NOVEMBER 2020

Destructive Cyber Operations and Machine Learning

CSET Issue Brief

CSET
CENTER *for* SECURITY *and* EMERGING TECHNOLOGY

AUTHORS
Dakota Cary
Daniel Cebul

Executive Summary

Cyber operations that impact the physical world rely on attacks against industrial control systems. These are the operational systems that control production lines, electrical plants, and critical infrastructure. Industrial systems' distinct structures, proprietary communication protocols, and blend of operational technology and information technology make attacking such systems a tall-order for cyber operations, yet machine learning could alter the nature of offensive operations.

Machine learning may change cyber operations against industrial systems in three ways.

First, modeling the industrial process using machine learning may decrease the number of failed attacks by advanced actors. This capability will make the good attackers better, but not improve the operations of less sophisticated attackers.

Second, machine learning models can serve as weapon for attacking industrial control systems (ICS). Fake sensor readings, generated by a model trained on the network's data, can cause the system to adjust itself in ways that cause damage to the system or the goods it is producing.

Third, adversarial machine learning can falsify data that hides ongoing attacks from ML-based anomaly detection systems—allowing some attacks executed by traditional malware to proceed without being detected. This same attack methodology can also create false alarms that desensitize human operators to alerts of an actual attack.

This issue brief offers three policy recommendations.

1. **Protect the data historian.** If attackers use machine learning models to prepare attacks against industrial systems, the data historian—a repository for some industrial data—will become more important to defend, as its contents are used to train such models. Collaboration between critical infrastructure operators, national laboratories, the private sector, the Cybersecurity and Infrastructure Security Agency, and the National Security Agency may yield more tailored technical solutions that can be deployed widely across industrial sites.
2. **Field an ICS Hunt Team.** The federal government should train and use a corps of ICS experts dedicated to proactive threat hunting in

ICS environments. Proactive defense of operational technology networks by federal officers, in collaboration with willing critical infrastructure operators, reinforces the defender's advantage. Threat hunting capitalizes on the long periods attackers must dedicate to reconnaissance in industrial networks and uses intelligence collected under persistent engagement by Cyber Command and the NSA to detect attackers.

3. **Bolster defensive research.** The U.S. government should support additional research into the potential malicious uses of machine learning in attacks against industrial systems, with the specific goal of identifying weaknesses in attack methodologies. Findings from additional research could drive collaboration between industry and government and the development of technical solutions.

Table of Contents

Introduction

The fear of a devastating cyber attack on physical infrastructure is older than the malware itself. For decades, theorists imagined how malicious code could turn off the lights, manipulate industrial machinery, and cause explosions. These systems were the bridge between the digital and the physical, the way in which nebulous code could have a real-world impact.

Then it happened. The Stuxnet operation against Iran's Natanz nuclear facility, discovered in 2010, showed that cyber attacks *could* cause physical damage. The operation sabotaged centrifuges, which failed without apparent cause. As cybersecurity policy and scholarship began to emerge, Stuxnet provided an exemplar to theorists. Everyone wondered what piece of code was next.

The answer was CRASHOVERRIDE. Launched by a unit of Russian military intelligence, the malicious code manipulated the industrial processes of power systems in Kyiv, Ukraine, plunging the city into darkness for several hours. For as notable as the attack was—the first known blackout caused by code—its lack of lasting damage meant that the same question emerged: technically and geopolitically, what comes next?

Geopolitically, the answer is apparent. The pace of attacks has quickened, as nations have turned to cyber attacks to advance their interests. In early 2020, Israel defended a municipal water treatment plant from an industrial attack by Iran and responded by attacking the computers that administer one of Iran's ports; the attack reportedly overflowed canals and flooded roads leading to the port terminal responsible for more than half of Iran's seabound trade.[1] The most recent United States National Counterintelligence Strategy identified safeguarding critical infrastructure as one of its five priorities, specifically acknowledging the national security risks that cyber attacks on industrial systems pose.[2] Perhaps worst of all, the field of adversaries is growing. A leading U.S.-based industrial defense company, Dragos, identified three new threat groups actively targeting industrial systems in 2019, bringing their total to eleven.[3]

Technically, the answer is more complex. This paper considers one important possibility. It asks: How will artificial intelligence (AI) and machine learning alter the destructive power of such attacks? This paper answers the question by examining malware used in past ICS attacks and analyzing ongoing research in the field of AI and ICS attacks.

We find that, for the best threat groups, using machine learning to model a target's environment and simulate attacks will decrease the number of failed attacks. In addition, some machine learning frameworks will create new attack methodologies and help traditional malware circumvent defensive systems, enabling stealthy attacks. The ability to conduct these attacks with a lower chance of failure and detection may make them a more attractive tool of statecraft.

Background on Attacking Industrial Control Systems

A Brief Overview of Industrial Control Systems

Industrial control systems are the operational technology (OT) that underlie industrial processes. The OT system consists of machinery and software responsible for monitoring and controlling the industrial process. This system works by storing and transferring electronic data throughout the ICS network like an information technology (IT) system distributes data across a network of computers. In the same way computer scientists and IT technicians rule the IT domain, engineers oversee OT networks. ICS are as varied as the industrial processes they control, ranging from systems that maintain the correct acidity of an at-home pool to nuclear reactors with thousands of components. With the exception of pre-designed kits, like pool systems, ICS configurations are often unique to their purpose and location. A company that manufactures the same widget in two separate factories may use the same manufacturing process at both locations, but the ICS configuration could be vastly different. Individual attributes such as the layout of the factory floor, the differing regulations by nations or states, and the deployment of machinery built by different manufacturers all impact how engineers construct industrial systems.[4] The unique configuration of industrial systems creates a natural barrier to reproducing cyber attacks at multiple sites.

