to pass, that would mean a huge advancement for AI, and it would further increase its impact on the industry. When looking at these two tasks split between buy and sell data, it is interesting to see that 100% of respondents that sell products believe that AI will be able to predict threats in the future, and 90% selected that AI will be able to actively defend against threats.

[8] Histograms and Tables in Appendix H

Looking at all the tasks overall, it is easy to conclude that there is great potential for any of these five tasks to be one of AI's new advancements. There were at least 35% of respondents selecting each option, which leads to this conclusion. The task that seems the least likely of them all is the ability for AI to perform a counterstrike, or "hack back". 38% of overall respondents selected this option, while 62% did not (see Table 2) and it is understandable as to why that is. There could be many legal complications when it comes to hacking someone back, so if this advancement was to happen, it would have to be done with caution. Tables 1 and 2 below give the number of respondents that selected or did not select each option. See Appendix I for visual representations (histograms) of this data.[9]

Table 1 - Shows the number of respondents that selected each potential task that AI could evolve to perform in the future

AI Future	Buy		Sell		Total	
	Selected	Percentage	Selected	Percentage	Selected	Percentage
Predict Threats	27	79.41%	11	100.00%	38	84.44%
Write Code	24	70.59%	7	63.64%	31	68.89%
Hack back (counter)	13	38.24%	4	36.36%	17	37.78%
Replace Human Analyst	15	44.12%	3	27.27%	18	40.00%
Actively Defend	29	85.29%	10	90.91%	39	86.67%

Table 2 - Shows the number of respondents that did not select each potential task that AI could evolve to perform in the future

AI Future	Buy		Sell		Total	
	Not-Selected	Percentage	Not Selected	Percentage	Not Selected	Percentage
Predict Threats	7	20.59%	0	0.00%	7	15.56%
Write Code	10	29.41%	4	36.36%	14	31.11%
Hack back (counter)	21	61.76%	7	63.64%	28	62.22%
Replace Human Analyst	19	55.88%	8	72.73%	27	60.00%
Actively Defend	5	14.71%	1	9.09%	6	13.33%

[9] Histograms in Appendix I

Correlations Between Survey Results and Interviews

There were many similarities between the survey results and the responses given from the interviewees, in part because the survey questions were formed based on the interviews. Unlike the survey respondents, the interviewees were pretty split between companies that buy AI/ML products (Interviewees 1, 2, 3, 5, 7) and companies that sell them (Interviewees 3, 4, 6, 8). Threat detection was discussed in one way or another with every interviewee (specifically mentioned in interviews 1, 3, 4, 5, and 6), which aligns with it having 85% of respondents selecting that option in the survey. This feeds into the next question with the cyber defense mechanisms that are most enhanced through the use of AI/ML. Intrusion detection and prevention falls into the same overall category of threat detection, which had a part in every interview. User Behavior Analytics (UBA) was also talked about by a majority of interviewees (specifically mentioned by interviews 1 and 4), aligning with those two options being the top picks for the survey.

When analyzing the challenges, expertise and false positives (ranked top two in the overall survey dataset) were widely discussed among interviewees. All interviewees said that there is a gap in the industry in terms of skills and resources. Interviewees 1 and 3 said it was knowledge that was preventing them from hiring the right people, while Interviewees 2, 4, 5, 6, 7, and 8 just said it was a lack of security people in general, not specific to just skills. These responses align with the survey results where expertise was ranked the biggest challenge. False positives were mentioned by Interviewees 1, 4, 5, and 7, making it clear this is still a large challenge, and affirming its rank in second place on the survey challenge list. When looking at cost, most interviewees mentioned how this isn't a big problem anymore. Interviewee 7 discussed how the development of the cloud made it more affordable for users, and Interviewee 3 also said it was more affordable due to more players entering the vendor market. Interviewee 5 also said implementation costs are relatively low, it is the vendors that have the massive costs. The only conflicting perspective was from Interviewee 6 who stated that storing large amounts of data can still be really expensive. But, from the order of challenges from the survey, it is shown that the respondents agree with Interviewees 3, 5, and 7 and ranked cost as a lower challenge than the others listed. As for malware complexity, although Interviewee 6 gave a statistic that there was a 20-25% increase in complexity of

attacks, respondents still ranked this low. This could be because beating the hackers is the name of the game in the industry, and there is also a wide variety of complexity and sophistication across hackers today, which could lead this not being a major problem when it comes to the implementation of AI into one's environment.

Transitioning to the cyber-attacks, Interviewee 5 mentioned the ability for malware to avoid detection through detonation refusals and timeouts, while Interviewee 6 mentioned how malware is modular. 21 survey respondents ranked this as the topmost effective way hackers can use AI to their advantage, showing that many other people think the same way as Interviewees 5 and 6. What was interesting is that Interviewee 6 said "it's like pulling a cannon to shoot a fly" when referring to hacker uses of AI. They said that hackers don't need to develop AI to hack companies, they could just send an email. It seems like 18 of the survey respondents agree, ranking this option as the least effective, but 12 other respondents ranked this first, which hints there is a debate as to how these hackers are best using the technology.

The question about the future of AI in part stemmed from the Interviewees as Interviewees 2 and 8 both mentioned the ability for AI to write its own code in the future. Interviewee 8 says this already exists and will only improve. 69% of survey respondents selected this option, which is more than the 25% of interviewees that mentioned it. Interviewee 2 also said that AI will be able to predict threats, and although they were the only interviewee to mention it, a lot of the survey respondents (84%) agreed. Next, Interviewees 2 and 4 discussed how AI will be able to automatically block threats or disable users (some tools already can, but need improvement), and this also got a healthy response from the survey with 87% of respondents selecting that option. Lastly, a majority of survey respondents (60%) and a majority of interviewees (1, 2, 3, 4, 6, 7) believe that AI will not replace a human analyst in the future. Interviewee 8 said the industry is still ways away, and Interviewee 5 said that it has already replaced them in a way, but not fully. Interviewee 4's customers believe that it can happen, which aligns with the 40% of respondents who did select this option. Most of these responses come from the companies that buy the products, further aligning with Interviewee's 4 statement about the consumer companies believing this is possible.

DISCUSSION

Gartner is a $5.5 billion company that does cybersecurity industry research. They have over 19,500 associates in more than 85 offices globally. For over 40 years, they have been providing insights and expert guidance to their clients all over the world about the best security practices in place. They make sense of trends and anticipate obstacles to be overcome in the future. Their experience and knowledge in the field makes them a great resource to see where AI stands in the cybersecurity industry today, and where it could be headed. In this discussion, the findings from the survey and interviews will be compared and contextualized to what Gartner has said in recent publications.

To start, Gartner wrote an article talking about the top trends in cybersecurity from 2022. In that article, they say that hybrid work and the development of the cloud have introduced new cyber risks. 60% of knowledge workers are remote, and at least 18% of them will not return to the office. This greater use of a public cloud, among other factors, has exposed new surfaces for hackers to attack, leaving some organizations more vulnerable than they have been. Interviewees 3 and 7 discussed this same concept of the cloud changing the industry. In another portion of this article Gartner discussed that security products are converging. Vendors are combining security functions into single platforms and are introducing options that make the packaged solutions more attractive. This in turn should reduce complexity of implementation, cut some of the costs, and improve overall efficiency of the tools. This is in line with what Interviewee 3 was saying about the prices of these tools coming down and becoming more affordable in recent years due to competitive pricing, and also this was also shown in the survey data with cost being ranked quite low compared to the other options.

Interviewee 6 discussed the problems with transparency between AI vendors and their customers. Gartner has also done research on this, and their results are that some organizations have deployed hundreds of AI models that none of the IT leaders can explain or interpret. In turn, Gartner expects that by 2026, corporations that initiate AI transparency, trust, and security will see their AI models achieve a 50% improvement in terms of the adoption rate, business goals, and user acceptance. On this same topic, the need for transparency will only increase as more organizations are going to outsource their

cybersecurity. It has become far too complex for some organizations to manage on their own, either due to a lack of skills or resources, so many will begin to look to external consultants for their security. This concept of outsourcing was seen with Interviewee 3's responses, as they were an example of a company that only had one person on their security team and needed more resources.

An interesting up and coming concept Gartner also discusses is adaptive AI. Adaptive AI learns as it is being built, and unlike traditional AI, can edit its own code for real-world changes that were unknown when the code was first written. This type of AI can react much quicker and more effectively when it identifies threats. The Gartner Distinguished VP Analyst, Erick Brethenoux said "Flexibility and adaptability are now vital, as many businesses have learned during recent health and climate crisis." Gartner expects that by 2026, corporations that have adopted this adaptive technology will outperform their colleagues in the time it takes to operationalize AI models by at least 25%. While none of the interviewees specifically mentioned this concept, it was hinted at by Interviewee 2 who said that AI may be able to write its own code in the future, and also inferred by all the respondents who selected that option in the survey.

A good way to visualize the development of these AI tools that are circling today is through Gartner's hype cycle, seen in Figure 9. The first phase is called the innovation trigger. This signifies a potential breakthrough that kick starts a technology, but there are often no usable products that exist. The next phase is the peak of inflated expectations. This is when there are a lot of success stories, but still some failures. After that is the trough of disillusionment. This is when interest in the technology starts to fade as there are more and more failures. This is followed by the slope of enlightenment, which is when the interest begins to pick up again and more benefits are realized. Lastly, there is the plateau of productivity. This is when the technology becomes mainstreamed and is adopted by most of the companies in the industry.

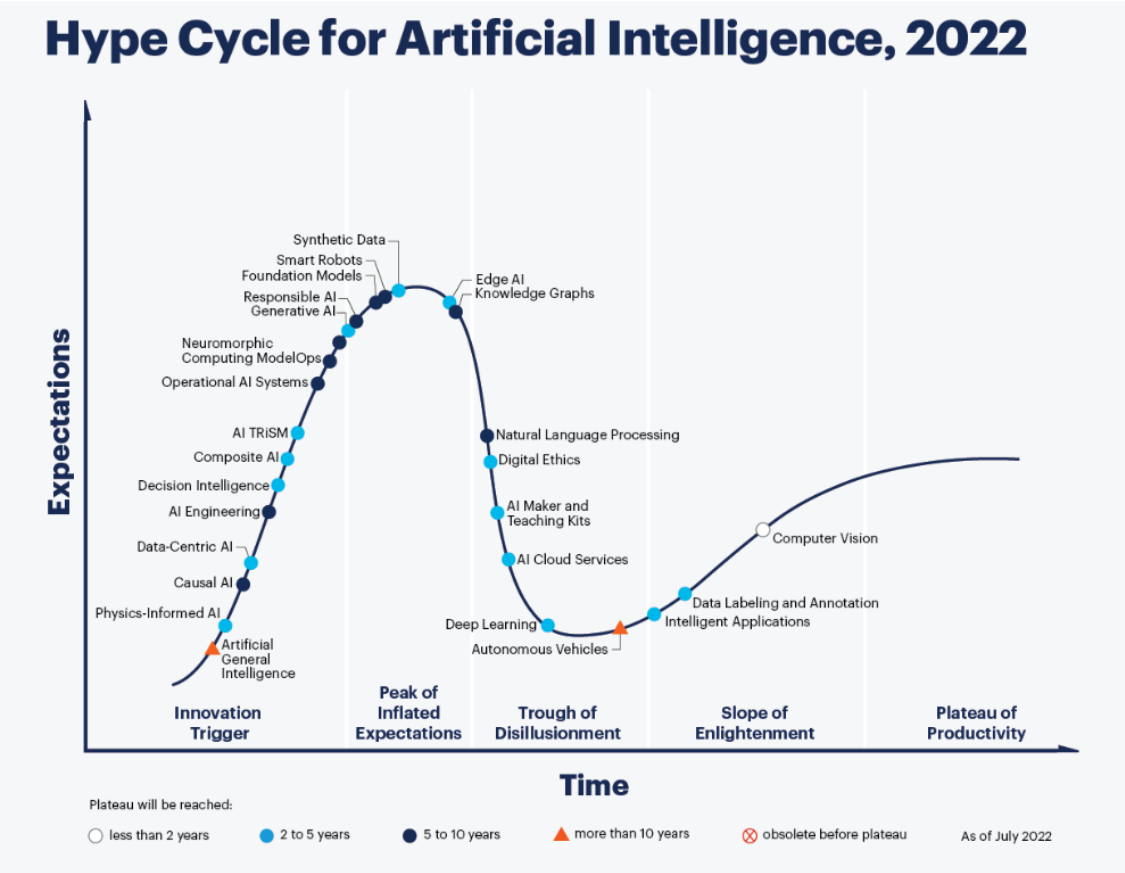

Figure 9 - Hype Cycle for Artificial Intelligence, created by Gartner

For the security industry, AI TRiSM (AI trust, risk, and security management) is the best to analyze. This point falls in the innovation trigger phase of the hype cycle as seen in Figure 9, with an estimate of 2-5 years for it to reach the plateau of productivity. The data collected in the interviews and survey could align with this position, but it could also be argued that it is in the trough of disillusionment phase. It can be said that AI has not yet reached its full potential, and that there is immense room for growth, suggesting that it has not yet reached the peak of inflated expectations. On the flip side, it can be said that the benefits of AI in the industry have been realized and there are workable products in circulation, but it is now in the phase where companies are trying to figure out if it is really worth it to put their effort and resources into developing these tools. As seen from the interviews and surveys, there are still a handful of challenges that need to be overcome before the technology will be mainstreamed, specifically expertise and false positives. Both of these challenges will fade once people understand the technology and fine tune it to reduce those false positives. It is unknown

exactly where AI in cybersecurity lays on this graph, but it definitely is not in the plateau of productivity just yet. Afraz Jaffri, Director Analyst at Gartner said, "AI innovations continue to deliver big benefits to business and adoptions rates will accelerate in coming years" (Gartner_Inc), suggesting that AI's plateau of productivity is in the near future.

In fact, Garnter predicts that by 2028, AI-driven machines will account for 20% of the global workforce, and 40% of all economic productivity. This shows that this technology is going to continue to grow and also expand into a multitude of other industries. Interviewee 6 also said this and gave examples of other ways their company uses the technology. Additionally, the spending on these technologies is going to rapidly increase. The leaders in the technology world like Amazon, CrowdStrike, Google, IBM, Microsoft, and more are prioritizing investments in AI/ML research. Microsoft spent $1 billion in cybersecurity research and development in the past year and committed to spending $20 billion over the course of the next five years on that same task (Microsoft's security business generates $15 billion annually) (Columbus, 2023). The overall market size for AI specifically in the industry is predicted to be $22.4 billion in 2023 and is expected to reach $60.6 billion by 2028 (Columbus, 2023). This seems like a large amount of money, but it is just over 10% of the Garnter forecast which is that overall information security spending will reach $187 billion in 2023.

As discussed in this thesis, hackers have also evolved and have used the AI/ML technology to their own advantage. Gartner provides insights on this as well. Gartner says that by 2025, these bad actors will have weaponized operational technology environments successfully to cause human casualties. Attacks on operational technology have become more common and more disruptive, forcing some corporations to focus on hazards to humans and the environment, and not just theft of information and money. Gartner also predicts that by 2025, 45% of organizations will experience attacks on their software supply chains, which is 3 times as many as there was in 2021. In fact, 2022 was a record-breaker for the immense volume of cyberattacks, data breaches, phishing scams, and crypto heists. In turn, it is predicted that 65% of the world's population will have personal data covered under modern privacy regulations, which is a major increase from the 10% in 2020 (Drolet, 2023).

Additionally, developments like ChatGPT, which is a generative AI technology, can also aid the hackers in their attacks, just like how it helps businesses and individuals. AI/ML tools like ChatGPT can do anything from designing malicious code that avoids detection, to writing large volumes of customized phishing emails, to finetuning algorithms that are designed to steal privileged access credentials. In fact, a recent survey by Blackberry found that 51% of IT decision-makers expect that ChatGPT will be abused for a successful cyber-attack within the year (Columbus, 2023). Over this year, it is expected for hackers to get a better handle on ChatGPT specifically, and how it can be used to reinforce their skillset. These trends in the hacker world were also found in the interviews and survey. Interviewees 5 and 6 discussed the concept of the attacks being adaptable and modular to be able to avoid detection, and Interviewee 3 discussed the phishing concepts. The survey responses to the malware evolution question also show that people in the industry believe AI can be used by the hackers for a multitude of reasons, all of which were deemed effective due to the varying ranks of the respondents.

LIMITATIONS & IMPLICATIONS FOR FUTURE RESEARCH

Artificial Intelligence is rapidly changing and evolving every day, which makes it difficult to truly understand where it is now, and to predict what could happen in the future. The technology in the industry is evolving so quickly that it may be hard to keep up with it and get the most recent information. This research is limited as only eight interviews were conducted, and 45 complete survey responses were recorded. Of the 45, only 11 were completed by companies who sell this product, so the results were skewed more towards the companies that buy them. There very well could be a difference in opinions and impacts between these two groups that couldn't be fully explored due to the lack of data. It is almost impossible to say that the information found through this research is applicable to the whole cybersecurity industry due to the industry's immense size and complexity, so it is merely only a snippet.

Cybersecurity has a part in every industry: retail, banking, government, manufacturing, healthcare, and much more. Everyone needs to be able to protect their assets from cyber criminals, who try to get any information they can, from wherever they can. More data needs to be collected from a larger number of cybersecurity professionals that range across these

different environments to fully be able to conclude the impact that AI has on security. This is an enormous task, so large that there are entire companies, like Gartner and Forrester, that focus on this research and charge their customers a lot of money to learn their analysis.