There are just a few leading industrial automation companies. According to one market report, just 10 companies held 67 percent of the market for global factory automation in 2019.[5] Siemens' Digital Industries division holds the largest market share (20 percent) for both the industrial software and factory automation sectors.[6] Some of the most advanced cyber attacks may work against different industrial systems and facilities that use the same underlying technology.[*] Attackers with experience conducting attacks against targets that use Siemens devices may find portions of attacks against other targets using the Siemens instruments easier to construct and execute. For example, a 2017 attack against an oil and gas company in the Middle East targeted a safety device used at 18,000 other locations, though the malware is not directly transferable to new targets without modification.[7] The relative concentration of market share at the top betrays the competition in the remaining 37 percent, which is divided amongst many more companies. These companies and their products, which are not as dominant as their competitors, reinforce the problem of uniqueness of industrial systems.

The remainder of this section introduces components of an ICS. The following terms are important to understand how the ICS technology works. In a plant producing the fabled "widgets" of economic textbooks, the operational technology of the ICS typically includes:

- *Sensors*: these devices monitor measurable aspects of a system. For example, sensors monitor the temperature of metal sheets used to stamp widgets or the pressure of gas inside a pipeline.
- *Actuators*: these pieces of machinery are responsible for changing some aspect of the industrial process. For example, an electric coil used to increase or decrease the temperature of the metal sheets is an actuator, as is the motor responsible for pumping gas at a specific rate into the pipeline. Another actuator could be a network-connected safety instrumented system, which alters part of the process when specific, unsafe conditions are met.
- *Programmable Logic Controllers* (PLCs): these devices receive data from the sensors and send commands to the actuators. The controllers' decision-making process is programmed by human operators. PLCs can receive inputs and issue commands to hundreds of sensors and actuators.[8] For example, if the sensors register a temperature that is too low for the sheet metal, the controller might command the actuator to increase heat. Similarly, if the sensors in a pipeline detect low

[*] Discussed in more detail later, the CRASHOVERRIDE malware seems transferable between systems which use similar infrastructure according to reporting by Dragos. Though, much rejigging of the particulars is required to do so.

pressure, the controller could command the pumps (which are actuators) to increase the flow of gas.

- *Human Machine Interfaces* (HMIs): these screens display the data from sensors and ongoing processes to operators monitoring or adjusting the industrial activity. If the boss determines that higher-quality widgets can be produced by using higher temperatures, the operators would use the HMI to re-program the PLCs, which in turn would send new commands to the actuators, thus increasing the temperature. There are many variants of HMIs and this paper considers any workstation designed for human inputs as an HMI.[9]

Figure 1. Screenshot of an HMI.[10]

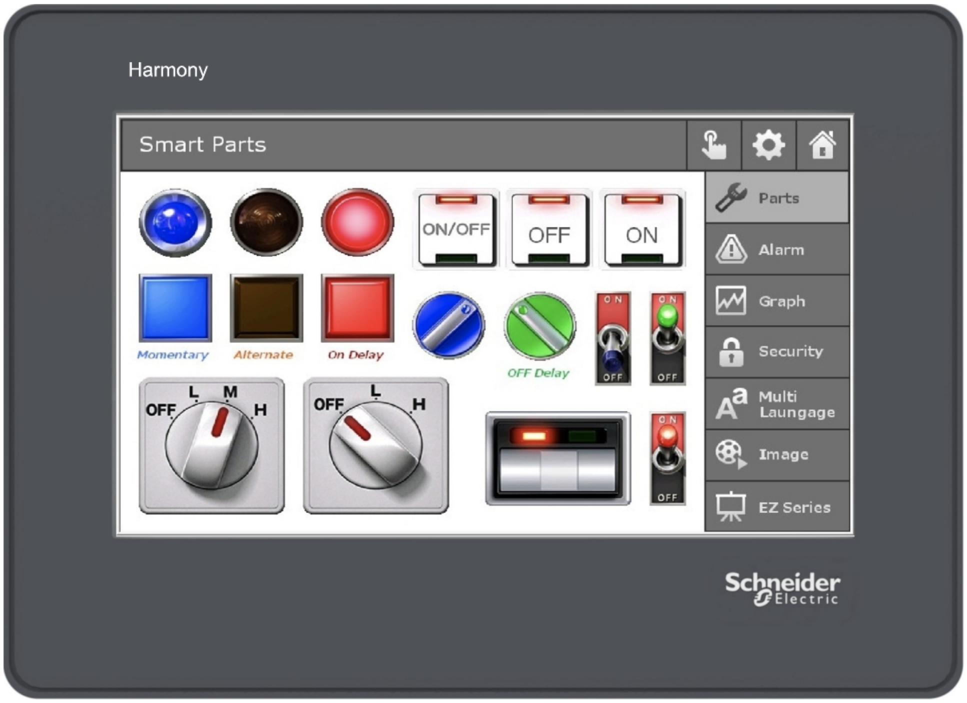

- *Data Historian*: this database collects and stores data generated by the ICS. Operators can select which data is collected and at what frequency. Sensor readings, PLC commands, actuator actions, and changes issued by an operator at the HMI can be recorded by the historian. Operators may examine the records in the historian to review past operations at the plant, search for more efficient ways to produce widgets, or determine if any changes have been made to the industrial process.

Attacking an ICS

Cyber attacks on industrial systems vary by their objectives and sophistication, creating a wide spectrum of effects an attack could have on its target. Leading ICS defense researchers argue that conducting precise attacks requires more time, skills, knowledge, and resources than simple, disruptive attacks.[11] Some attacks, like ransomware that holds data for ransom or wipers that overwrite critical software, may only require the attacker to gain network access before activating their malware.[12] Attacks with more complex goals require more intelligence collection. Attackers who want to impact the ICS in specific ways (e.g., Stuxnet) must be willing to endure the significant costs of preparing to execute such an attack.

Cyber attacks against an ICS occur in two distinct stages. Attackers must first conduct a traditional cyber operation to gain access to the IT network before moving into the industrial system. The second phase uses collected information to develop and execute the attack against the OT network.[13] To attack industrial systems successfully, attackers must possess the skill of a traditional hacker and the knowledge of an industrial engineer. The multi-stage attack process contributes significantly to the difficulty of conducting attacks against industrial systems.

All ICS attacks designed to have a destructive impact share some basic characteristics. At the most fundamental level, threat groups must manipulate the industrial process to have a destructive effect and prevent defenders from responding in time to halt the attack.[14] The way in which attackers achieve these two objectives varies across attacks. Figure 2 illustrates the options available to attackers; each has its own set of intelligence and capabilities requirements.[15]

Figure 2. Destructive Cyber Attack Concept Tree

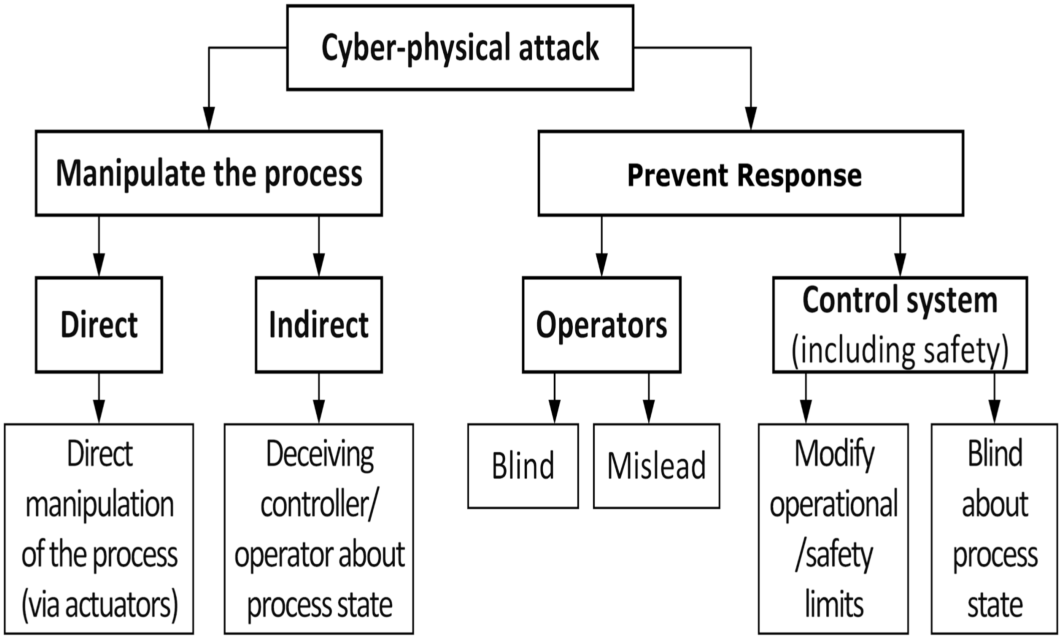
Cyber-physical attack
Manipulate the process
Prevent Response
Direct
Indirect
Operators
Control system (including safety)
Direct manipulation of the process (via actuators)
Deceiving controller/ operator about process state
Blind
Mislead
Modify operational /safety limits
Blind about process state

To cause physical damage, threat groups have two avenues of attack: direct and indirect manipulation of the industrial process. Direct manipulation relies on attackers either sending commands to the actuators that the attackers know will cause harm to the system, or changing the pre-programmed logic of the controllers to cause harm. Attackers disrupted normal operations at the Natanz nuclear facility by issuing direct commands to the PLCs, which changed the speed and pressure of the centrifuges.[16] Indirect manipulation involves changing or falsifying the data that the PLC or HMI receives, causing either the system or the operators to behave abnormally and send commands to actuators that will degrade or destroy the industrial system. An easy-to-grasp example of an indirect attack is one that tries to damage a pipeline by feeding false pressure sensor readings to the PLC, thereby causing the system to increase the pressure inside the pipe beyond safe levels, leading to destruction or triggering a safety mechanism.[17]

To prevent response by the human operators at the industrial site, attackers have a number of options. They can blind the operators by overwriting key files on the HMIs or otherwise shutting those systems down. The attackers can also mislead the operators by sending the HMIs false data about the true state of the ICS. Both avenues have advantages and disadvantages. An operation that chooses to shut down the HMIs must quickly cause physical damage, as the blacked-out computers will be a clear indicator to defenders that something is amiss. Misleading the operators with false data may buy more time for the attack to take place or allow for multiple iterations of the attack, but the attackers will need to anticipate and allow operators to issue commands through the compromised HMI, otherwise unresponsive systems may trigger an investigation.