Another issue is the competition in the industry. A lot of companies are reluctant to fully disclose where they are with their products due to the possibility of losing a competitive advantage or giving away company secrets, so there is a slight lack of transparency when trying to get information out of these companies.

Between the rapid-changing nature of AI and the competition in the industry, more data needs to be continuously collected to truly follow AI's evolution and impact that it has on cybersecurity.

CONCLUSION

In conclusion, AI has a seemingly positive impact on the industry, though some of the challenges need to be resolved for it to truly be mainstreamed. Currently, the industry is focusing on detecting threats, but as the technology grows it will reach into other areas of security and move towards more automation. Expertise and false positives will become less of a challenge as the technology continues to evolve and advance. Once it is more effective, and people can better understand how it works, it is almost certain that the technology will officially take off in cybersecurity, more than it already has. It may reach a point where it can predict new threats or actively defend against attacks. With that said, hackers will continue to use the technology for their own benefit and advantage, so it is important for security analysts to keep developing these engines, as it is going to be the best bet to beat these criminals in the long run. To finish off with a quote from the VP analyst at Garnter, Katell Thielemann, "The rise of artificial intelligence (AI) is a double-edged sword for CISOs. Enterprises are facing a deluge of automated cyber-attacks, which are exponentially rising in velocity, variety, and complexity. However, AI is simultaneously supporting security teams in detecting and responding to threats, fundamentally changing organizations' defense paradigms." As we move into the future, the cybersecurity industry is going to have to continue to develop and grow these AI tools to stand a chance in this arms race against cyber criminals to see who can best the other.

APPENDICES

Appendix A – Survey Questions

Title of Study: Artificial Intelligence Impact on the Cyber Security Industry
Student Investigator: Lindsey Shearstone
Faculty Supervisors: Francis Varin and Michele Varin

You are invited to participate in a research study of the impact of artificial intelligence on the cybersecurity industry. We hope to learn the current trends and potential future advancements of the technology, and the impact it will have on the overall industry. You were selected as a possible participant in this study because you have knowledge of the technology, or work in the field of interest.

Description, Including Risks and Benefits: If you decide to participate, you will be asked to answer a few questions about what your company is doing in terms of advancing cybersecurity technologies and how it is being implemented into different areas. Additionally, you will be asked about some challenges that may come up in the process, and what may be in store for the future. There are 7 questions, and you only need to take the survey once. It should not take longer than 10 minutes. There are no anticipated risks or discomforts to your participation in this study. Although you may not directly benefit from this research, by participating in this study you will help make the public aware of the current trends of this evolving technology, what is in store for the future, and the impact it may have on them and their privacy.

Confidentiality: Any information obtained in connection with this study will remain confidential and will not be disclosed to the general public in a way that can be traced to you. In any written reports or publications, no participant other than the researchers will be identified, and only anonymous data will be presented.

Compensation: You will not be paid for participating in this study.

Statement that Participation Is Voluntary: Your participation is totally voluntary, and your decision whether or not to participate will not affect your future relations with Bryant University or its employees in any way. If you decide to participate, you are also free to discontinue participation at any time without affecting such relationships. However, it is requested that you notify the researcher of this.

Persons to Contact: If you have any questions, please contact Lindsey Shearstone at lshearstone@bryant.edu. If you have any additional questions later, we will be happy to answer them.

Signature Indicating Informed Consent: Please check the box below if you have decided to participate. This will act as your signature and indicates only that you are at least 18 years of age, have read the information provided above, and are qualified to answer questions on this topic.

- ○ I agree and am qualified to fill out this survey

In what information security area do you think AI/machine learning has the most benefit?

- ○ Threat Detection
- ○ Vulnerability Management
- ○ Training and Awareness
- ○ Automatic Penetration Testing

Please expand (optional)

Please select the below cyber defense mechanisms that can be most enhanced through the application of AI/machine learning.

- ○ Intrusion Detection and Prevention
- ○ Phishing Protection
- ○ Vulnerability Identification
- ○ User Behavior Analytics
- ○ Spam Filtering

Please expand (optional)

Which statement best describes your company?

- ○ We only develop AI/machine learning and sell it as a product
- ○ We develop our own AI/machine learning engine and incorporate it into our own cyber security product that we sell
- ○ We buy an AI/machine learning engine and incorporate it into our own cyber security product that we sell
- ○ We buy an AI/machine learning cyber security product and implement it in our environment

Please expand (optional)

Please rank the following challenges of implementing AI, 1 being the most challenging and 5 being the least challenging.

- False positives
- Cost
- Expertise
- Effectiveness
- Malware complexity

Please expand (optional)

Please rank how cyber criminals can use AI for malicious reasons, 1 being the most effective and 3 being the least effective

AI can be used to identify patterns in computer systems that reveal weaknesses in software or security programs

Use AI to create large numbers of phishing emails to spread malware or collect valuable information.

AI can also be used to design malware that is constantly changing, to avoid detection by automated defensive tools.

Please expand (optional)

How sophisticated do you think AI is currently? (select all that apply)

- ☐ Accomplishes mundane, non-complex tasks
- ☐ Able to analyze and correlate numerous events beyond a human's capability
- ☐ Takes the place of advanced human analysis

Please expand (optional)

What do you think is in store for the future of AI/machine learning in cyber security? (select all that apply)

- ☐ Be able to predict threats
- ☐ Be able to write it's own code and improve itself
- ☐ Replace human analyst
- ☐ Be able to actively defend against threats
- ☐ Conduct a counterstrike against a cyberattacker (hack back)

Please expand (optional)

Appendix B – Figures and Table from Area Most Benefit Survey Question

Figure B1 – Histogram showing the information security area that AI/ML has the most benefit based on surveyed companies that buy a cybersecurity product and implement it in their environment

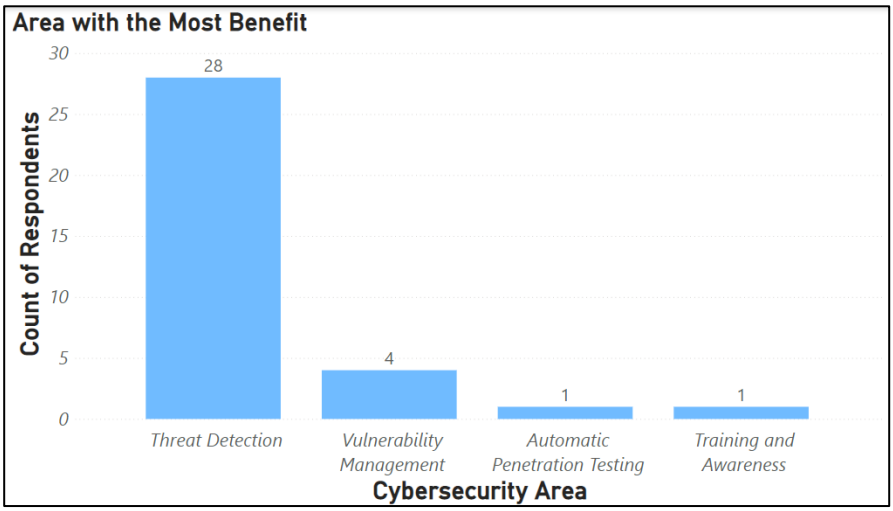

Figure B2 – Histogram showing the information security area that AI/ML has the most benefit based on surveyed companies that sell an AI/ML product

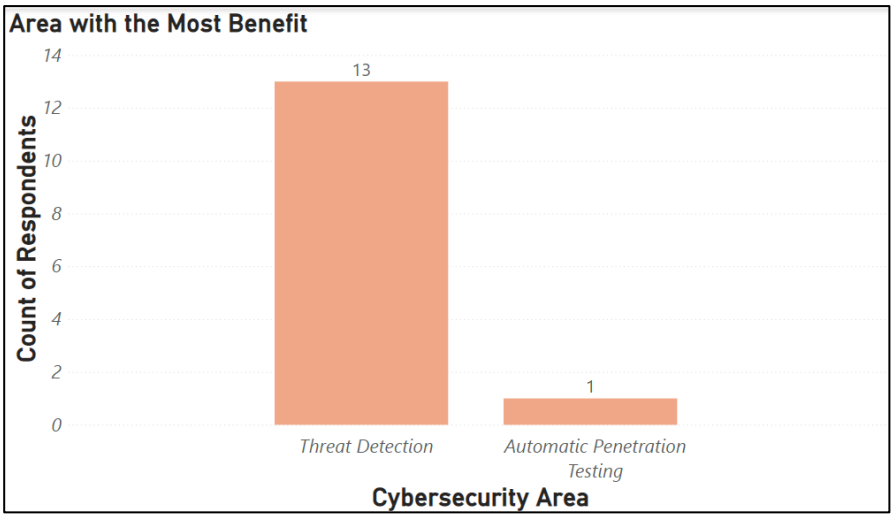

Table B1 – Information security area that AI/ML has the most benefit in according to survey data

AreaMostBenefit	Buy		Sell		Total	
	Count	Percentage	Count	Percentage	Count	Percentage
Threat Detection	28	82.35%	13	92.86%	41	85.42%
Vulnerability Management	4	11.76%	0	0.00%	4	8.33%
Automatic Penetration Testing	1	2.94%	1	7.14%	2	4.17%
Training and Awareness	1	2.94%	0	0.00%	1	2.08%
Total	34	100.00%	14	100%	48	100.00%

Appendix C – Figures and Table from Defense Most Enhanced Survey Question

Figure C1 – Cyber defense mechanism that can be most enhanced through the application of AI/ML according to surveyed companies that buy cybersecurity products

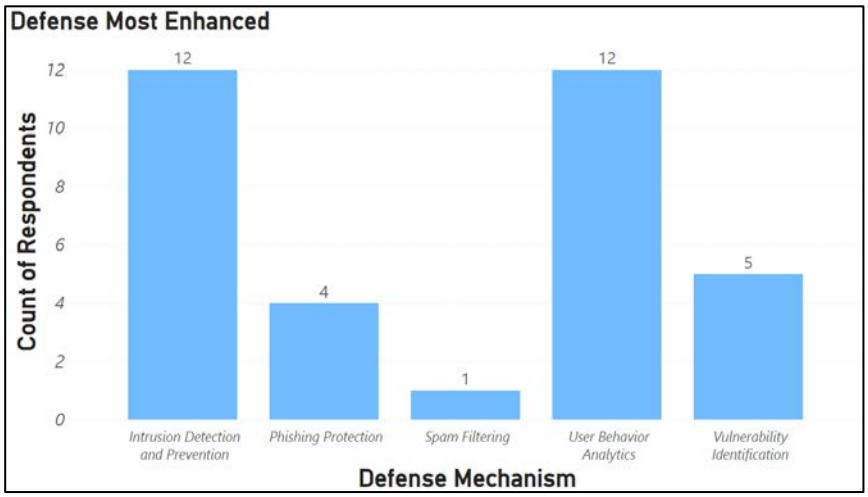

Figure C2 – Cyber defense mechanism that can be most enhanced through the application of AI/ML according to surveyed companies that sell AI/ML products

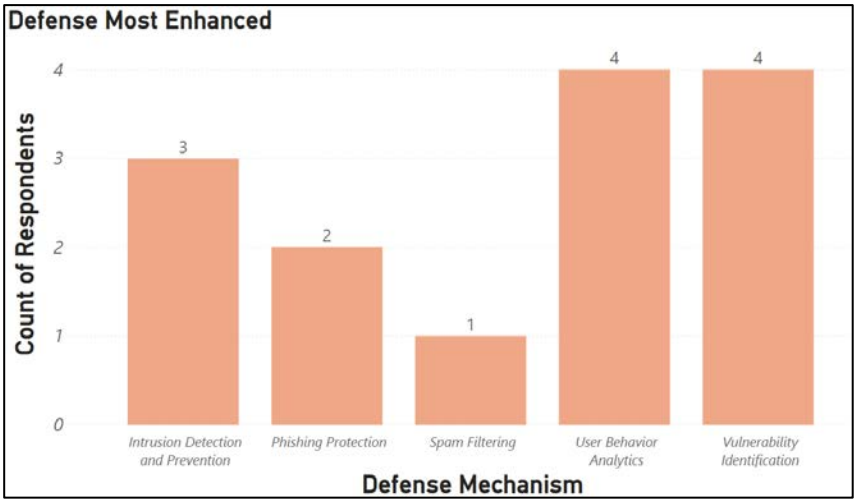

Table C1 – Cyber defense mechanisms that can be most enhanced through the application of AI/ML according to survey respondents

DefenseMostEnhanced	Buy		Sell		Total	
	Count	Percentage	Count	Percentage	Count	Percentage
Intrusion Detection & Prevention	12	35.29%	3	21.43%	15	31.25%
Phishing Protection	4	11.76%	2	14.29%	6	12.50%
Spam Filtering	1	2.94%	1	7.14%	2	4.17%
User Behavior Analytics	12	35.29%	4	28.57%	16	33.33%
Vulnerability Identification	5	14.71%	4	28.57%	9	18.75%
Total	34	100.00%	14	100%	48	100.00%

Appendix D – Figures and Table for Challenge Ranking Survey Question

Figure D1 – Histograms representing where all survey respondents ranked each challenge of implementing AI, with 1 being the most challenging, and 5 being the least

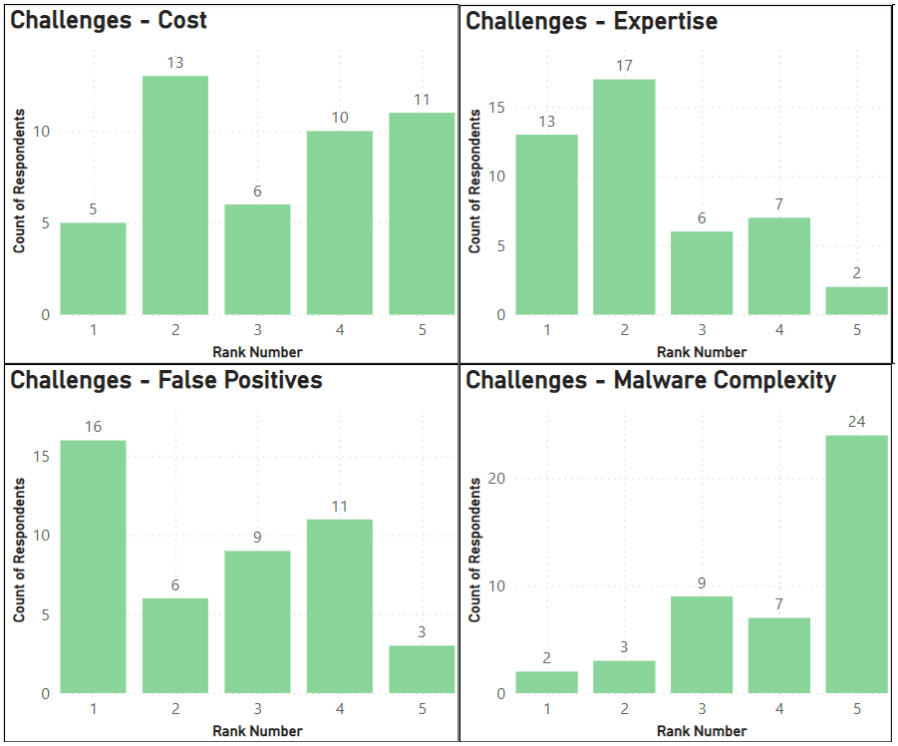

Figure D2 – Histogram representing where all survey respondents ranked the challenge of effectiveness when implementing AI, with 1 being the most challenging

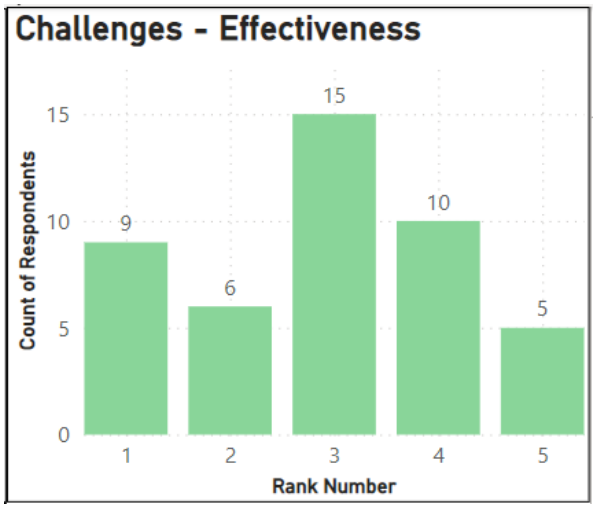

Figure D3 – Histogram representing where the survey respondents that buy products ranked the challenge of effectiveness when implementing AI, with 1 being the most challenging

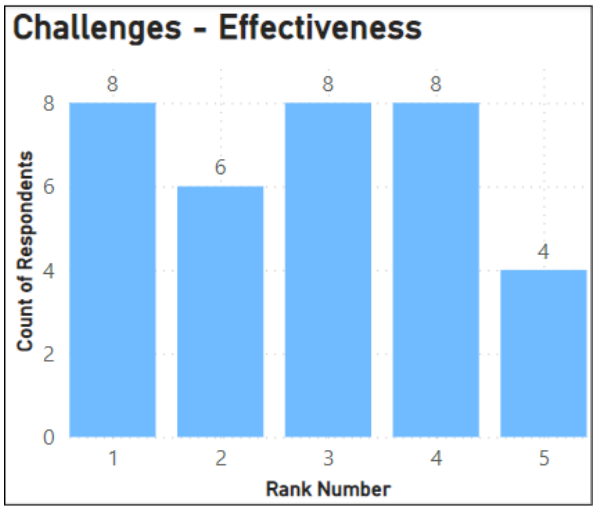

Figure D4 – Histograms representing where the survey respondents that buy products ranked each challenge of implementing AI, with 1 being the most challenging, and 5 being the least