Misleading defenders can be complex. In addition to human operators, attackers must also bypass or trick the safety components of the industrial process if they wish to have a destructive effect, as these systems have redundant fail-safe measures to prevent accidents and attacks. Malware deployed in previous attacks, like the 2016 Ukraine blackout carried out by Russian military hackers, attempted (and failed) to disrupt the normal function of protection features built into the grid.[18] Not all safety features can be impacted by cyber operations, however. Mechanical systems disconnected from the OT network, like a pressure release valve on a pipeline, can limit the amount of destruction achievable and are not vulnerable to cyber attack. Collecting the intelligence required to attack network-connected safety equipment or avoid attacks that trigger mechanical fail-safes can be a daunting task.

Even advanced actors encounter problems when attacking industrial systems. Understanding the target's industrial system is critical to a successful attack, but piecing together a target's system is a challenge, since seemingly small mistakes can have significant consequences for an operation's effectiveness. The first stage of an attack, reconnaissance, can even be perilous. An instance of the Havex malware family, which can gather information about ICS networks, was reportedly detected after it accidentally triggered an emergency shutdown at an industrial plant in 2014.[19] Seemingly successful attacks are not exempt from failures; post-mortems have often discovered unrealized potential. For example, the late 2017 attack against an oil and gas plant in Saudi Arabia was derailed by incorrectly written malware.[20] After-action analysis showed that an error in the malware code forced the network-connected safety system to enter a fail-safe state and initiate a shutdown. If not for this error, attackers would have gained access to safety controllers and could have inflicted physical damage rather than the unintended short-term shutdown without lasting damage that resulted from the incident.[21] The complexity of industrial systems can undermine the attacks of even the titans of cyber operations.

The Significance of Reconnaissance

Reconnaissance is important to virtually all cyber attacks, but especially attacks against industrial systems. In cyber operations against IT networks, reconnaissance helps attackers determine which malware to develop or deploy against a target. In cyber operations against OT networks, there is a wide variety of machines, many manufactured by different vendors and communicating with obscure or proprietary protocols.[22] Understanding these various components is critical to designing the actual attack.

Reconnaissance is further complicated by the need to achieve *process comprehension*—industry jargon for having a nearly complete understanding of the industrial system, critical to designing an attack that achieves the desired impact.[23] Making a pipeline explode under pressure requires knowledge about how much pressure is required, where the pipeline is located on the network, which pumps are responsible for that section of piping, what communication protocol is required to make those pumps increase pressure, and whether there are any safety features to disarm. The only way to reach complete or nearly complete process comprehension is by a significant investment in target network reconnaissance. Past cyber operations against industrial systems have required attackers to stay on target for as long as seven months before conducting their attack.[24]

Destructive attacks on industrial systems with precise effects rely on two approaches to solving the process comprehension problem. Attackers can either pursue full process comprehension and potentially build a model of the target's environment, or they can create malware designed to leverage their collected intelligence and overcome gaps in that intelligence with guesses and approximations. Under a full comprehension approach, attackers can identify each piece of equipment that will be impacted and explicitly program each step of the attack. On the other hand, malware that relies on inference and automated mapping tools does not have its exact targets in advance but is able to make determinations about which devices it is supposed to target. Both approaches have success cases.

The Stuxnet attack against Iran's Natanz nuclear facility demonstrated the advantages and disadvantages of the full process comprehension approach. In one of its versions, Stuxnet destroyed centrifuges by changing the speeds at which they were rotating.[25] Before the attack was executed, Stuxnet was reportedly tested on many machines involved in the nuclear enrichment process.[26] This was undoubtedly costly and laborious. The process comprehension required to create a viable sequence of tests is both the most strenuous and most assured way to prepare a cyber attack. Such preparation and understanding of the target yield greater confidence that the attack will work when launched.

CRASHOVERRIDE, the malware behind the 2016 Ukraine blackout, overcame gaps in process comprehension by using a blend of collected intelligence and searching capabilities. By constructing the malware with specifically-purposed modules, the attackers could upload new tools to the malware and expand attack capabilities as necessary. Two modules relied on information collected from the attackers' first strike against Ukraine in the year prior.[27] Another pair of modules employed techniques to circumvent the attackers' intelligence gaps; these modules used brute force to guess and identify the names of devices connected to the ICS network. With this capability to manage uncertainty about the target's configuration, the malware identified and attacked systems on the network of UkrEnegro—the Ukrainian power company—causing a short blackout of the Kyiv electrical grid. The highly automated process of finding and attacking targets on the industrial network without full process comprehension represented a significant step in ICS attack capability, though the attackers still had to place the highly automated search modules onto specific parts of the targeted network to enable their function.[28] The modularity of CRASHOVERRIDE and its capacity to overcome gaps in intelligence is a sign of things to come.

How Machine Learning Can Aid ICS Attackers

Machine learning and AI may help attackers plan and rehearse destructive attacks, facilitate long-term attacks that disrupt services or the normal production of goods, and decrease the effectiveness of intrusion detection systems. While public research in this area remains limited, this report offers a preliminary, technically grounded analysis of how ML may amplify ICS attacks.

Planning and Rehearsing Attacks

Machine learning may help attackers model their target's environment, reducing the amount of time and resources dedicated to process comprehension. Attackers with better machine learning-enabled modeling capabilities could shift focus from just collecting intelligence about the target network's architecture and its components towards collecting data that would be used to train a model representing the whole system. While the scope of reconnaissance would not narrow, time spent stitching together collected information would decrease. Such a capability would also allow attackers to plan and rehearse their attack sequence on their own networks, determining which kinds of attack commands have the desired effects and increasing the probability of successful operations. Given its complexity, likely only the most advanced attackers would be able to adopt an improved modeling capability to improve operational outcomes.

A group of researchers modeled the industrial system of a water treatment testbed using two machine learning frameworks and the data historian.[29] This digital recreation of the industrial system allowed the researchers to use a second algorithm to execute random combinations of settings, observing which combinations led to the quickest overflow of the targeted water tank.[*] The researchers identified 27 different attacks that caused the simulated tank to overflow. Further testing demonstrated the transferability of the attacks from the simulated environment to the actual testbed, though the testbed in question focused on computer systems and did not incorporate offline mechanical protection features that might have mitigated or prevented destructive effects.

It is notable that the researchers were able to train their machine learning model of the industrial system with just the information stored in the data historian. This reduces the need to conduct further reconnaissance in other parts of the industrial network.[†] In short, the data from the data historian gave a machine learning system the information it needed to understand how the target system worked, and how it could break. While attackers would still need to contend with any hidden mechanical industrial protection and safety features, the proof of concept shows how machine learning can enable a

[*] The authors used a genetic algorithm to create a fuzzing tool for the LSTM and SVR models generated with the data historian. The authors defined the fitness function for assessment of the results by representing the height of the water table tracked by ICS sensors.

[†] Commercial tools to model industrial processes exist, but none offer the ability to simulate an environment with just data from a network historian. This is likely due to the fact that industrial system operators prefer more complete system models with comfortable user interfaces and interoperability with the system, like that of an HMI. Moreover, these tools are built to be used by consumers with legitimate access to industrial facilities, and therefore access to all relevant data.

better understanding of target ICS networks and a better capacity to attack those targets.

New internet-connected devices will add layers of complexity to ICS. Often called Industry 4.0, the increasing number of "Internet of Things" devices will increase the attack surface and overall vulnerability of industrial systems if not managed properly. Ericsson—a maker of 5G equipment—is deploying 5G technologies in its own factories. Manufacturers are widely expected to do the same, increasing the number of potential vulnerabilities in internet-connected devices. Attackers may turn towards machine learning to aid in process comprehension just because of the sheer number of interconnected devices coming online. ML modeling of the target's ICS might therefore become more important to ensuring process comprehension.[30]

Only states that can already conduct attacks on ICS will be able to reap the benefits of using ML in this way. Even with machine learning models, attacks will need to write code that manipulates the many proprietary communication protocols used by ICS machines. Moreover, attackers must still be capable of exfiltrating data from the historian or industrial network. Implanting that malicious code onto the network at the correct locations will require further skill. Modeling the industrial process using machine learning may decrease the number of failed attacks by advanced actors, but it will not obviate the need for industrial expertise or reduce the skills required to conduct an attack. This capability will make the good attackers better, but not improve the operations of less sophisticated attackers.

Long-Term Degradation and Stealthy Attacks

Machine learning may create new indirect attack methodologies that support longer, hard-to-detect campaigns. An indirect attack can achieve an ideal strike against an industrial system—one where the alarms do not go off, the system still responds to operator commands issued at the HMI workstations, and any investigation does not immediately conclude that a cyber-physical attack is underway. Stuxnet has long served as the best example of a stealthy attack, as no other reported attack has been comparably stealthy in duration or degradation. Research into the application of machine learning and industrial attacks may help future attacks build on Stuxnet's enduring success.

It is possible to conduct an indirect attack by training an AI agent to produce fake data and cause an industrial system to attack itself. In 2017, a team of researchers trained a variant of the well-known framework responsible for deepfakes—known as a generative adversarial network, or GAN—on the data sent between the PLC and the actuators and sensors.[31] GANs are composed of two competing machine learning systems: a generator and a discriminator. In the research, the generator produced false sensor readings, while the discriminator simulated the defender's intrusion detection system. The two networks competed against one another until the generator developed the ability to produce false sensor readings that fell beneath the real intrusion detection system's threshold for triggering an alarm. Such a capability allows attacks to proceed without being detected. To demonstrate the capability, the attackers specified a desired attack and the GAN spoofed the data sensors sent to the controller, causing it to respond with commands that ultimately changed the industrial process in a stealthy and degrading way.