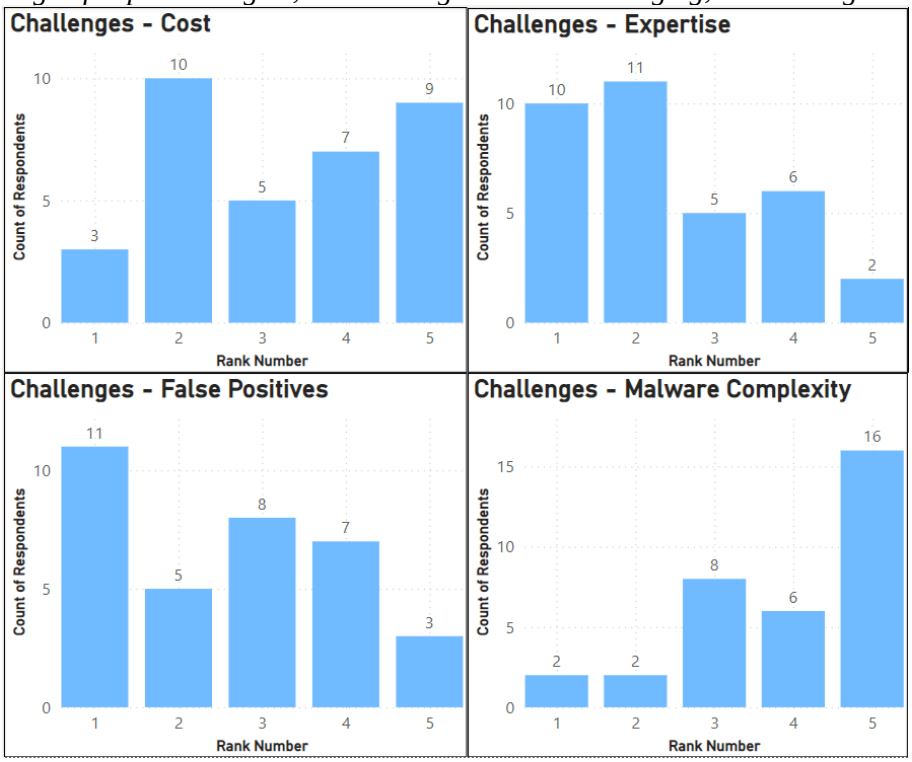

Figure D5 – Histograms representing where the survey respondents that sell products ranked each challenge of implementing AI, with 1 being the most challenging, and 5 being the least

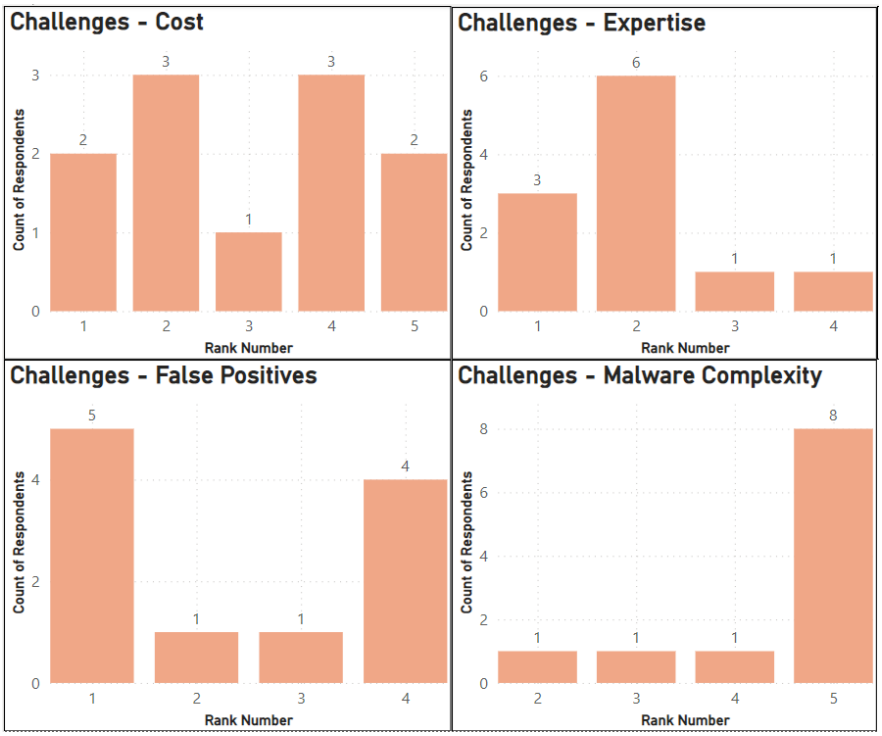

Figure D6 – Histogram representing where the survey respondents that sell products ranked the challenge of effectiveness when implementing AI, with 1 being the most challenging

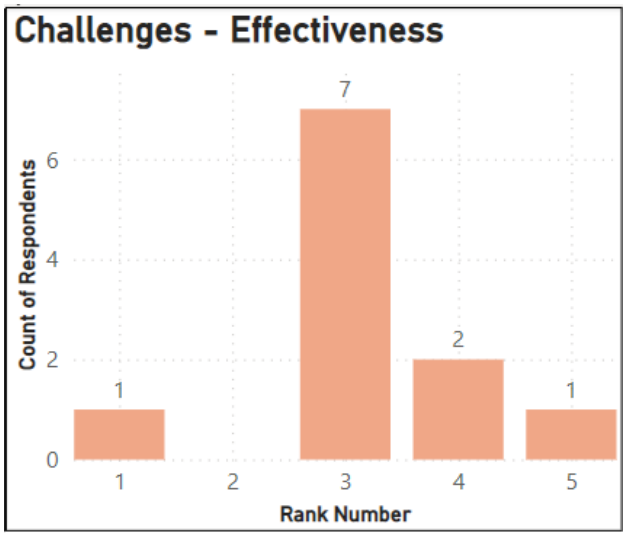

Table D1 – Shows the mean and standard deviation of each challenge that was ranked from 1 to 5, looking at segmented data and the overall data, with 1 being the most challenging

Challenges	Buy		Sell		Total	
	Mean	SD	Mean	SD	Mean	SD
False Positives	2.588235294	1.37	2.36	1.43	2.53	1.38
Cost	3.264705882	1.38	3.00	1.48	3.20	1.39
Expertise	2.382352941	1.26	2.00	0.89	2.29	1.18
Effectiveness	2.823529412	1.36	3.18	0.98	2.91	1.28
Malware Complexity	3.941176471	1.23	4.45	1.04	4.07	1.19

Appendix E – T-tests in Excel for Challenge Ranking Survey Question

T-Test Against Challenges - All Data	**Alpha**
Degrees of Freedom (df) = 45 + 45 - 2 = 88	0.05 / 10 = 0.005
Ho: False Positives - Cost = 0 Ha: False Positives - Cost ≠ 0 t-statistic = -2.285968768 p-value = 0.024657847 Fail to reject	Ho: Cost - Expertise= 0 Ha: Cost - Expertise ≠ 0 t-statistic = 3.350174955 p-value = 0.001191 Reject
Ho: False Positives - Expertise = 0 Ha: False Positives - Expertise ≠ 0 t-statistic = 0.905028686 p-value = 0.36792134 Fail to reject	Ho: Cost - Malware Complexity = 0 Ha: Cost - Malware Complexity ≠ 0 t-statistic = -3.169960354 p-value = 0.002098008 Reject
Ho: False Positives - Effectiveness = 0 Ha: False Positives - Effectiveness ≠ 0 t-statistic = -1.350893738 p-value = 0.180192782 Fail to reject	Ho: Expertise - Malware Complexity = 0 Ha: Expertise - Malware Complexity ≠ 0 t-statistic = -7.10266627 p-value = 3.02187E-10 Reject
Ho: Cost - Effectiveness = 0 Ha: Cost - Effectiveness ≠ 0 t-statistic = 1.026428881 p-value = 0.30750265 Fail to Reject	Ho: False Positives - Malware Complexity = 0 Ha: False Positives - Malware Complexity ≠ 0 t-statistic = -5.646674376 p-value = 1.97336E-07 Reject
Ho: Expertise - Effectiveness = 0 Ha: Expertise - Effectiveness ≠ 0 t-statistic = -2.401783061 p-value = 0.018418703 Fail to Reject	Ho: Effectiveness - Malware Complexity = 0 Ha: Effectiveness - Malware Complexity ≠ 0 t-statistic = -4.4345749 p-value = 2.65511E-05 Reject

Figure E1 – T-tests comparing the means of the ranked challenges from all the survey respondents

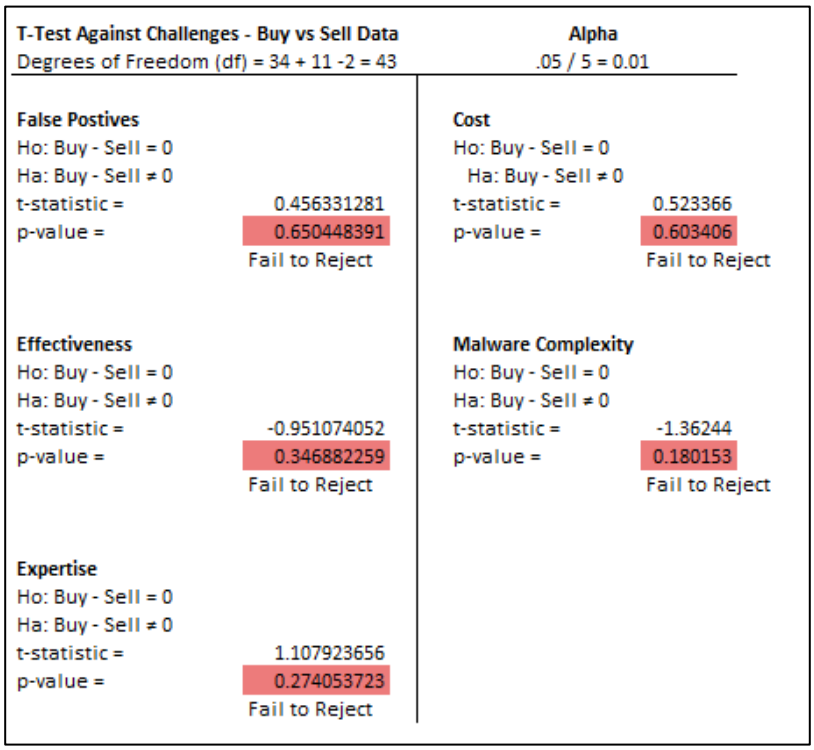

T-Test Against Challenges - Buy vs Sell Data		Alpha	
Degrees of Freedom (df) = 34 + 11 -2 = 43		.05 / 5 = 0.01	
False Postives		**Cost**	
Ho: Buy - Sell = 0		Ho: Buy - Sell = 0	
Ha: Buy - Sell ≠ 0		Ha: Buy - Sell ≠ 0	
t-statistic =	0.456331281	t-statistic =	0.523366
p-value =	0.650448391	p-value =	0.603406
	Fail to Reject		Fail to Reject
Effectiveness		**Malware Complexity**	
Ho: Buy - Sell = 0		Ho: Buy - Sell = 0	
Ha: Buy - Sell ≠ 0		Ha: Buy - Sell ≠ 0	
t-statistic =	-0.951074052	t-statistic =	-1.36244
p-value =	0.346882259	p-value =	0.180153
	Fail to Reject		Fail to Reject
Expertise			
Ho: Buy - Sell = 0			
Ha: Buy - Sell ≠ 0			
t-statistic =	1.107923656		
p-value =	0.274053723		
	Fail to Reject		

Figure E2 – T-tests comparing the means between the companies that buy products versus the companies that sell the products for each ranked challenge

T-Test Against Challenges - Buy Data	Alpha
Degrees of Freedom (df) = 34 + 34 - 2 = 66	0.05 / 10 = 0.005

Ho: False Positives - Cost = 0	Ho: Cost - Expertise= 0
Ha: False Positives - Cost ≠ 0	Ha: Cost - Expertise ≠ 0
t-statistic = -2.027897308	t-statistic = 2.760306495
p-value = 0.046610872	p-value = 0.007468299
Fail to reject	Fail to reject
Ho: False Positives - Expertise = 0	Ho: Cost - Malware Complexity = 0
Ha: False Positives - Expertise ≠ 0	Ha: Cost - Malware Complexity ≠ 0
t-statistic = 0.645148213	t-statistic = -2.136333674
p-value = 0.521067577	p-value = 0.036364663
Fail to reject	Fail to Reject
Ho: False Positives - Effectiveness = 0	Ho: Expertise - Malware Complexity = 0
Ha: False Positives - Effectiveness ≠ 0	Ha: Expertise - Malware Complexity ≠ 0
t-statistic = -0.710139419	t-statistic = -5.172270387
p-value = 0.480119381	p-value = 2.33971E-06
Fail to Reject	Reject
Ho: Cost - Effectiveness = 0	Ho: False Positives - Malware Complexity = 0
Ha: Cost - Effectiveness ≠ 0	Ha: False Positives - Malware Complexity ≠ 0
t-statistic = 1.329449632	t-statistic = -4.279946651
p-value = 0.188276715	p-value = 6.18047E-05
Fail to Reject	Reject
Ho: Expertise - Effectiveness = 0	Ho: Effectiveness - Malware Complexity = 0
Ha: Expertise - Effectiveness ≠ 0	Ha: Effectiveness - Malware Complexity ≠ 0
t-statistic = -1.390356192	t-statistic = -3.556191336
p-value = 0.16909172	p-value = 0.000701861
Fail to Reject	Reject

Figure E3 – T-tests comparing the means of the ranked challenges from the survey respondents who buy the products

T-Test Against Challenges - Sell Data	**Alpha**
Degrees of Freedom (df) = 34 + 34 - 2 = 66	0.05 / 10 = 0.005

Ho: False Positives - Cost = 0 Ha: False Positives - Cost ≠ 0 t-statistic = -1.023234356 p-value = 0.309932987 Fail to reject	Ho: Cost - Expertise= 0 Ha: Cost - Expertise ≠ 0 t-statistic = 1.914854216 p-value = 0.059846696 Fail to reject
Ho: False Positives - Expertise = 0 Ha: False Positives - Expertise ≠ 0 t-statistic = 0.71383061 p-value = 0.477848764 Fail to reject	Ho: Cost - Malware Complexity = 0 Ha: Cost - Malware Complexity ≠ 0 t-statistic = -2.666666667 p-value = 0.009625304 Fail to Reject
Ho: False Positives - Effectiveness = 0 Ha: False Positives - Effectiveness ≠ 0 t-statistic = -1.561972803 p-value = 0.123076993 Fail to Reject	Ho: Expertise - Malware Complexity = 0 Ha: Expertise - Malware Complexity ≠ 0 t-statistic = -5.948810765 p-value = 1.13559E-07 Reject
Ho: Cost - Effectiveness = 0 Ha: Cost - Effectiveness ≠ 0 t-statistic = -0.339031752 p-value = 0.735662161 Fail to Reject	Ho: False Positives - Malware Complexity = 0 Ha: False Positives - Malware Complexity ≠ 0 t-statistic = -3.921467452 p-value = 0.000212054 Reject
Ho: Expertise - Effectiveness = 0 Ha: Expertise - Effectiveness ≠ 0 t-statistic = -2.95149796 p-value = 0.004374625 Reject	Ho: Effectiveness - Malware Complexity = 0 Ha: Effectiveness - Malware Complexity ≠ 0 t-statistic = -2.958039892 p-value = 0.004293619 Reject

Figure E4 – T-tests comparing the means of the ranked challenges from the survey respondents who sell AI/ML products

Appendix F – T-tests in Excel for Cyber Criminal Use Ranking Survey Question

T-Test Against Cyber Criminal Uses - All Data	**Alpha**
Degrees of Freedom (df) = 45 + 45 - 2 = 88	0.05 / 3 = 0.0167

Ho: Identify Patterns - Create Phishing Emails = 0
Ha: Identify Patterns - Create Phishing Emails ≠ 0
t-statistic = -0.532795428
p-value = 0.595518096
Fail to Reject

Ho: Create Phishing Emails - Constantly Changing = 0
Ha: Create Phishing Emails - Constantly Changing ≠ 0
t-statistic = 1.761294
p-value = 0.081662
Fail to Reject

Ho: Identify Patterns - Constantly Changing = 0
Ha: Identify Patterns - Constantly Changing ≠ 0
t-statistic = 1.292954958
p-value = 0.19941084
Fail to Reject

Figure F1 – T-tests comparing the means of the ranking of Cybercriminal uses of AI for malicious reasons from all survey respondents

T-Test Against Cyber Criminal Uses - Buy vs Sell	**Alpha**
Degrees of Freedom (df) = 45 + 45 - 2 = 88	0.05 / 3 = 0.0167

Identify Patterns
Ho: Buy - Sell = 0
Ha: Buy - Sell ≠ 0
t-statistic = -0.69184429
p-value = 0.492754089
Fail to Reject

Constantly Changing
Ho: Buy - Sell = 0
Ha: Buy - Sell ≠ 0
t-statistic = 0.016268445
p-value = 0.987095483
Fail to Reject

Create Phishing Emails
Ho: Buy - Sell = 0
Ha: Buy - Sell ≠ 0
t-statistic = 0.644413265
p-value = 0.522730488
Fail to Reject