The researchers further showed their attack's ability in two testbeds: a water treatment plant and a gas pipeline.[32] In the water treatment plant testbed, the researchers were able to change the quality of the water produced by altering sensor data used to inject treatment chemicals. The same attack process was used to generate artificially low pressure readings in the pipeline testbed, which caused the controller to increase pressure within the pipe beyond safe limits. The attacks succeeded without setting off the testbed's intrusion detection system.[*]

This method of attack has two distinct limitations, however. The first limitation is imposed by the intrusion detection system. Researchers were able to remain undetected because the fake data remained under a certain threshold of statistical deviation from normal operations that would have triggered the alarm. However, if the alarm on the real target's intrusion detection system was more sensitive or the system did not already perform near unsafe levels, the GAN may be unable to produce fake data that can cause the desired attack state.[33] Attackers using this method will have to calibrate their efforts to have the maximum desired effect and not trip alarms. This balancing act may limit this methodology's potency.

[*] The GAN framework was selected for its ability to both generate and discriminate data—the discriminator component of the GAN trained the generating component to only spoof data in a way that would not be detected while achieving the stated goal of the attackers. In this way the discriminator played the role of an intrusion detection system.

Second, this attack methodology is constrained by the destructive potential of the ICS itself. Most industrial attacks are carried out simultaneously at many points in the network.* Yet, because this attack method relies on spoofing data at one controller instead of many, its physical impact is limited to the systems controlled by that PLC. It is technically feasible that many attacks like this could happen in tandem, leading to a system-wide impact, but the requirements of coordinating such an attack may be insurmountable. These diffuse ML agents would have to communicate over machine-to-machine protocols, which many defensive systems and engineers monitor for anomalous behavior. If two machines which have no need to communicate begin sending messages to one another, defenders will know something is amiss. Moreover, an uncoordinated attack at many points could trigger the intrusion detection system, pique engineers' interest, or create anomalies in process monitoring systems.

There are sure to be a number of attack scenarios that could cause significant damage with just one PLC, however. As with the water treatment testbed, this may take the form of an attack that causes pumps to inject too much chlorine—something achieved by adjusting the flow rate of the chlorine pump with false data injections. Such an attack objective is realistic; Israel reportedly experienced an attack that tried to adjust chlorine levels at a water treatment plant, though the attack did not seem to use false data injection.[34]

Adversarial Machine Learning Attacks on Intrusion Detection Systems

Defenders of industrial control systems often deploy anomaly detection software to spot abnormal behavior on the network.† Some of this software uses (or may use in the future) machine learning to model normal operations and detect intrusions.[35] Intrusion detection systems recreate a statistical baseline for normal operations and flag any anomalous behavior which

* This is seen in CRASHOVERRIDE and Stuxnet. Attacks which have been disruptive, but did not cause physically destructive effects, often impact just one part of the industrial process.

† These anomaly detection systems are different from the systems discussed below, as they monitor for abnormal network behavior (packets traveling in greater volume or abnormal patterns) rather than the performance of the industrial system.

might suggest that the system is not operating within its normal bounds.* A group of researchers demonstrated that these models may be vulnerable to manipulation through adversarial attacks, just as other machine learning systems are.[36] An adversarial attack uses information about the target model—which can be learned by probing the target with inputs and recording the responses—to create an input (like ICS data) that causes the target to fail at its job.[37] The researchers trained an ML model using adversarial learning techniques to generate false sensor data which could deceive the defensive systems and help enable an attack.

The adoption of ML-assisted intrusion detection systems will likely increase, as the market for AI solutions and industrial automation is estimated to reach $17.6 billion in 2021.[38] As a case in point, IBM announced in 2019 the sale of an AI-infused edition of its industrial systems management software, Maximo Asset Performance Management, which the company claims improves predictive maintenance, anomaly detection, and performance optimization of industrial systems.[39] As ICS integrate more connected devices, operators may derive more financial gains by using machine learning to inform industrial design and automate operations, speeding the adoption of the technology and pushing human defenders out of the loop.[40] Indeed, easing IoT integration and optimizing system performance is central to the marketing literature for IBM's system.[41] Attacking these integrated machine learning applications with adversarial techniques may yield a new vector of attack against industrial systems, just as the management of these processes is being handed over to software.

Though the specific vulnerabilities and defenses may vary across targets, the researchers demonstrated two general approaches to deceiving ML-assisted intrusion detection systems. The first approach falsified sensor data that was designed to fool the intrusion detection system and conceal an ongoing attack. The researchers trained their adversarial attack model on the normal data passing between the components of the ICS. After a period of time, the malicious ML model was able to discern what fake data might or might not pass as normal. After training on the target's data, the adversarial framework was ready to be tested. While the adversarial attack blinded the intrusion

* The defensive models and their potential impact on ICS will be discussed in a forthcoming paper on the impact of ML/AI on ICS defense.

detection system, other components of the malware made malicious changes in the industrial process. The fake data created by the adversarial attack acted as a camouflage for the real attack being carried out simultaneously.[42] If human operators were to examine the screens of the HMIs at the plant, nothing would seem abnormal. The proof-of-concept provides clear evidence of the vulnerabilities inherent to intrusion detection systems based in machine learning.