Figure F2 – T-tests comparing the means between the buy and sell survey responses for each rank option for how Cybercriminals can use AI for malicious reasons

Appendix G – Histograms and Table for Cyber Criminal Use Ranking Survey Question

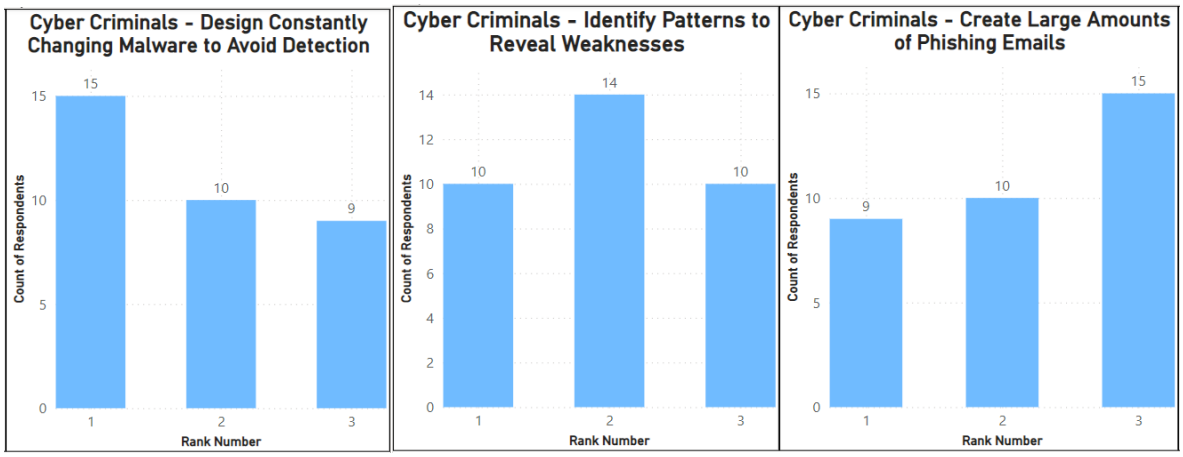

Figure G1 – Histograms that show how many respondents that sell products ranked each option of how Cybercriminals can use AI for malicious reasons, 1 being the most effective

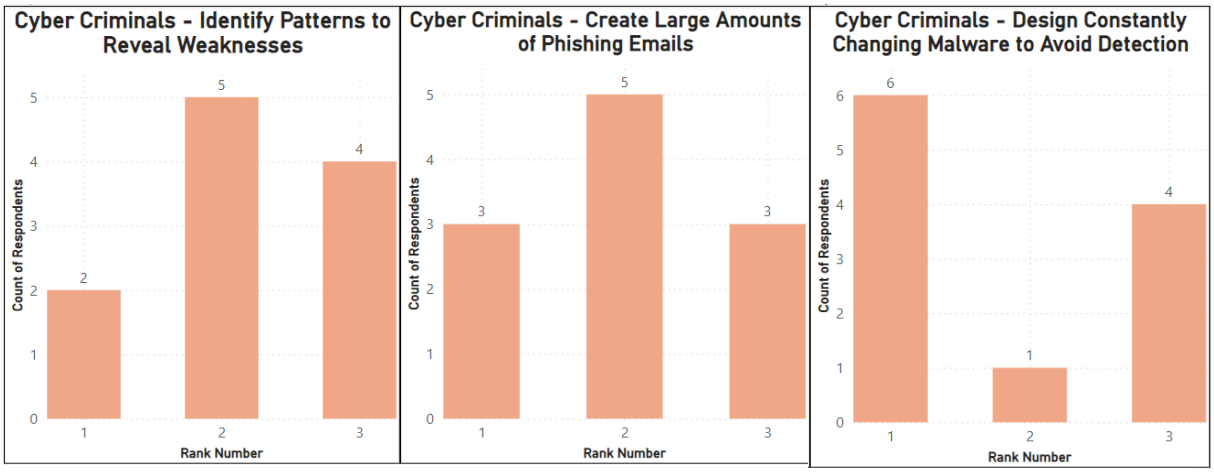

Figure G2 – Histograms that show how many respondents that buy products ranked each option of how Cybercriminals can use AI for malicious reasons, 1 being the most effective application

Table G1 – Shows the mean and standard deviation of each Cybercriminal use of AI that was ranked from 1 to 3, looking at segmented data and the overall data, with 1 being the most effective

Cyber Criminals	Buy		Sell		Total	
	Mean	SD	Mean	SD	Mean	SD
Identify Patterns	2	0.78	2.18	0.75	2.04	0.77
Create Phishing Emails	2.176470588	0.83	2.00	0.77	2.13	0.81
Constantly Changing	1.823529412	0.83	1.82	0.98	1.82	0.86

Appendix H – Histograms and Table for Survey on where AI is Currently

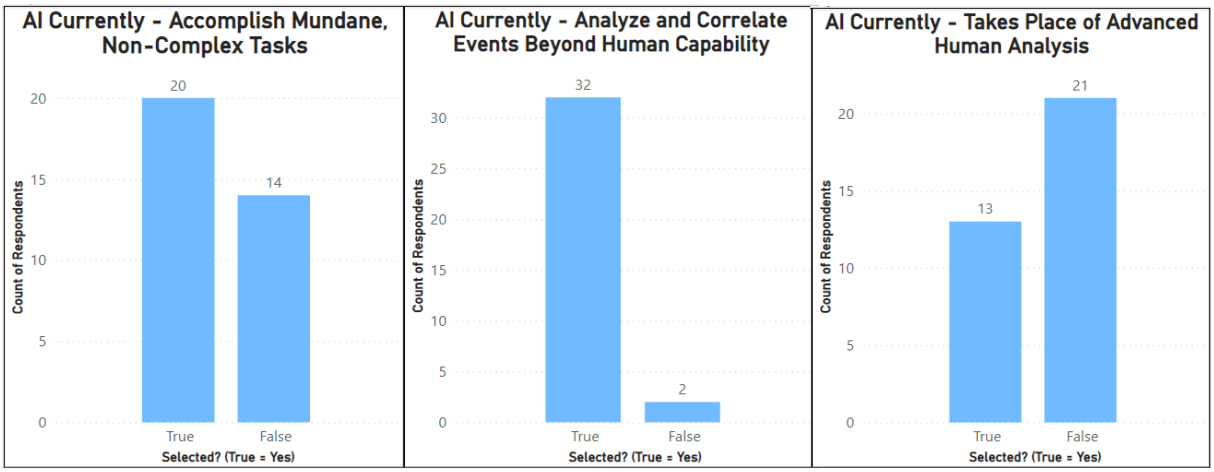

Figure H1 – Histograms that show the number of survey respondents that buy products that selected (True) each task that states where AI currently stands in the industry.

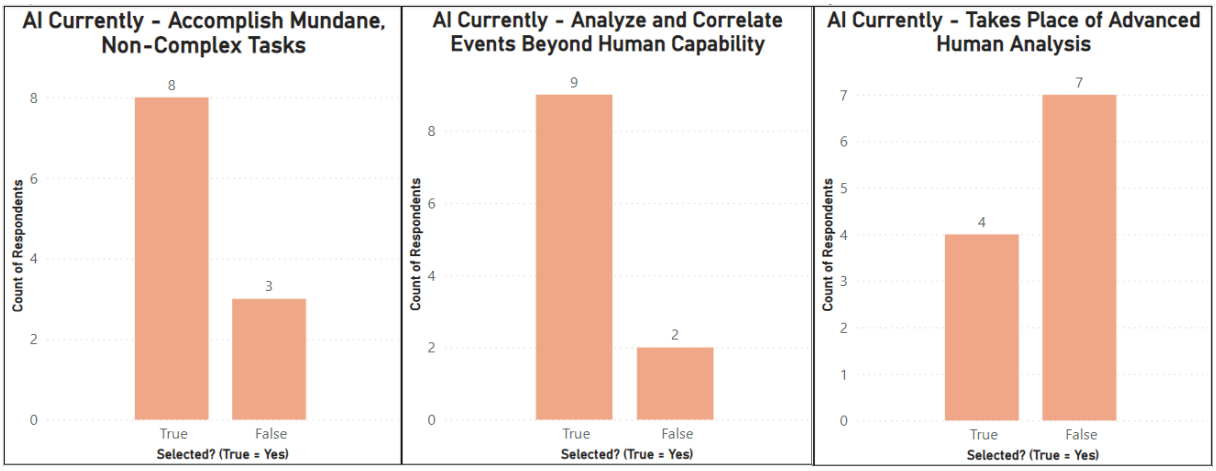

Figure H2 – Histograms that show the number of survey respondents that sell products that selected (True) each task that states where AI currently stands in the industry.

Table H1 – Shows the number of responses that selected (and the percentage) each option as to how sophisticated AI is today

AI Current	Buy		Sell		Total	
	Selected	Percentage	Selected	Percentage	Selected	Percentage
Mundane Tasks	20	58.82%	8	72.73%	28	62.22%
Analyze Beyond Human	32	94.12%	9	81.82%	41	91.11%
Advanced Human Analysis	13	38.24%	4	36.36%	17	37.78%

Table H2 – Shows the number of responses that did not select (and the percentage) each option as to how sophisticated AI is today

AI Current	Buy		Sell		Total	
	Not-Selected	Percentage	Not Selected	Percentage	Not Selected	Percentage
Mundane Tasks	14	41%	3	27.27%	17	37.78%
Analyze Beyond Human	2	6%	2	18.18%	4	8.89%
Advanced Human Analysis	21	62%	7	63.64%	28	62.22%

Appendix I – Histograms and Table for Survey on where AI is Headed in the Future

Figure I1 – Histograms show the number of total respondents that selected each potential task that AI could perform as it evolves in the future

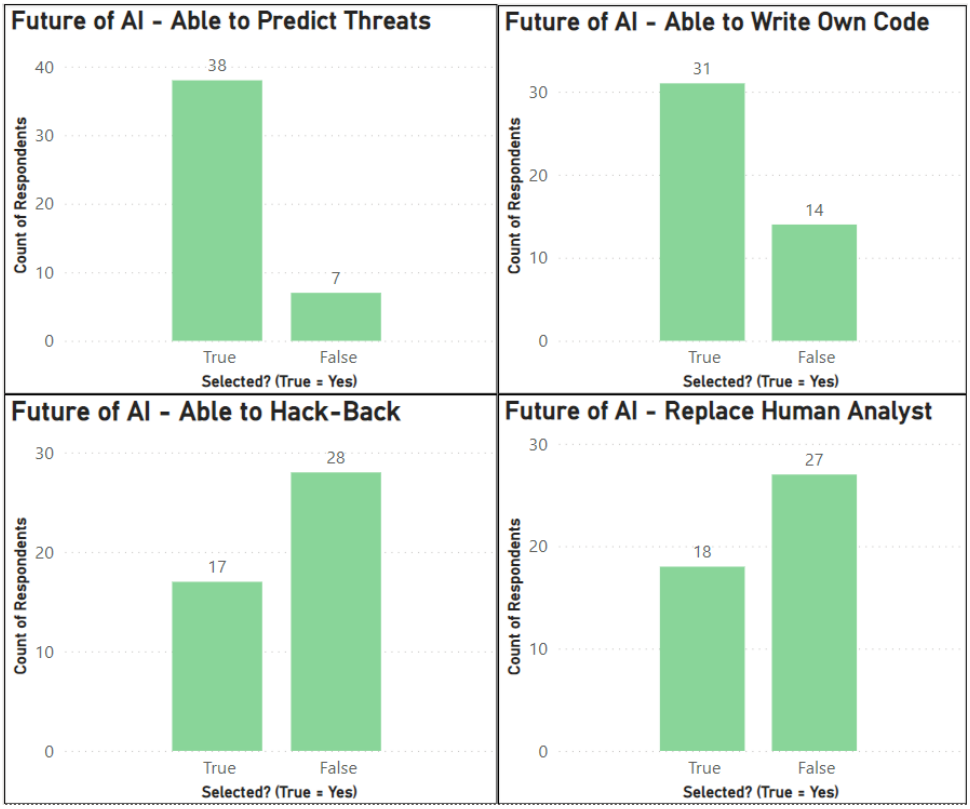

Figure I2 – Histogram shows the number of total respondents that selected that AI will be able to actively defend against threats as it evolves in the future

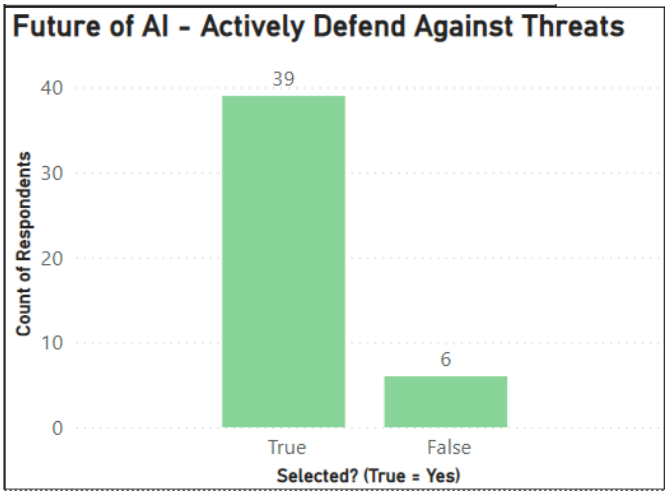

Figure I3 – Histograms show the number of respondents that buy products that selected each potential task that AI could perform as it evolves in the future

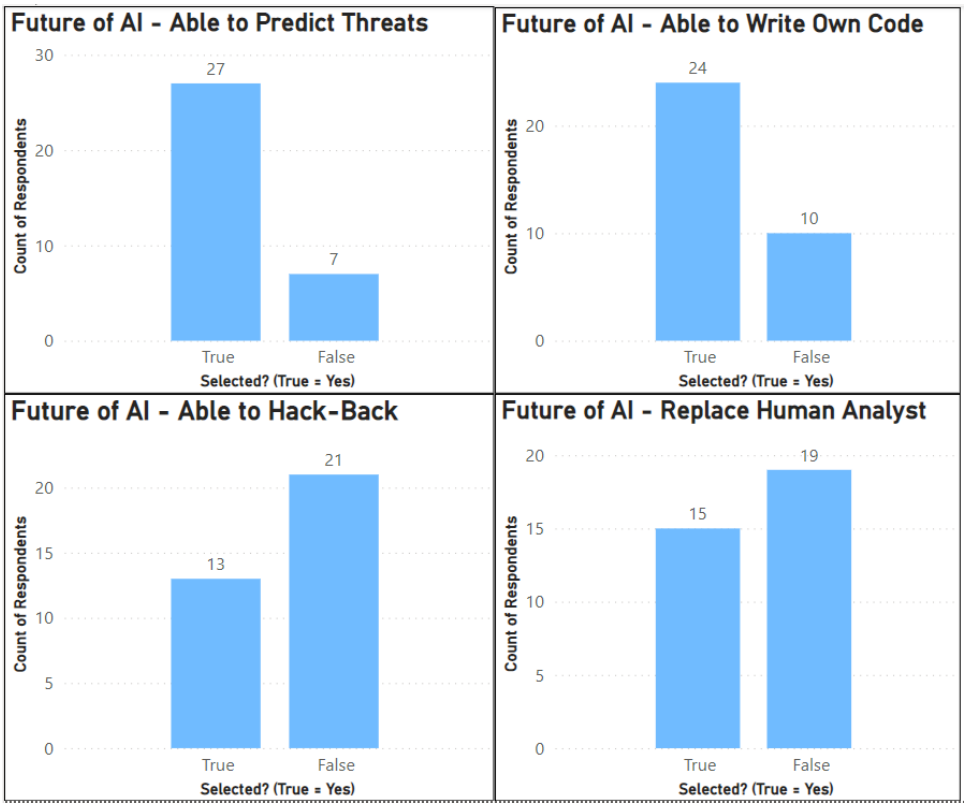

Figure I4 – Histogram shows the number of respondents that buy products that selected that AI will be able to actively defend against threats as it evolves in the future

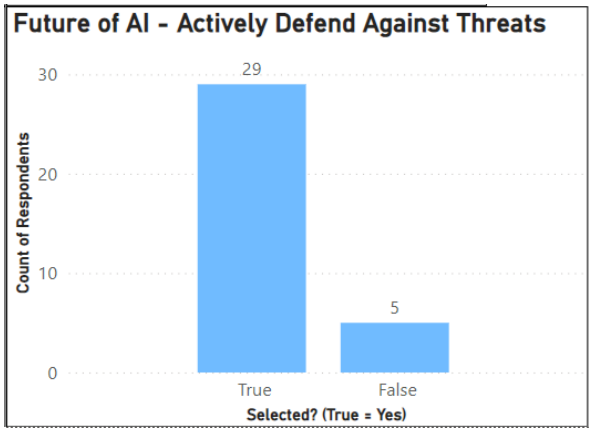

Figure I5 – Histograms show the number of respondents that sell products that selected each potential task that AI could perform as it evolves in the future

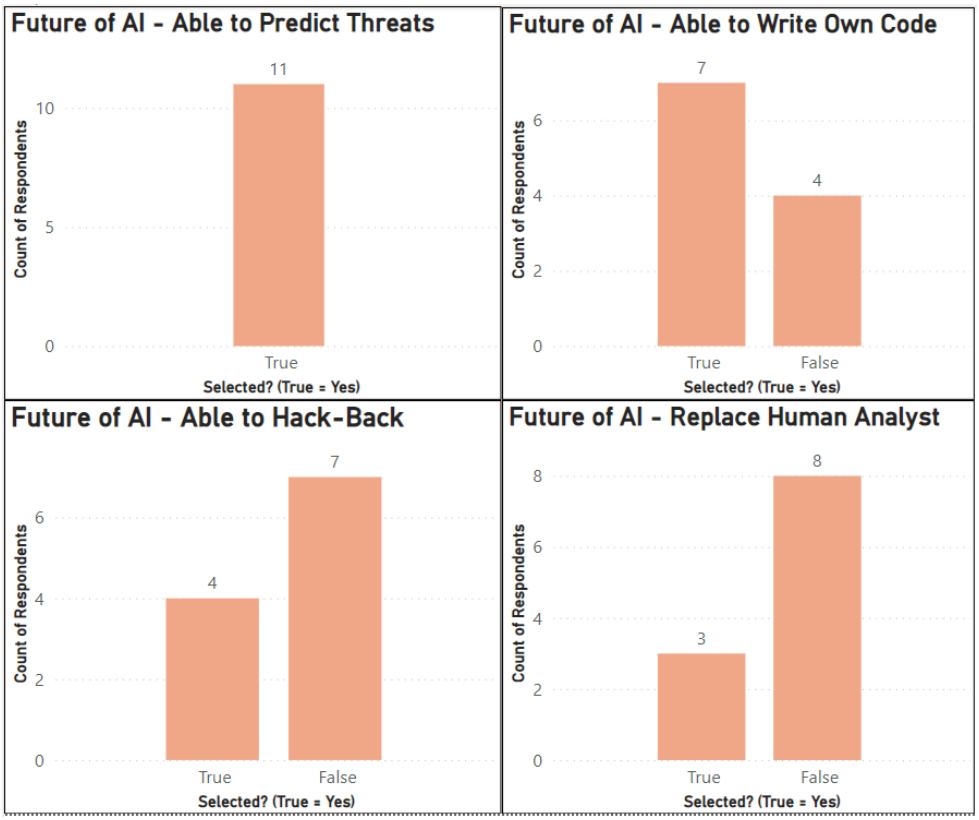

Figure I6 – Histogram shows the number of respondents that sell products that selected that AI will be able to actively defend against threats as it evolves in the future

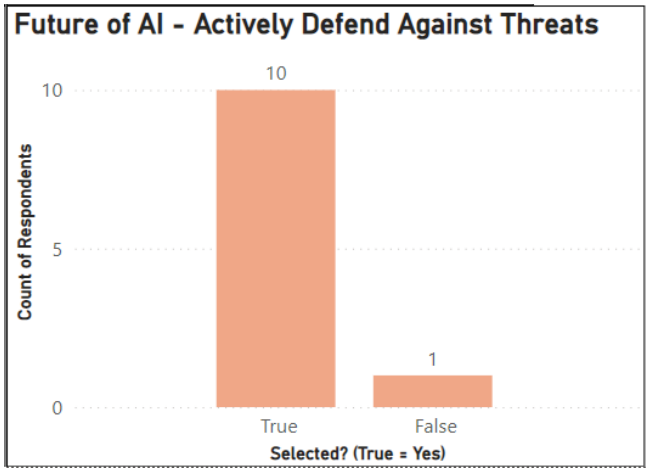

Appendix J – Data Collected from Interviews

Information Given Before Questions were Asked

- Signature based vs. behavioral based
    - Signatures (hashes of known malicious files)
        - How did it get there, why is it there, who brought it, what remediation needs to be done
    - Behavioral (AI, much larger data sets – harder for humans to write specific detection string to query)
- User Behavior Analytics (UBA) – thousands of users to analyze, next to impossible to do it individually/manually, AI can do it
- AI – learns environment, takes data, and normalizes it, then can make judgements on where the abnormal behavior is being created, and then send an alert to your soc team or tier 2/3 analyst to take a look at it
- Penetration Tester
    - Hacking companies to test their defenses
- Gartner Analyst
    - Industry analyst who rates and evaluates markets, industries, and vendors
        - AI was a big up and coming topic there – a lot of vendors are saying they use AI now
- AI helps make the security program more agile on the different pillars within the typical security team

1. **How is Artificial Intelligence being implemented into different products at your company? Do you collaborate with companies focused on AI advancement and products, or do you develop your own technology?**

Interview 1:

- Mostly vendor-based relationships
    - Lean on their partners (can't disclose products)

- AV [Antivirus] products use AI to make determinations on behaviors of processes
    - Double edged sword
        - Get utilities that have process, arguments, and commands in them that can be used for potentially malicious activity – they get flagged and stopped (can stop user from completing their day-to-day job)
        - Also, can be a legitimate activity, and because it can be used maliciously sometimes it will get flagged and stop the process even though it was a legitimate process anyways
- Pain Point:
    - It is good to have AI behind it, but it can have a lot of false positives
        - Judgment calls from the human at that point

Interview 2:

- AI is playing a huge role in the cyber field
    - Started looking at next gen antivirus
        - Next gen [next generation] meaning that it uses AI to detect malware viruses, ransomware, etc.
- Vendor based relationships – doesn't develop AI

Interview 3:

- No bullet proof security, the way you do it is to put layer after layer and hope that one of these layers will plug the hole
    - Have to take a proactive approach in the application development as well as when building an infrastructure to make sure it adheres to the best security practice
    - Any time they create an infrastructure, they go through this practice
- New project – building a B-B and a B-C website (ecommerce website)
    - A lot of infrastructure is being built in Microsoft Azure, should have gone live around October/November 2022
    - Brings in a security firm to do a full assessment of the infrastructure

        - Let the developers and infrastructure engineers run free to create anything they want, and create it based on best practice
        - Bring in a security infrastructure company to evaluate and recommend on what to do
- Security starts from knowing your assets – know what you own (good asset management is the first step in cyber security, a lot of people don't recognize that – you need to know exactly what you have, then you can understand how to protect it)
- Need protection on the desktop level (end point protection), the network level, and for e-mail, which is the biggest threat
- End point protection
    - Switch solutions every 3-4 years
    - Used to be with Norton security, then switched to Carbon Black
    - History of endpoint protection
        - Used to be signature based, knowing the virus from its signature
            - First MacAfee and Norton used to look at the Assembly code and binary code. The code looks at the parent and it can recognize if this file/executable is a virus
    - A few years ago, when AI became available, everyone turned to behavioral protection – looking at the behavior of the executable/document, and based on the behavior, determines whether or not it is a virus.
        - Carbon Black is one of the leading companies in the AI space, and one of the first that really adapted the behavioral piece, and then everywhere else started catching up with them
    - A couple of years ago, switched to Sophos Intercept X, who also uses AI to determine if it is a virus
        - Provides literature about how they use AI to determine if it is or isn't a virus, how they score it, etc.
        - Give their clients a sandbox/warranty box where they can test things and look at the solutions
- They key is to make sure solutions are up to date

- Network level
    - SIEM (Security Information and Event Management) solution
        - One of the pillars you need for your environment
        - Been around for about 15 years
            - Look at the event, set your policy on the events, and then the solution can identify based on your policy what is good and bad
        - One of the leading companies that handles this is called Logarithm
            - According to the magic quadrant, they are one of the leads in visionaries of that space
        - What it does
            - Takes every log and every event that is produced by every single network node you have connected to your network and forwards it to the SIEM solution
                - Example: They can tell the solution that if someone is trying to log in with credentials 3 times on Monday, alert them. So, when someone comes in and they are trying to authenticate their domain and they hit the domain controller, and they try to authenticate and fail 3 times, every failure is being forwarded through the log of that domain controller operating system, but because they installed a small client that forwards these events to the SIEM solution, the SIEM solution becomes aware of it and based on the set policy, alerts them immediately. The alert will pop up and say someone is trying to log in as so and so.
                    - Help desk will take the alert and contact the person to see if it is just them failing their log in 3 times, or if it is a hacker attempting to get in
                - Can also track what devices you sign in on

        - Normally companies only allow login on the company laptop (little variety), so they know the devices that are normally used. So, if there is a login on a new device, the SIEM solution will immediately alert them
    - The logs are huge, ideally you have millions and millions of lines – the tools weed through them
        - Now SIEM solution does what endpoint protection does, they are cloud-based
            - It uses your system for data, does machine learning on your events, and does the same thing from other systems from other clients, and from things that are going on in the world
            - SIEM solution went from just collecting logs and trying to sort what's good and what's bad based on the policy, to not just basing it on the policy, but it also basing it on real world events that are going on

- Email
    - Only way to protect yourself is dissipating filters
    - Spam Filters
        - In the beginning (Symantec and Barracuda started like this but have evolved to AI)
            - See who it is coming from (if they are on a blacklist)
            - See what brand the email is coming from (AOL, Gmail, etc.)
            - Looks at the subject of the email
            - Some looked for certain key words in the content (job opportunity, mortgage, etc.)
        - Now
            - Looking at behavior
                - Scanning the attachments

        - Right now, the number one attack is cryptologic (attacks networks straight from the e-mail attachment)
            - Cryptologic is encrypting everything on your network
                - 99% of the time, it is an email attachment – Employees don't know what they are doing and open the attachment, and then the virus is in the network and on the PC
    - Use the standard spam filter of Microsoft Office 365, but on top of that, use a product from a company called Mimecast
        - Mimecast is huge on AI
            - Every email goes through their engine before going to the recipient
            - Engine looks at every attachment, scans it, if it found a macro in an excel or word document box, or if it's a zip file, they open it in safe mode to make sure it's okay (zip files are a popular attack)
                - Engine is not just based on policy, it is machine learning (inputs from other clients, because it is also cloud based) – when there is a phishing e-mail, they will add it into the knowledge of the machine learning and flag it as a fishing e-mail for everyone using the solution, not just those on that network
- Used to learn how to hack things
    - Writing assembly code – learned how to write viruses
    - Attacks in the past used to require certain skills (reverse binary code, assembly, C plus), which are not allowed to be taught today
    - In today's world, you still need the skills, but you can also just do it from sending an e-mail attachment

        - Can rent an SMTP server for fractional costs or even for free, and start spamming people with emails with infected malware
            - Some attacks aren't that sophisticated – social attacks (just have to call somebody)
- Social Attacks
    - An employee can get a call saying, "this is your help desk, we saw you are having a problem logging in, can you please verify your user ID and password" and some people will give away their password freely
        - They have policies in place to prevent this (help desk will not call and ask for password) and they educate their employees (security training program online courses) on what to look for
- Penetration test
    - Internal test every week
    - Hire a company to hack their system and give them back a report
        - They do this 1-2 times a year
    - Use a company called Qualys
        - Qualys is a famous company that has a scanning engine
            - Do this almost every month
            - Scan their entire network to look for certain things and get back results (what's going on, if there is a vulnerability)
- Firewalls
    - Traditional Firewalls
        - They still use them – still a valid solution
        - Setting network rules (who can talk to the outside, who can talk to the inside, which machines have access, where is the DMZ (demilitarized zone), etc.)
    - Cisco
        - AI engine, they call it FMC (fire management control)
        - Still have the network rules, but they now build on top of intrusion detection (AI cloud-based engine that's connected to Cisco)

Interview 4:

- Develop their own products
- On a security perspective, organizations are trying to understand what is really happening across their organization, employees, devices, and basically all avenues
    - Security teams have a responsibility to stop external threats from hackers that are trying to get what they call the crown jewels.
        - Ex: At a hospital, crown jewels would be electronic medical records, or for a bank it would be credit card information
    - Security operations teams in today's businesses have the responsibility to really cover a lot – monumental task of finding things and not just finding them but stopping them.
        - Because of this, they need help from different tools to be able to do this.
- In their case and across the industry, there are tools that stop a lot of different types of things (firewalls stop external threats, different endpoint management or EDR (endpoint detection response) stops a lot of user-based detection)
- They pull the feeds from all the security devices, the routers, switches, and lot of the networking type devices
    - Do a lot of machine learning type detection based on behaviors
        - If you see behaviors across the organization or the security that could be a potential security threat, you want to alert on that and send an alarm to the security administrators to get their attention
    - In the end they are a tool, they are not going to essentially be a silver bullet that cures every kind of possible concern that you ever have
        - They bring alerts to the forefront
    - Their product has machine learning functionalities, but they don't delve a lot into the artificial intelligence side of things, it is more machine learning

Interview 5:

- Progression over a period of time
- Most products for the last 20-30 years have been based on signatures

  - Whether it be network detection or antivirus, historically it's always been signature based (take a hash value of a file and check to see if it is in a database full of millions of hash values, if it was, it was deemed a virus and would be shut down)
  - Network intrusion detection systems would monitor network traffic and then look for things that break signatures (based on hex values of network traffic)
    - If it saw a particular set of hex code in a network packet, it would identify it and drop the traffic – very static, if it's in the list, then stop it
- Where AI/machine learning has come in
  - AI doesn't want to know about the signature (it is important, but easy to get around now), it analyzes behavior (identify patterns of behavior to detect malware based off the behavior and not the signature, or at least a combination of both)
    - The idea is to have a higher detection rate and a lower false positive rate
- The company bases all of their offerings off compliance or framework
  - From a program-based approach – you can't just choose a product
    - When looking at it from a program-based approach, or a program level you start to realize that the products themselves don't really matter as much as long as you know what control they are satisfying and how well they can satisfy it (also what it takes to actually implement and monitor)
      - Then you start to put technology in place for services or resources in other areas
    - "You're better off buying the cheapest thing off the shelf at Walmart and putting some consistent and capable resources behind it than buying the most expensive product you could find and understaffing it"
    - From a high level, every security program has to consider all sorts of things related to policies, procedures, documentation, training, compliance aspects of it, framework, and legal components

- The company doesn't onboard or purchase new vendors often because they have a lot of people that are delivering their product (have to be trained and certified)
    - Comes down to business discussion more than a technical one (how well are they going to work with us, what sort of support resources are they going to have available, how does their pricing work, how quickly can we turn pricing around to customers)
    - Everything they use is going to be on the Gartner top 5 list or Forrester top 5 (industry analytics lists – two different companies that analyze different product sets like the top products, their pros and cons, and they do this year after year)
        - If you are choosing a top five vendor off of those lists for any particular category, they are going to cover the bases, and then their company is filling in the gaps around that with other services and technologies
            - Can't say the AI component or machining learning has been a massive driver for them other than the fact that the whole industry is basically mandating that you have to have the ability to look for threats based on behavior, however that happens
- Another big thing that has been said for a long time is that you spend all this money on technology that has AI to analyze and create a baseline for your environment and then alarm based on deviations from that, where at the same time most businesses in the country are under say 500 users, even 100-200 users
    - So, it is quite easy to have someone spend 30-40 hours looking at each department seeing what normal things the department does and write up basic rules that's proprietary to that company (might be just as effective as spending all the time on an AI algorithm that you don't have a whole lot of info on)
- UBA (User Behavioral Analytics) – how do you baseline what your users normally do and then create alarms for deviations from that
    - It might take 3 weeks for the machine learning algorithm to train itself that your executive team logs in between 7:30 and 8:30 AM every single day and if

they log in at 6:00 AM then it is weird (this might be a level 3 alarm) or if they log in at 3:00 AM then it is really weird (this might be a level 10 alarm)

- You could, yourself, write a rule that says anything outside of 7:00 AM to 6:00 PM is an alarm that someone needs to go and look at – you don't need machine learning to train itself for three weeks to do what you could just ask a person/ department to write into a rule set
    - Obviously this doesn't work for large organizations with very large data sets and complex models
        - But what he is seeing is people are leaning on AI technology to solve problems even though they don't know how they are being solved (don't understand how they were created to begin with)
        - It ends up causing issues in other areas (laziness, lack of documentation) because they are leaning on a product that uses AI/ML rather than actually building up their own kind of environments and creating better procedures and policies themselves

Interview 6:

- Develop their own technology
    - As a company, started out in User Entity Behavior Analytics
    - AI-based technologies normally focus on anomaly detection
        - When you are trying to catch hackers, you don't always know in advance what you're looking for
            - They might change the tools they are using, change internet addresses that they are attacking from, change the malware (piece of malicious software)
        - If you don't know in advance what you are looking for, you can't have a simple search filter, you have to start looking at behavior

- Only way to look at behavior is to use machine learning and statistical analysis because you are trying to identify deviations from baseline, rather than a specific action

- Vendors use this AI/ML to implement these behavioral capabilities at a network level and at the process executable level (laptops, workstations)
  - It is being leveraged to try to generate some sort of baseline of what normal behavior is, and then identifying deviations from that baseline, but then also tying that to specific types of behaviors
    - Ex: Microsoft PowerShell runs everywhere in the whole world, but how do you know when it is being used maliciously
      - Baseline this on hundreds of systems with some sort of algorithm to see how it is normally being used across those systems – if they see any deviations from that baseline, from now on, they can generate alarms

- From their point of view, AI machine learning is a good concept because it does increase their detection rates (increases the ability to find things that are actually malicious and then block them)
  - At the same time, it's creating a bigger divide between people using technology and between the vendors, because all the vendors are considering this proprietary technology and they're not sharing it
    - When they develop a really good algorithm for analyzing these types of things, of course they're not going to make it open source, they're not going to share with other vendors, and a lot of times they're also not going to share with the customers either