The second attack methodology used adversarial machine learning to decrease human operators' responsiveness to the anomaly detection system.[43] By causing slight, but abnormal, perturbations in the data, the researchers were able to create many false alarms without actually altering the industrial process. Whereas the first attack camouflaged the ongoing assault in fake data, this second method tricked the intrusion detection system into sensing attacks where there were none. All these seemingly false alarms would probably lead defenders to either decrease the sensitivity of the defensive system or start ignoring the alarms out of "alert fatigue." Attackers could then exploit any cognitive overload introduced by the false alarms. Excessive alarms have contributed to accidents and their inadequate response; most famously, Three Mile Island's nuclear accident was in part facilitated by many concurrent alarms, obscuring operators' understanding of the system and contributing to cognitive overload.[44]

Conclusions

The decision to adopt ML to conduct or support ICS attacks will be based on similar criteria for any new technology: will the technology make the task easier to achieve or cheaper to execute, or meaningfully increase the user's ability? In this case, the answer is likely yes when advanced attackers want to achieve a specific outcome rather than cause general disruption. Using machine learning in attacks on ICS still requires significant knowledge of industrial systems and even *increases* the requisite technical competencies to create an attack.

For the small group of attackers capable of using machine learning for cyber operations, the technology could enhance their efforts by increasing success rates, speeding up attack development, and simplifying the conduct of sophisticated attacks. First, attacks may not fail as frequently. As demonstrated earlier, there are benefits to modeling the targeted industrial system with machine learning. If machine learning models give attackers more confidence in their ability to launch successful ICS attacks, they may choose to do so more frequently.

Second, machine learning may enable attackers to design and deploy operations against ICS more quickly. Attacks against industrial systems have most often required dedicated reconnaissance and the successful interpretation of pilfered data.[*] By modeling the ICS with machine learning, attackers might reduce the amount of time spent consuming and processing intelligence about the ICS's operations, speeding up attack development and execution.

Third, long-term degradation and stealthy attacks may become easier to conduct. Though constrained by their scalability and the effectiveness of intrusion detection systems, machine learning systems offer the ability to spoof industrial data while accounting for the regular operations of the industrial system. Research has demonstrated the ability of attackers to outmaneuver machine learning-enabled intrusion detection systems, a development which could increase the persistence of long-term attacks.[†] The ability to conduct long-term degradation or stealthy attack campaigns may prove an attractive investment for advanced attackers.

Not every advanced attacker will have the motivation or resources to use machine learning to aid cyber attacks against industrial systems, however. Some attackers may prefer to develop – or continue developing – attacks that are easily transferrable between targeted ICS. A lack of competency in machine learning, the nature of the targeted systems, or the geopolitical goals of the attackers themselves may prompt attackers to develop transferable attacks. Attacks that work across industrial sites would enable attackers to strike industrial facilities quickly and without conducting additional time-consuming reconnaissance and development. CRASHOVERRIDE demonstrated attackers favoring malware that could overcome gaps in intelligence collection and be used against other electrical grids after significant modifications. This pursuit of transferability may reduce the

[*] CRASHOVERRIDE was notable in its ability to combine what was known with tools to overcome gaps in intelligence about the target's network.

[†] The topic of Adversarial Machine Learning attacks against intrusion detection systems is covered in the next paper in this series about ICS Defense and ML.

certainty of having a specific effect. As discussed earlier, the wide variety of industrial configurations may constrain transferable attacks to those which are disruptive, but not destructive.

Of the possible factors that could usher machine learning into the attack process against industrial systems, one stands out: the widespread adoption of Industry4.0 technologies. Defined broadly as adding many IoT devices, connecting industrial systems to the cloud, and using big data to improve everything from the factory floor to the supply chain, Industry4.0 is considered to be the next phase of industrial manufacturing.[45] All of these new internet-connected devices and computers portend a larger attack surface with new vulnerabilities and opportunities for attackers, but also new challenges. As ever-more devices increase the amount of data produced by a single industrial system and asset management software imbeds AI to make use of new data, attackers may also have to turn to machine learning to aid process comprehension or camouflage their attacks. Ultimately, the target's environment will be the final arbiter of what skills an attacker must possess to be successful. If industrial systems produce more data and use AI to manage operations, attackers may be forced to turn to the technology as well.

Recommendations

Attackers' use of machine learning to augment or conduct new attacks on industrial systems demands innovative defensive strategies. To that end, this report offers three main recommendations. First, defenders must pay closer attention to protecting the data historian, as access to its data could significantly improve the chance of a successful destructive attack. Second, critical infrastructure operators will need to deploy updated defensive strategies and cybersecurity solutions, and a more proactive, collaborative role with defenders should be in the offing. Third, further research must be conducted on the use of machine learning in cyber attacks against industrial control systems. Additional study may illuminate new strategies for defending against discovered attack methodologies.

Protect the Data Historian. Industrial system operators are responsible for protecting the data historian. The federal government has laid out detailed steps to achieve this task.[46] In addition to following federal guidance, critical infrastructure operators may benefit from closer collaboration with federal organizations, particularly the National Security Agency and CISA, to develop and implement novel cybersecurity solutions. The NSA offers a suite of software under its Commercial Solutions for Classified Programs that provides tools to secure the data historian and some components of the IT/OT networks. Though built specifically for classified information assurance, the solutions are available to critical infrastructure operators. Collaboration between critical infrastructure operators, national laboratories, the private sector, CISA, and the NSA may yield more tailored technical solutions that can be deployed widely across industrial sites.

Field an ICS Hunt Team. The federal government should support a corps of ICS experts dedicated to threat hunting in ICS environments. Threat hunting relies on knowledgeable professionals actively investigating network data for evidence of tactics, techniques, and procedures which match known attackers.[47] These threat hunters could be a department of federal officers, a sector-specific team of defenders from participating operators, or a hybrid approach of government and private sector employees. No matter the structure, proactive threat hunting will yield benefits for defenders.