        - You are at the mercy of the product vendor to say "well I can tell you this is bad because of our algorithm, but I'm not going to tell you how we figured out it was bad"
            - Start to get the issue of only being as good as what the vendor can do, and you have to make sure you're cognizant of that and be able to fill in the gaps in other areas with other technologies
    - A good thing from the detection perspective, but from an industry in collaborative perspective, it is creating other challenges
        - Same thing happened with the overall intelligence segment, just indicators/signatures in general
            - Back when it was McAfee vs. Symantec, they had whole teams of researchers that were analyzing things and creating signatures
            - The same is true now for the teams that are developing and analyzing these types of algorithms, but what ultimately happened was saturation of the market
                - Other vendors came into play that may not have had as big of a team, so they would be licensing the signatures from other vendors (all back-channel licensing and sub licensing relationships)
                - Going to start seeing the same with AI – it doesn't matter what the product is if the intelligence is the same
                    - Researchers and consumers are going to start mandating that there is openness about what's being shared around the algorithms and how they are being analyzed and tuned because people are going to get more and more frustrated with this proprietary notion

Interview 7:

- Security is tough to develop your own technology

    - Don't see many companies build out their own security technologies because it is cost prohibitive
- In the past, used enterprise technologies through the leading security vendors
- AI has evolved on many fronts
    - Incident response within the security program
        - This pillar is designed to identify malicious traps that are happening in the environment
        - Trying to stay ahead of sophisticated hash packers like nation state (folks with a lot of money that try to fly under the radar)
- The value proposition with AI in that space is that it can help identify anomalous behavior, things that may not look like a threat on the surface, but be able to constantly analyze what's happening in the environment 24/7 – in the event that it identifies something that looks suspicious, it will alert you
    - Within the context of incident response, basically it is all about trying to stay ahead of the bad guys who are leveraging AI as well
    - Identify and keep track of what normal is and then identify abnormal (which deviates from day-to-day business) and AI does that better than a human
        - Obviously because humans just can't look at that amount of data and make a conclusion on it, whether it is an outlier or not from normal behavior
- In the context of vulnerability management, some AI is used to quantify vulnerabilities in a better way (look at different attributes of vulnerabilities and classify it in terms of how high of a priority it should be)
- Biggest place that AI exists is in incident response
    - That is where you're actually looking for the bad guys and they're leveraging AI at the same time while corporations are leveraging it, and it is a battle between the two to see who can stay ahead of the AI at this point in time

Interview 8:

- Develop their own technology
    - Sell their product for a variety of use cases

- AI for good
    - Predictive analytics in Africa for water wells
        - Charities build water wells for small villages, but sometimes it can take about 2 months to get notice if something goes wrong.
        - Use their platforms to say predictive analytics of what well was going to fail next
    - Department of Corrections
        - Data on the criminals (background, education, zip code, etc.)
        - Predict the criminals that would reoffend
        - 90% accuracy
        - Not used to just keep people in jail, but so they could allocate their resources better
    - How many bananas does Walmart need
    - Where is the likelihood to be fraud
    - Who banks should loan money out to
    - Large IT customer used the software to do predictive analytics on outages within the network
        - Saved $2-3 million a year with it
- AI – you can use it for absolutely everything

2. If your company is the one developing the AI software, how do you go about that process?

Interview 6:

- Have a Data Science team that works on the basis of hypotheses
    - Read up somewhere about a certain type of attack
    - They work out some AI approach that they could use
    - Then they start testing it to see if its viable solution

  - If it works out to be something usable, the engineering component comes in where they try to build it into a product

Interview 8:

- The engine is there, it is just iterations of the engine that are being built
- 300-400 engineers that are developing software
  - A lot are in Eastern Europe – low-cost sensors for labor
- 300 data scientists
- The cost is in hours of labor for building the use cases (anywhere from 10 to thousands of hours depending on the complexity of the use case)

3. What are the considerations when they license security software in terms of cost?

Interview 1:

- Depends on the vendor, product, and size of the company – can get really expensive
  - A lot of vendors will bill by seat (the actual individual license they use for the product)
  - Others will bill by throughputs (amount of data you are pumping through the solution
- AI driven Splunk can be up to millions of dollars a year depending on how much data you are pumping through it [Splunk is a tool that collects logs, can run queries on them, and alert on them. AI driven Splunk also incorporates anomaly detection algorithms]
- Any cutting-edge technology is not going to be cheap, but it is a must have
  - Companies are not afraid to open their wallet if it is going to help them have more actionable alerts in the future
- Depends on the size of the company
  - Huge corporations could see a $6 million bill and not blink an eye
  - Smaller corporations could see a $1 million bill and people will question if they really need it

Interview 2:

- Cost is usually the biggest consideration – want to know that you're getting all the features you want, how many seats you get, how many computers you can install it on, how it's going to be supported, how long is the contract for, if you have issues with the software who do you talk to, etc.
- Paid around $80,000 for 300 users for a couple years
    - Renew it after the period is up

Interview 3:

- In the past few years, things became more affordable
    - The first SIEM solution he looked at 7-8 years ago were close to half a million dollars and up each year
    - Carbon Black was really expensive a few years ago but has become more affordable
        - It is becoming more affordable because more players are jumping in and there are more solutions
            - Still an expense
    - He spends somewhere between 12-15% of his IT budget on cyber security

Interview 4:

- You get what you pay for
    - There are some great technologies out there that are great because they are open source, and it is a collaboration of people working and building different things within open-source tools
- In his experience, you can get some pretty lousy software and thinking you can get apples to apples
    - Theres a lot of different third-party analysts that do reviews of these different regions
        - SIEM magic quadrant, there is also a firewall magic quadrant that Gartner does
    - If you want to get a really high-end, top-quality tool, you are going to see who the leader in that space is
        - They are competitive and are competitively priced

- You can get different tools that are a lot cheaper, but you will not get the level of incident detection and functionalities that you are getting elsewhere
    - At the same time, some people don't want that. Theres some people that just want a tool that collects, and stores logs and they don't care about security and running through threat hunting. For these people, they should get the cheapest tool that just collects and stores logs and does some searches [like a non-AI Splunk] and that is fine for them
        - Call them the box checkers, who go into compliance, and they say, "yep I checked the box, I have a tool, it does what I want it to do" (store logs and search)
        - There is a cost if you get hacked that you have to consider

Interview 5:

- They leverage their vendors, and most vendors have a goal to make implementation as simple as possible
    - They are not implementing AI, they are implementing a tool set
- Endpoint technology advanced – next gen antivirus like Carbon Black and Cybereason
    - Their implementation is as simple as building out a standard set of policies, which for the interviewee's company takes about 15 minutes, and then they give an install file to an assistant admin at an organization to roll out
    - There are hundreds of thousands of endpoints - tune algorithms based on alarms that get generated
        - Tune it on the fly – usually takes 2-3 weeks
        - This takes dozens of hours, not a massive amount of time, but it does take time to tune
            - Tuning is not waiting for the algorithm to train itself, it is "hey, this alarm was generated, and this user is complaining something was blocked, let's just put a policy in the whitelist" – and that takes 5 minutes
- Implementation from a cost perspective is very low

- Vendors have huge data analytics teams that are analyzing these data sets and then working with engineers to build these algorithms and adjusting it
    - The vendor costs are massive, but as the consumer, it is easy to implement
    - Still a lack of visibility – industry is eventually going to mandate this (you don't know how much is really going into the data sets and the training of the algorithms and understanding how they are actually being implemented
        - Could just as easily have a couple thousand behavioral rules that are just as successful as an AI engine
        - It is a combination right now, you cannot rely on an algorithm for everything, so you have to have protection rules built in that are "hey if this happens, make sure and generate an alarm and then allow policies to bypass it"
            - A lot of the times, it has nothing to do with any sort of AI or machine learning, it is literally just a set of rules of here are a couple thousand scenarios and if this happens, generate an alarm, and that can get updated fairly easily

- Buying the software
    - If you have 100 computers, you buy 100 licenses of the antivirus with the service, and pay for the years' worth
    - Could be anywhere from $50 a year to 100
        - Basic Symantec is 25-30 bucks a year
        - The more next gen, industry-leading stuff isn't that much more expensive
        - Pricing has come down in the market a lot over the past couple years
            - A lot of vendors doing the same thing, collecting data, making rules and algorithms to analyze it, and they are all tested against each other

- The top 5 today, are probably the same top 5 from last year, and will probably still be the same top 5 next year – if you pick any of those, you will be in a good spot

Interview 7:

- The cloud flipped the paradigm of AI being cost prohibitive on its head
    - In the past, let's say you had on premise technology [hardware and software applications that are hosted on-site], and you wanted to leverage AI. You would maybe have some module on premise, have the server capacity, have the infrastructure supported, have the technologies that integrate, and have on premise that integrates with an AI
        - Building out this structure was cost-prohibitive
    - When looking at a cloud environment, the vendors will build out a solution, overlaying the AI across the solution or datasets, and they carve out virtual environments for the different corporations based on size
        - Paying per user (same amount of money per user)
        - Only have to have internet connection
- Today, there is not a cost issue anymore, everyone can leverage it and that is the beauty of the cloud

4. Do the executives of the company understand and value AI in terms of security software?

Interview 1:

- Executives care about the cost aspect
    - Cutting-edge technologies can be really expensive, but there are also ulterior costs like resource costs
        - If you can use AI to really cut through the amount of noise in your environment and get real actionable alerts to your analysts, then you might actually be saving money in the long run (resource hours – people looking at garbage alerts that they won't have to look at in the future)

- C-Suite adoption is all cost-driven
    - They do care about security posture (they would have to spend a lot of money trying to clean up after a data breach or if some other malicious event hits the news, they may need to spend billions of dollars' worth of remediation to try to fix the customer view of the company
- When trying to force the adoption, you need to be able to sell the ability to reduce costs in the long run from either collapsing tool sets or buying a tool that has AI capabilities that once it is fully running in your environment, you could turn off and stop using several other products – in turn saving money
    - Making resources more efficient at their job also helps with cost savings

Interview 2:

- President was on board
    - When they were moving to CrowdStrike (next gen antivirus), they had a presentation put together explaining all the terms and benefits
    - The minute they say ransomware and they can stop it fairly quickly, they get on board

Interview 3:

- Interviewee is part of the executive team; thinks peers understand AI but not from a security perspective
    - It is a constant education to get them to understand what security is and why you have to have it
- 4-5 years ago, he implemented an education program, and they hired a company (paid them like $5,000 a year for the subscription for 250 people) that has thousands of online courses
    - They enroll their employees in these courses (ex: about phishing emails or social attacks, how they occur, etc.)
        - They had some opposition from people, thinking it was a waste of time
            - The training did help some people in their personal lives (ex: dating apps)

- Produces KPIs [Key Performance Indicators] of the new tools to show how they are working
    - When they first invested in Mimecast, interviewee did a report showing how many phishing emails were caught, how many got through
    - Interviewee works with "techies" [technologically skilled people], and they don't know how to communicate in business language, so interviewee converts their language to a business presentation

5. Will the purchased AI software replace a method already in place or will it be in addition to preexisting security measures?

Interview 1:

- Always going to be an additional piece
    - Still a relatively new technology
    - Products are still adopting it to help their own use case
    - From personal experience, there hasn't been a plug and play AI product where you just buy it, test it, then you are all set – usually it is a lot of tuning and creating processes around that of having human eyes look at the alerts and either take action on them or suggest further tuning
        - AI can be relearned if you are getting super noisy alerts that aren't helpful for anything – you have to go through the process of learning the environment again to see if you can get better results
    - In the short term, he doesn't see any industry where AI is an all-in-one product that does everything for you
        - Always going to be adding context to your other tool sets and hopefully be able to improve accuracy and speed to remediation and visibility in your environment

Interview 2:

- Wanted it to be in addition to the normal antivirus, Windows Defender (not nearly as expensive as going to CrowdStrike), but it depends on the company

    - Some like to run both to have the extra protection, some just run one to cut the cost of running the other, but doesn't think it is a huge cost to run both

Interview 3:

- No magic bullet solution
    - You just layer your solutions and keep adding on every year
        - They either come up with another layer, or replace a layer (replaces endpoint protection every 3-4 years, he reevaluates every solution to either keep it, not keep it, replace it, etc.)
- AI is built into every security solution these days whether it is a spam filter, endpoint protection, network protection engine, network scanning, or penetration tests
- No comprehensive one solution that covers everything – companies may claim that but, in his opinion, no

Interview 4:

- Security in general is leaning towards machine learning specific, there is always a level of what is marketing and what is reality
    - Marketing oftentimes says, "yeah we are AI", but when you really peel back the onion, you can see that it's not, and for that reason, the interviewee's company doesn't claim to be AI, they claim they do machine learning functionalities
        - Based on what he is seeing in the industry and in the space, the tools that incorporate these functionalities like machine learning are definitely getting ahead of the legacy tools

Interview 6:

- Primarily enhancing existing needs or existing technology
    - Security monitoring has been around for a long time, they are just starting to automate a lot of the actions, and trying to detect behavior rather than using a static indicator (ex: a file name: malware.exe)
- They sell AI but they also buy it (AI is everywhere now), they are consumers as well
    - Use AI when they develop stuff, use it in marketing, they use AI all across the company

        - Examples
            - Canva – AI based presentations (Helps with suggestions)
            - Analyze why you win/lose against competitors
            - Grammarly – language suggestions (lots of mistakes)
    - It is really valuable for certain topics
    - There isn't anyone who just sells AI but doesn't use it, at least not that he is aware of
- They have been using it primarily to improve something which people have been doing, and especially because the attackers are evolving as well. While they are not building AI malware, they are using the same kind of AI tools companies are using to speed up their game and to improve what they are doing

6. AI needs a large amount of data to be trained, how are you going about getting that data and diversifying your data set?

Interview 1:

- Their environment is "a little bit weird" – they have a lot of networks converged together, which presents the challenge of getting as much visibility as possible from all their different networks (application logging, event logs from user processes, service accounts)
    - Service accounts – pain for security
        - Service accounts are accounts that are used to run some process that are not associated with a human user (automated processes run off service accounts)
            - Presents issue with accountability – If you have a service account that is taking an action that you are unfamiliar with, it is hard to trace down where that is coming from if there isn't good documentation as to what this account is used for and why it is used
                - Run into this issue a lot

        - Service accounts are being used as a work around for some other issue with permissions
            - AI can determine what is or what isn't a normal process for that service account and can help alleviate some of that problem

Interview 2:

- A lot of companies take data right from their install base [their own network/environment]
    - CrowdStrike (major player in the next gen AI field)
        - Crowd source (hence the name CrowdStrike) all the data they need from their customers
            - Their sensors collect data from each customer (their unique experiences, the attacks that are happening to them) and they aggregate it all together
- Other companies have some sort of trial and beta period that they use
    - If you have a product that's installed on thousands of clients across 10s of thousands of PCs, running 5-6 times every day that's a pretty good data set that you're going to get

Interview 3:

- Traditionally, solutions used to live in your network [some centralized computer with installed software that manages applications on endpoints]
    - You would have an antivirus server that pushes to the client endpoint protection
- Now, it is Cloud Based solutions
    - AI solutions gather information from any feedback from a new virus that comes in within the cloud community
    - Today, AI takes input from a lot of their clients, versus having to use your own server and knowledge base (whole world is now feeding information to the AI)

- The AI brings this information together to evaluate what is going on in your specific environment
- Machine learning cannot be isolated for one client, it has to be an aggregation of the whole community

Interview 4:

- Biggest challenge with this space is the fact that there is so much data that is going around
  - What they try to do is understand what normal behavior is in the data, so that if possible incidents happening in that data there can be a response
  - They try to capture an anomaly in the data
    - Do this by understanding normal based on employee information (across your organization – HR, finance, sales) and see where deviants or changes in the behavior is happening
      - Ex: Sales department normal behavior would be to access different databases related to sales, so if they start accessing an HR server with employee social security numbers, or some other application that wasn't suitable for the role, there would be an alert

Interview 5:

- Doesn't have a whole lot of visibility into it
- Generally, the products and vendors tie into their requirements that they are sending data back to their "clinics" (collecting all the data, centric)
  - The more customers they bring on, the larger their sample size gets

Interview 6:

- Depends on the type of learning
  - Deep learning – need large amount of data
  - There is an increasing focus on types of machine learning which don't need that much training data because in cybersecurity, they don't always have the data

        - Coming back to the point about you don't know what in advance you are looking for
- Data needs to be labeled – this is the bigger problem
    - Labeled means that you tag it as good or bad
    - Problem in cybersecurity is that generally, a lot of people don't like sharing true positives
        - If they have been attacked, they don't give that data – this is data they don't have a lot of (don't have a huge representative sample of it)
        - Have lots of benign data (data that is business as usual)
- Another problem – not just getting the data, but paying to maintain the data
    - Large amounts of data still cost a lot of money to store
        - The more data overtime, the more expensive it gets
        - There are people who spend millions of dollars a year on storage costs
            - Typically, in the range of $10,000-20,000
        - Storage types
            - Warm Storage
                - Can immediately search and query it
            - Cold Storage (cheaper)
                - Means you are compressing the data and putting it onto a cheaper storage medium for longer term
                    - Ex: If you find out you were attacked 9 months ago, you will probably have to get those files out of cold storage and decompress them, it is not just available at your fingertips
    - They try to, where possible, use approaches that don't need years' worth of data, rather they can work with weeks' worth of data
- Get their data from their customers and by capturing samples of malicious activity
    - Have a honeynet

        - Have researchers who will have fake systems on the Internet waiting for them to be attacked so that they can see how it was done and use that to train models

Interview 7:

- There is a technology that can integrate the environments of say vulnerability management and incident response – SIEM [Security Information and Event Management]
- Leverage a SIEM, which is a next generation SIEM
    - Difference between the next generation SIEMs and the traditional SIEMs written prior is that they leverage AI to go through the data and identify trends, and can bring in vulnerability data
        - These technologies have feeds, and you can feed data from different sources (from your computers, your servers, your cloud technologies, building management type technologies – i.e., temperatures and card access systems)
        - As it ingests that data set from all these different environments, it learns the kind of patterns that your employees leverage within the environment (how they interact with a computer, how they interact with the rooms in terms of access, etc.)
            - Learning is done in an automated fashion that's designed to learn – the more you feed it, the more it will learn
- Goal of the organization is to feed it as much data as possible
    - Look at all the different areas within the organization and set up the feeds
        - Once you set up the feeds, it takes about 30 days for it to learn the environment and then any minor deviants will be picked up
        - If you bring in a new solution, technology, or building management system, you integrate all of these feeds to keep the AI up to speed with how your environment/corporation is changing from that perspective and it will learn

- Vendors are constantly updating the AI to make it learn better, so as they're improving it, you'll feed that as well, giving it more and more data sets (solutions are opening up their technology more and more to provide more data)
- The data feed is what allows you to teach the AI

Interview 8:

- Customers
    - Customers use their own data set
    - They consult with them and help them, but it is from their perspective on the particular use case

7. Has your company experienced a skill resource gap as the technology advances?

Interview 1:

- Cybersecurity in general has a gap – starting to change
    - There are curriculums that are cybersecurity centric
    - Interviewee's undergrad and even their masters, wasn't cybersecurity centric
        - IT major that does cybersecurity on the side
    - Two ways of looking at it:
        - People who have worked in IT and they have a good foundational knowledge of IT, and they just need to learn security parts
        - They are security students that have a risk analysis impact mindset, but they don't have the technical knowledge coming through IT ranks and knowing how those processes work
    - For someone who is skilled in IT in general, as they move on in their career, they pick up the security mindset
        - Normal (traditional) path: work in IT, gather expertise, and then lean your resume towards security
            - How this interviewee came through the ranks
    - Some schools are starting to offer security as a major
        - Programs dealing specifically with cyber security and are pumping out students with that skill set

- This creates problems:
    - Just learning about the security side of things, you don't know how it works under the covers/IT side of things
    - Harder to manage when you are just looking through the security lens (risk levels)
- Hot topic – news articles saying there is 0% unemployment rate in cyber security
    - All of a sudden, everyone wants to be in cyber security and floods in all at once
- When hiring
    - If you have somebody that has the IT background, and they have security knowledge – they are a unicorn (rare) – hire immediately
    - It is usually picking one or the other, you just have to teach them the other side

Interview 2:

- Not with skills, most of the teams he has been on have been pretty well trained, keeping up with the latest skills
- They would love to have more people on their team (there was 3), but that might have been a resource thing or business issue

Interview 3:

- Absolutely, they still do
- Gap on multiple levels
    - First level – whole world changed how you build an application and how you deploy it (cloud application)
        - Applications used to be built on one VPN [Virtual Private Network] or around one box, now they build applications on web services, API [Application Programing Interface] services (database services) and put it all together. Each of these services has its own vulnerability, so if you don't pay attention, you are left vulnerable

        - Gap in understanding how to build a secure application using cloud services
    - Second level – to become a hacker, you don't need to be that sophisticated
        - The volume of those is increasing, and the sophistication of some attacks is also increasing
        - As a small business, they try to build people in house and complimenting that with an outside contractor (Logarithm is managed primarily by Logarithm, secondarily by interviewee's company)
            - He only has one person, not a whole team, so having the external team that does it (sharing resources) is becoming very common for small and mid-size organizations to close those gaps

Interview 4:

- Without a doubt, one of the biggest things that has brought out the need for these technologies that are constantly and consistently evolving and changing
- When they first got in this space doing security event management, they were at a company where you had to write all your own rules
    - Biggest talks with customers today are that we are not in the day and age anymore that you can log into this application and see a to-do list and have 200 alerts that you need to go through – not feasible because you have maybe 3-4 team members that are tasked with this
        - Not only do they have to work in their platform, but they also have to be using for example a CrowdStrike or Palo Firewalls – they have other fish to fry, they can't just be flipping through false positives all day, and all the alerts that aren't real in order to respond
- With the skill shortage, companies are much more reliant on solutions like theirs that actually prioritize different alarms and different alerts
    - Ex: Janet is talking to a competitor and accessing files she normally doesn't use that aren't suited for her role, they use different processes to get those to bubble to the top of the surface for the analyst so they can address it first

instead of it being item 737 on the priority list – try to make life easier for the security analyst

- Also, different tier levels of incident response
    - For something that would be a low-level alert would be like somebody talking to a competitor, it wouldn't be a level of an incident that a high-level security analyst should be dealing with, should be more entry level type responses
        - Their product also helps with that by prioritizing the alerts so the more entry level type responses can be handled by different teams

Interview 5:

- Resource shortage in the industry is a bigger issue than the skills gap
    - A massive skill gap based on new technology – doesn't think that's happening as much as there are not as many people that are spending time learning these technologies, which is the bigger issue
        - You have a lot of people that are learning only cyber and maybe learning some tools, but to use a lot of the tools properly, it requires a lot of foundational infrastructure background
- What they found, is people that have background in systems administration or network administration, and have done that for years, then transition into cyber, can pick up the concepts of the newer capabilities much quicker than someone that's only focused on just the cyber part
    - Ex: your mail server crashes at 10PM and everyone is going to be there at 9 AM in the morning, so there is pressure to have it fixed (hundreds of people are relying on it)
        - Understand what it means to be responsible for that scenario and translating that into security tools, it's been helpful for a lot of people as opposed to ones that haven't had that background

Interview 6:

- It's always hard to find good security people

    - Tend to recruit people who come from IT, development, or tech, and they train them the security part
- This process makes sense because you can do a degree in cybersecurity, but you can't really just study cybersecurity
    - Security is always predicated on understanding something about what you are securing (could be infrastructure, web services, code), that is a bottleneck in terms of training people
- Amazon, Google, Microsoft
    - Between them they have about 200 something thousand security people
        - Really starved the market
- Vendors Game
    - Lots of vendors selling security software
        - Also starves the market (end users of security people because vendors pay better)
- Other problem – have a lot of people in the middle (not necessarily new entrants)
    - An entry level job, you don't tend to do for very long, you can get a new position quickly because there is a skill shortage – makes it worse
- Almost everyone has experienced the skill gap to a degree, but depending on where you work in different ways, it has a different impact

Interview 7:

- It is tough to find people that actually understand security
    - Some people think security is just about clicking buttons or implementing technology, but there is also a business element of it, staying ahead of the malicious actors and being able to research and see what is happening in the field – those types of individuals are hard to find
    - If you are looking for a firewall administrator or somebody that is dealing with a particular technology, it is a little bit easier to find people, but those people tend to not understand security and the different pillars of security
- Shortage of people at the higher end of security

    - Because of automation and some other things that put the people at lower levels out of work
- Overall, definitely a shortage, but also thinks there is a shortage in every field
    - Talked to his vendor who said that there was a shortage in finding a good security person with 10 years of experience and with cloud experience
        - Large percentage of companies just started to migrate to the cloud
        - COVID forced companies to embrace the cloud and embrace innovation
            - Need for agility in a lot of areas (quick turnaround to working at home) sped up the movement towards the cloud strategy
- Where we are today is a good thing in terms of being able to have a situation where companies across the border are as agile as possible, leveraging the latest technologies and potentially will help with the gap of not being able to find staff because they are not doing things in house anymore, they are instead leveraging the cloud

Interview 8:

- Yes, absolutely
    - They used to have 4 of the top 20 data scientists in the world, and now they have one
    - It is hard to compete with other companies
    - A lot of their customers are struggling with their product because they can't hire data scientists
    - Shortage of staff right now – great market to be in
- In the past, there were no developers, so when there was one, they would get paid double, or a lot more. Then, there is a flood of people to that market, which leads to pay cuts
    - A man he knew was the highest paid person in the company, and the company said he could renegotiate his salary, or they would have to let him go
        - Were so desperate two years before

- It is like a pendulum swinging – everyone goes towards the market, flooding it, then it goes back to a shortage

8. Have you ever participated in a red team/ blue team exercise? If so, how did it go, and do you have any suggestions for improvement?

Interview 1:

- Mentor to students trying to change career paths, or are fresh out of college trying to get into cybersecurity
    - Curriculum has the red team and blue team exercise
        - In their experience – it is antiquated [outdated]
    - Normal red team versus blue team engagement – there isn't that cooperative factor
        - Red team is going through their playbook of what they're trying to get done, they will go through their process of trying to get in and trying to get lateral movement to see where they could get
        - When that is done, they have this big, giant documentation that they send to the blue team
        - Blue team does their best to remediate those issues that were discovered during the red team attack
        - Then the blue team goes back to the red team to see if those solutions worked
    - Purple Team [newer concept] – doing it side by side together red team and blue team
        - Red team is constantly working while in communication with the blue team saying hey, we are finding this can you take a look, and then blue team fixes it as they go along
            - Much more cooperative and agile
            - Prevalent in the industry today
    - "If you are in an interview and they ask you about red vs blue, you better start talking about purple because that is where it's going"

- Suggestions – purple team and transparency
    - As an Incident Response Analyst:
        - They would get events that would look bad, and they would waste resource cycles looking into the artifact that was found and alerted on
        - It ended up being a red team penetration test that no one was aware of so then they spent real time looking at fake incidents when they could have been looking in real stuff in their environment
        - Pit fall of traditional red vs blue (not transparent)
            - Purple team doesn't have that resource waste

Interview 2:

- He has not participated in an actual red team blue team
- He has done tabletop exercises, but never a full-blown red team blue team
- They planned for them, but never had the resources to pull it off
    - Always been on smaller teams, so he hasn't had the chance to do it yet
    - Wants to do one in the future
- At one of their jobs, they were a team of 3 people
    - Would need one person dedicated to full-time pen testing (red teaming) and the other two could do blue team, but they are all generalists, so they didn't have the full-time pen tester on staff

Interview 3:

- He did it a couple years ago on a practice base, not an actual incident base because they always have incident response policy
- Learning experience – having what policies are in place helps (knowing what to do, who to notify)

Interview 4:

- His company has done them in the past
- None of his customers have
- Other regions have had their solution deployed in a red team blue team exercise

- Biggest challenge they see when in the presales stage (working with a customer trying to get them to buy the product, showing their value as a software and solution) is that they have a short window of being able to have the solution deploy
    - The way their application works, is the longer that is has time to ingest data and understand normal behaviors and generate baselines, the better it can distinguish abnormalities
        - If you're just looking at a small sample, everything would pop up as weird because there is not enough sample time to see that
            - Normal testing period is 30-60 days, and red team blue team exercises normally take a lot longer than that
- He himself has not been included in the exercise, but other regions have been, and they have been successful with that
    - Simulation of a potential incident, and them being able to catch that – have been successful doing that

Interview 5:

- Yes, they do these all the time
    - They have a full red team, and their SOC (Security Operations Center) is considered their blue team
    - Call the results purple teaming (taking the results of a pen test/red team type exercise, presenting that to a blue team who should have been the ones that were actually responding or monitoring that environment, and comparing notes to see how far one team could get without the other team detecting it)
- Suggestions for improvement
    - Results are intended to be how much did you not find and how do we make sure
        - Detection rules
            - Ex: I did this at 3 in the morning and here are the 3 steps I did to get there, well let's write a rule that looks for those three steps
                - Have to have feedback – maybe it could have been 20 steps or 20 steps, how do you make sure you write the

rule where it's flexible enough to actually detect the behavior you're looking for rather than just exactly what was presented

    - Being able to understand what all the possible permutations would be for a particular scenario and not just focus exactly on what happened
- Done at different times in different ways depending on the different purposes
    - Might do this engagement in tandem with the blue team and every step of the way they are aware of what's actually happening so you can build your rule sets as they go
    - Quite often what happens is you work until the red team engagement is completely finished before you bring in a blue team and do a purple team or tabletop exercise and run through all the scenarios
        - Usually that is after the fact, could be a month or weeks after, and to go back weeks from a blue team perspective is sometimes challenging (depending on what your data retention is, some systems only have data for 3-4 days)
        - If you are trying to make sure you have detection capability for something that happened a month ago or even longer than that, a lot of times you're not going to have the data available
            - Depending on what the outcome is, you might have to do it in lockstep [progress at same speed and direction] to make it considerably more effective (or when you go through those scenarios, make sure that if there are things you don't have data for, those get called out and either identified if it needs to be done in lockstep in the future)

Interview 6:

- Yes, as a pen tester (red team, attacking side)
    - Company they are on advisory board for automates that part – called a breach and attack simulation solution
        - Tries to look a lot like a hacker would

        - Whole idea is to see if you can detect it or if you can block it
- It is hard for a lot of companies to have a red team
    - The average company does not have a pen tester in house, so using some kind of automation tool is quite common
- Best thing for improvement is to mix what you can automate with what people can do to increase how much you can test in a certain time
    - If you automate it all, it isn't as smart as a hacker
    - If you have a person do it all manually, it takes forever
        - At that point, people won't be paying for 2-4 weeks of a consultant doing this, because they will cost somewhere between $1,200-2,000 a day, which adds up
    - So, combining the person with the automation is the way to go, and that is the way that most of the market is going
- Hacker One – has a friend who works here
    - Does automated pen testing
- Humans have a thinking adversary, and that can adapt to what you are doing, and they can be stealthy; Automation is really brutal – you can see it immediately from a detection point of view because it is not designed to be stealthy
    - A human can really take their time
    - Companies can't afford to hire somebody full time to do this pen testing (budget is already stretched enough)
        - Just do a small test – usually not enough to actually validate that what they are doing is enough

Interview 7:

- Yes, definitely, participated on both red team and blue team
    - Constantly improving your process and understanding where things fall down
- Improvement – couldn't think of anything it needed to improve on

Interview 8:

- On the red team yes – tens of times
- Most times, they have an internal red team, someone full time

- Judged on how successful they are in finding bugs
- There are companies that earn money off being the red team
  - One guy will do everything from cyber security to breaking into buildings
    - Use technology to pick locks (all electronic)
    - Knows how to bypass systems
      - Get into a building's HVAC [Heating, Ventilation, and Air Conditioning] system – which if it was for real, can cost a company millions of dollars' worth due to physical damage
- When doing a red team exercise, it doesn't matter how you get in, you just have to get in

9. What is your experience with the evolving malware, and what are you or your company's responses to that? Have you experienced any attacks that have used AI?

Interview 1:

- Big topic in the industry – threat intelligence driven monitoring detection scheme
- Signature based was good, but not good enough to beat zero-day attacks and new methodologies (like AI capabilities to help infiltrate)
- Should start to use threat intelligence feeds to source into those detection tools to make those algorithms smarter and more accurate
- AV [Antivirus] Product that uses AI - Was recently going through detections and comparing what that product flagged as potentially malicious through its AI calculations – took those hashes and compared them to the community (Open-Source Intelligence)
  - Community agreed that about 40% of them were malicious and the other 60% were benign/no one has reported issues with those hashes
  - AI is nice to have – to throw attention at things that could possibly be bad
  - But it is not the "be all, end all" - it has a human component

    - Human Component - Training the algorithms trying to make them more accurate; connecting into as many threat intelligence feeds as you can to improve that accuracy can go a long way

Interview 2:

- Doesn't know if they have experienced any attacks using AI
- Malware attacks at the bank they worked at
    - Log4J – hackers took a piece of Java code and inserted malicious code into it
- A lot of attacks happen with phishing – bad guys go after the human element

Interview 3:

- Have not experienced any attacks that use AI, but they don't know how they could tell if the hacker used AI or not
- Education to desktop users – they catch a lot of malware with their tools that they have in place
- Haven't had any major incidents, few incidents here and there
- There was an email that looked like it came from the CEO asking an employee to do an urgent bank transfer, and he almost did it, but they caught it in time

Interview 4:

- Their solution doesn't directly stop malware or prevent malware, but they can detect it
    - If they had a user that clicked on a malicious email or downloaded a bad file, they have the ability to see all the other employees or users that may have also touched that file
        - Do analytics and detections based on that

Interview 5:

- They see loads of stuff, as for anything that uses AI, that's a good question – they don't know, but it is very possible
    - Usually what they do is tie it back to the particular groups that are actually generating them, so generally try to at least identify what the variant is and then there is usually quite a bit of information you can dig into in terms of how it gets used, what technology is being used, whether it is AI or not

    - A lot of stuff out there that is using detection avoidance (complex concept), avoiding being detected in sandboxes
        - There are variants that use algorithms to identify it it's trying to be run in a sandbox and then it will refuse to detonate itself so that it can't actually be analyzed
    - Some things to consider as data points to go into an algorithm:
        - Running a VM [virtual machine] that cannot detonate other ones
        - Is it running on the same system that was downloaded originally?
        - Was it moved off to a different system and not downloaded directly?
        - How many times has it been detonated from this IP in the last period of time?
- In the last years they have definitely seen a huge uptick in the number of complex systems that are "living off the land" or living off file list types (PowerShell or built-in windows encryption-decryption libraries)
    - Can be less likely to be detected if it's using built-in functions and not downloading it or bringing in its own code – these things are getting more common and harder to detect
        - Presumably, they're building these with some sort of intelligence that is able to avoid detection
        - Timeouts are a big one – might detonate code for a period of time initially to avoid detection, or it might wait before doing that to avoid detection
            - Seen this more recently, where they might wait until the middle of the night - if someone opens a link or clocks on something at noon on their lunch break, it might wait until 2 in the morning to perform a second step (could be hard coded or could be based off algorithms)
                - Say no one's been sitting at the keyboard for 8 hours, so I have another 6 hours for this to be run, so it is going to

wait for a period of inactivity – this happens a lot, generally is for the purpose of detection avoidance
    - Even if it gets detected, if it happens at 2:00 in the morning, if you don't have someone watching your environment 24/7, chances are someone is not going to look at it until 9 or 10 the next day, or even later – at that point, the damage has been done and doesn't matter

Interview 6:

- Have not experienced any attacks that use AI
    - Not going to get a piece of malware that has AI in it because it requires a large size
        - Not something that could be hidden easily on the system
    - You could use the machine learning to create malware, but the malware itself wouldn't be AI, it is AI developing the malware
- On the Advisory Board of a company focused on Adversarial AI – they do ML testing
    - They build adversarial AI, and they are always looking for real world examples
        - There was an image recognition attack where somebody faked an image to fool an image recognition software
            - Wasn't a data breach for money, but it could be used with some sort of theft (ex: breaking into someone's bank account)
        - Other thing he heard was that someone was using machine learning to create fake companies, including fake people on LinkedIn with fake generated faces (people are all connected, so it looks legitimate)
            - Not the same as actually attacking
- Attackers don't need to develop an AI to hack companies, all they need to do is find somebody who isn't paying attention and send them a well-crafted email, so the AI would be overkill
    - "It's like pulling a cannon to shoot a fly"
- Another advisory board

- Do an annual report called a Director Report – look at the complexity of malware and measure it by the number of actions that the malware can do
  - Compared to last year, it was about a 20-25% increase in terms of complexity
    - Companies are getting better at catching the hackers, so as a consequence, the malware is starting to focus more on evading detection
  - Malware is modular
    - One piece of malware can be installed on a PC, but that malware can download other modules to do various things, depending on what you want to do (ex: steal login credentials, mine Bitcoin), and they can change these modules whenever they need to – it is adaptable malware, which is what makes it dangerous
      - Not like it is just one piece of code and it does what it can do, the hackers have control over it, telling it to load new functionality or just develop new functionality

Interview 7:

- Can't get into any attacks on them
- In his opinion as it related to evolving threats:
  - Corporations are doing a pretty good job at staying up to speed with what's happening in terms of traditional malware (clicking on things, file attachments that come in via e-mail, trying to spoof an address and make it look like the e-mail is coming from somebody at your work
- End users (weakest link) or employees are more educated now than they ever been in terms of knowing what not to click on, etc.
- Malicious actors are now trying to leverage corporations that they've already attacked and compromised, and use that as a pivot point to attack other organizations
  - Once they get a hold of somebody (a vertical that doesn't have a lot of money to spend on security – school systems, police force, municipal entities, fire

department) and compromise their corporation, they are able to use their email to send e-mails to other corporations, so it looks completely legit (coming from a legitimate e-mail address)
    - Then leverage this to launch the attacks and typically instead of sending the full-blown malware out of the gate, they will send bits and pieces of it, and do some analysis and evade your system
- These are the biggest threats right now – when bad guys are leveraging a compromised corporation's e-mail account to launch attacks
    - You don't trust anybody now – can't trust your business partners; scrutinize traffic from everybody

Interview 8:

- AI is a double-edged sword
    - Can be used for both good and bad

10. What do you think is in store for the future of AI and its impact on the cyber security industry? Do you think AI will replace humans in the future or will humans still be needed to develop the software?

Interview 1:

- Hasn't seen a solution yet where AI is a plug-and-play thing
    - Usually a long learning curve, not for the users but for the AI to learn the environment and baseline all the data
    - Even when that is done, you can still get a lot of false positives where it will flag some random users as high risk, when there is no harm in their behavior at all
    - In the future, the real trick is trying to make these algorithms and these AIs more accurate and that goes back to the idea of having more threat intelligence driven alerting
    - Went to a Splunk Conference in Vegas – every other session people were talking about trying to use threat intelligence to drive your SOC (Security

Operations Center) and then have more useful actionable alerts and hopefully add more accuracy

Interview 2:

- More tools are going to be AI driven, maybe even code that writes itself, or systems that build themselves – there is a lot of use for AI
- Predicting analytics behind what the attackers are going to do, how you can stop it
- Always going to need someone to make the phone call, pull the trigger, respond to something – there are some analytical things that only humans can do

Interview 3:

- Doesn't think it will replace humans in the future
    - Doesn't like to call it AI, likes to call it machine learning
    - There is no intelligence in it, it is about patterns: learning a certain pattern and recognizing that pattern
- He believes that humans are still needed to challenge the machine learning or AI
- The engines are as good as what they learn, but on the other hand you have better creativity in a human, they will come up with something the machine didn't think of
- "As long as the enemy of the AI and machine learning is a human, the human will always win"

Interview 4:

- Customer said they are trying to get to a place where they have basically no analysts – opinions on this vary
    - This customer was under the impression that you can get a tool like this
    - A tool called SOAR (Security Orchestration Automation Response)
        - A software that will either disable a user (put it into a quarantine type thing) so they can't access the network, or it will block certain things, like if an e-mail is shown to be compromised
            - This capability would essentially be able to shut that user off, but that creates some challenges

    - Say the CEO is suddenly potentially being targeted as a malicious user – you don't want the CEO to be cut off from his people, that wouldn't end well for anyone
  - There are those, yes, that think these SOAR type tools will be able to solve everything that humans would, however, in the interviewee's opinion, humans always need to have some kind of level of interaction and involvement
    - Because of scenarios like the example with the CEO, or if it was an IT analyst
- The technology has come a very long way, they can do a lot of neat things and shut down things automatically when they are happening, but there always needs to be a level of analyst involvement

Interview 5:

- AI is going to continue to get used and it's having a lot of success, so it will be continuously developed and will be used in a lot of other areas
- See it in content filters and content distribution networks, and be able to shape where content is stored
  - If there is an effort to reduce storage traffic volumes, shaping where traffic is set to or making sure the data is appropriate for the volume of traffic based on analyzing different algorithms – allows us to do more with less already, which is something everyone is obviously striving for
- Starting to see a lot of companies that are promoting the concept of you can do more with fewer people
- Carbon Black and Cybereason [leading EDR (Endpoint Detection & Response) vendors]
  - One analyst can handle about 200,000 end points as opposed to most other vendors which is a fraction of that (how true that is, again – that is great, but if you don't know how they're detecting it how do you know where the gaps are in your process)

        - Still a big issue you will see in the industry is people demanding to understand better how these things are being trained (what those data sets are)
            - Going to start seeing more organizations that are a lot more open about those things
- Ways away from fully replacing humans
    - Already past the point where it is fully automated, and it is replacing humans for certain actions to be able to block things
        - Network IDS (Intrusion Detection System) and IPS (Intrusion Prevention System)
            - IDS is passive – can generate alarms all day, but it is not going to block anything or prevent anyone from doing their work
            - IPS is the opposite – going to block things automatically, might be blocking something that prevents someone from doing their work
                - It's much tougher to block all the malicious things and have it automated
                - IPS has been around for a long time – had a whole scenario of vendors that were very cautious about what they were going to block for that reason
                    - Now you have all these algorithms and endpoint technology that's blocking stuff all the time based on behavior, and they are removing the human element because they are automatically blocking these things
                    - Also making it very easy to see what was blocked, and then white list it and say "well if it is based on this application, or ran from this directory, or if it was this user or this set of computers, then you could ignore it next time"

- Manually bypassing the algorithms

- Replacing the human element by automatically blocking things – we don't know if its bad or not, and whether they are on the side of caution
    - But it makes it very easy to make it run again if something was blocked

Interview 6:

- No, doesn't think it will replace humans anytime soon
    - Not going to come from the security industry – don't invest enough in AI
    - Looking at other industries (finance, healthcare), they invest a lot more
    - Still 20 years away from full AI
        - Not sure the world is really trying to solve that problem
- There are academic studies that show AI excels at certain types of problems, while humans excel at certain types of problems, and if they work together, they are far more efficient than each one individually.
    - This hybrid is going to become more common and better
- AI doesn't come up with anything new
- There are certain areas (natural language processing, image recognition) where AI is really efficient, but if you look at recent data (ex: Tesla - still lousy at driving autonomous cars) it is still bad at some things
    - It is intended to free up humans to do what a machine can't do
        - Machine and human combination is going to be really interesting and will change everything
- Haven't yet understood what can be done with AI in its fullest potential
    - Normal because of recursive innovation – until you have a certain foundation in place, new innovation doesn't really happen
- With security, when they talk AI, everyone talks about detection (it's the obvious thing and what most people focus on)
    - But, they are pretty good at detecting stuff, it is doing stuff with the detections (automating incident response containment) that they are not doing a lot
        - It will be automatically capturing knowledge from monitoring what people are doing in an investigation and training the machine learning

engine that way so that in the future it can replicate the easy stuff – still ahead of us, there is a bright future

- "Something so powerful like this [AI] takes on power that is hard to predict… We are just at the cusp of something that's going to change everything."

Interview 7:

- AI has impacted and has a role in pretty much every vertical (security, financial services, healthcare)
- Not replacing people, it is doing the job that no one wants to do better than a human, which in turn allows the corporation to leverage the intelligence of a human to do things they weren't able to do prior
    - Get more out of the employee – if a person was spending 50% of their time on repetitive tasks and 50% of their time thinking and helping the business and driving innovation; once you eliminate the repetitive task (real painful thing that you can't do good because you're trying to sift through huge amounts of data, you pass it off to the AI) now you can spend 100% of the time being innovative and bringing ideas to the table
- Don't think it is going to replace people, just maybe get rid of the people who like to do repetitive tasks (ex: people who like to work in a toll booth – replaced by automation so if you don't retrain yourself then you're out of a job)
    - Security is one of those industries that evolve and there are always things you need to do in terms of understanding the threats that are out there and really digging into where a particular attack came from, and the techniques used
- AI is going to do a better job at helping detect malware, weeding out false positives, not at the point where it is going to replace people – "I don't have any fear of that, I'll embrace it"
- Example
    - In financial services 7 seven years ago, they started the question, "Do you need a mutual fund manager?"
        - Mutual fund manager manages a bunch of stocks – say they have 100 stocks that they manage in a portfolio for their customer

        - Along comes AI, and they look at how the fund manager picks stocks (method that every fund manager leverages – maybe they look at their sales, expenses, news in the market to see how a company performs, and if there is something in the news that no one else knows about) and does it itself
            - AI could do things faster – act fast if you see news that could impact purchases
- A lot of things can be automated, but people do fear it
    - You have to fight a lot of uncertainty and doubt when it comes to AI, machine learning, and automation
    - AI has broken the culture of "AI is going to fail because of this or this" and is now at the forefront to the point where if you aren't using AI you can't help an organization much with security
        - Has taken about 10 years to get to this point
            - A lot of people have gotten out of IT because of it
            - The ones that are in IT now are the "thinking people"
- With security, you have to stay ahead of it and constantly be learning
    - Some folks want to learn one thing and never evolve

Interview 8:

- No more signature detection
    - Can't do it anymore because the attacks are so sophisticated and you can't stop them with humans anymore, so it is all AI
- Every modern security team uses AI in one way or another
- It already has in a way replaced humans
- Humans will still be needed to develop the code
- A lot of code has been written in the past 5-10 years, so it is just taking jigsaw pieces and putting it together/replacing it to build more code
- Automation in code building is available, and will only improve

REFERENCES

Addo, A., Centhala, S., & Shanmugam, M. (2019). *Artificial Intelligence for Risk Management*. Business Expert Press. http://ebookcentral.proquest.com/lib/bryant/detail.action?docID=6134046

Addo, A., Centhala, S., & Shanmugam, M. (2020). *Artificial Intelligence for Security*. Business Expert Press. http://ebookcentral.proquest.com/lib/bryant/detail.action?docID=6134047

AI likely to replace humans in cybersecurity space by 2030. (2021). *FRPT- Telecom Snapshot*, 20–20.

Diogenes, Y., & Ozkaya, E. (2018). *Cybersecurity, attack and defense strategies: Infrastructure security with Red Team and Blue Team tactics*. Packt Publishing.

Erdogan, G., Hugo, Å., Romero, A., Varano, D., Zazzeri, N., & Žitnik, A. (2020). An Approach to Train and Evaluate the Cybersecurity Skills of Participants in Cyber Ranges based on Cyber-Risk Models: *Proceedings of the 15th International Conference on Software Technologies*, 509–520. https://doi.org/10.5220/0009892105090520

Hall, D. (2021). 4 Benefits of Using AI in Cybersecurity | CIO Insight. *CIO Insight*, N.PAG-N.PAG.

Halsey, J. (2021). Future of Control System Cybersecurity Built Upon Industry Standards. *Pipeline & Gas Journal*, *248*(10), 33–35.

Hautamaki, J., Karjalainen, M., Hakkinen, P., & Hamalainen, T. (2019). *CYBER SECURITY EXERCISE – LITERATURE REVIEW TO PEDAGOGICAL METHODOLOGY*. 3893–3898. https://doi.org/10.21125/inted.2019.0985

Hunter, A. (2020). Ai May Be the Solution to Defence Industry Cybersecurity. *Armada International*, *1*, 34–34.

Kaloudi, N. & Jingyue Li. (2020). The AI-Based Cyber Threat Landscape: A Survey. *ACM Computing Surveys*, *53*(1), 1–34. https://doi.org/10.1145/3372823

Labs, W. (2021). Advanced Intelligence: Machine learning and AI help processors outsmart cybercriminals. *Food Engineering*, *93*(4), 44–49.

Maguire, J. (2022). Tech Predictions for 2022: Cloud, Data, Cybersecurity, AI, and More. *EWeek*, N.PAG-N.PAG.

McClurg, J. (2020). Mind the Gap: Diversity & Other Challenges in the Age of AI. *Security*, *57*(12), 33–37.

Simonovich, L. (2021). Balancing AI advances with robust cybersecurity solutions. *World Oil*, *242*(9), 55–58.

Sophos Announces 4 New Artificial Intelligence Developments. (2021). *Reseller Middle East*, *283*, 39–39.

Xin, Y., Kong, L., Liu, Z., Chen, Y., Li, Y., Zhu, H., Gao, M., Hou, H., & Wang, C. (2018). Machine Learning and Deep Learning Methods for Cybersecurity. *IEEE Access*, *6*, 35365–35381. https://doi.org/10.1109/ACCESS.2018.2836950

Yamin, M. M., & Katt, B. (n.d.). *Inefficiencies in Cyber-Security Exercises Life-Cycle: A Position Paper*. 3.

Yampolskiy, R. V. (2017). AI Is the Future of Cybersecurity, for Better and for Worse. *Harvard Business Review Digital Articles*, 2–4.

ADDITIONAL RESOURCES

Columbus, L. (2023, February 24). *Experts predict how AI will energize cybersecurity in 2023 and beyond*. VentureBeat. Retrieved April 17, 2023, from https://venturebeat.com/security/experts-predict-how-ai-will-energize-cybersecurity-in-2023-and-beyond/

Drolet, M. (2023, January 3). *Council post: Six cybersecurity trends you can expect in 2023*. Forbes. Retrieved April 17, 2023, from https://www.forbes.com/sites/forbestechcouncil/2023/01/02/six-cybersecurity-trends-you-can-expect-in-2023/?sh=1bdfb49a4c97

Gartner_Inc. (n.d.). *Beyond chatgpt: The future of generative AI for enterprises*. Gartner. Retrieved April 17, 2023, from https://www.gartner.com/en/articles/beyond-chatgpt-the-future-of-generative-ai-for-enterprises

Gartner_Inc. (n.d.). *Gartner top security and risk trends in 2022*. Gartner. Retrieved April 17, 2023, from https://www.gartner.com/en/articles/7-top-trends-in-cybersecurity-for-2022

Gartner_Inc. (n.d.). *What it takes to make ai safe and effective*. Gartner. Retrieved April 17, 2023, from https://www.gartner.com/en/articles/what-it-takes-to-make-ai-safe-and-effective

Gartner_Inc. (n.d.). *What's new in Artificial Intelligence from the 2022 Gartner Hype cycle*. Gartner. Retrieved April 17, 2023, from https://www.gartner.com/en/articles/what-s-new-in-artificial-intelligence-from-the-2022-gartner-hype-cycle

Gartner_Inc. (n.d.). *Why adaptive AI should matter to your business*. Gartner. Retrieved April 17, 2023, from https://www.gartner.com/en/articles/why-adaptive-ai-should-matter-to-your-business

Gartner highlights three ways security leaders can prepare for the evolution of cybersecurity strategy, roles and Technology. Gartner. (n.d.). Retrieved April 17, 2023, from https://www.gartner.com/en/newsroom/press-releases/2022-06-07-gartner-highlights-three-ways-security-leaders-can-prepare-for-the-evolution-of-cybersecurity-strategy-roles-and-technology

Gartner Hype Cycle Research Methodology. Gartner. (n.d.). Retrieved April 17, 2023, from https://www.gartner.com/en/research/methodologies/gartner-hype-cycle

Gartner unveils the top eight cybersecurity predictions for 2022-23. Gartner. (n.d.). Retrieved April 17, 2023, from https://www.gartner.com/en/newsroom/press-releases/2022-06-21-gartner-unveils-the-top-eight-cybersecurity-predictio

What we do and how we got here. Gartner. (n.d.). Retrieved April 17, 2023, from https://www.gartner.com/en/about