Proactive defense capitalizes on the long periods attackers must dedicate to reconnaissance in OT networks. The U.S. government reportedly collects intelligence on adversaries' capabilities while carrying out its "defend forward" strategy. The National Cybersecurity and Communications Integration Center (NCCIC) offers critical incident response teams for ICS, but these teams are reactive.[48] A structure that either tasks federal officers with threat hunting, or permits greater collaboration between the government and the private sector would operationalize this collected intelligence.

The ability of these officers to operate on multiple OT networks at different industrial sites is crucial to a joint-team's efficacy. Since the variation between OT networks can be significant and almost all are unique, a few solutions exist. First, threat hunting teams—under any eventual jurisdiction or structure—could be organized by sector. By specializing in one specific area of industrial systems, like oil and gas pipelines, officers may be able to provide better advice and more effective services. Second, creating visualizations of the industrial process and network traffic will aid communication among defenders—increasing defense against many types of attacks.[49] With better tools, defenders could quickly understand the industrial process and network communications in new industrial environments. This would allow critical infrastructure operators and government experts to work across many different sites with relative ease by greatly reducing the amount of time spent learning new systems. Private sector industrial system defenders have also advocated for the use of such technologies.[50] Researchers at Pacific Northwest National Laboratory (PNNL) originally identified the visualization tool as a way to bridge the gap between IT and OT network defenders, though it would serve to help any coalition of proactive threat hunters.[51]

There is already precedent for such tools. The National Security Agency has published an open-source tool for ICS network defense called GRASSMARLIN, which helps OT defenders map network connections.[52] Acting on this recommendation would provide operators with new resources to secure our nation's infrastructure. The U.S. National Counterintelligence Strategy advocates for using federal agents to defend critical infrastructure networks and developing visualization tools to aid communication; the CISA's five-year strategy to defend ICS advocates for threat hunting programs in collaboration with critical infrastructure operators.[53] The structure of any such program should be formed in consultation with ICS operators, whose consent and commitment is necessary to operationalize this recommendation.

Bolster attack research. The U.S. government should support additional research into the potential malicious uses of machine learning in attacks against industrial systems, with the specific goal of identifying weaknesses in attack methodologies. Public research on this topic is minimal, but if the findings from the research cited here are any indication, integrating machine learning into destructive cyber attacks against industrial systems could prove a significant development for attackers. PNNL and the Idaho National Lab are studying this question from the defensive perspective. Research conducted from the offensive perspective may yield as many results. Findings from additional research could drive the development of technical solutions by collaborative development between industry and government.

The recommendations here seek to meet the stated strategic goals of U.S. government agencies regarding machine learning-enabled attacks on industrial systems. As is often the case in cyber defense, basic system maintenance and software patching can often make the biggest difference. The NSA and CISA have outlined the steps critical infrastructure operators should take to defend themselves, many of which are acknowledged best practices.[54] Though past policy documents have not had to grapple with the effect machine learning will have on destructive or stealthy cyber attacks, the research suggests that the need for such considerations is not far off. In the coming age of machine learning-enabled cyber attacks on industrial systems, a comprehensive, collaborative approach will be necessary.

Acknowledgments

For feedback and assistance, we would like to thank John Bansemer, Ben Buchanan, Melissa Flagg, Bryson Bort, Jack Lucas, Joe Slowik, Daniel Hague, Drew Lohn, Katya Sedova, Rebecca Gelles, and Cindy Martinez.

Document Identifier: doi: 10.51593/2020CA003

Endnotes

[1] Jiyar Gol, "What Is behind Mysterious Fires at Iran Sites?," *BBC*, July 6, 2020, https://www.bbc.com/news/world-middle-east-53305940; Ronen Bergman and David M. Halbfinger, "Israel Hack of Iran Port Is Latest Salvo in Exchange of Cyberattacks," *The New York Times*, May 20, 2020, https://www.nytimes.com/2020/05/19/world/middleeast/israel-iran-cyberattacks.html.

[2] United States National Counterintelligence and Security Center, "National Counterintelligence Strategy of the United States of America 2020-2022," January 7, 2020.

[3] Dragos, Inc., "The ICS Landscape and Threat Activity Groups; 2019 Year in Review," February 2020. "EnergySource Innovation Stream: Industrial Cybersecurity Solutions - Atlantic Council," accessed October 20, 2020, https://www.atlanticcouncil.org/event/energysource-innovation-stream-industrial-cybersecurity-solutions/.

[4] Sean Moran, *Process Plant Layout*, Second edition. (Amsterdam: BH, 2017).

[5] Statista Ubs, "Estimated Global Factory Automation Market Share in 2019, by Manufacturer," *Estimated Global Factory Automation Market Share in 2019, by Manufacturer*, February 26, 2020, https://www-statista-com.proxy.library.georgetown.edu/statistics/728562/global-factory-automation-market-by-manufacturer/.

[6] Siemens, "Siemens – Business Fact Sheets," Siemens Investor Relations, 2019, https://www.siemens.com/investor/pool/en/investor_relations/equity-story/Siemens-Business-Fact-Sheets.pdf.

[7] Nicole Perlroth and Clifford Krauss, "A Cyberattack in Saudi Arabia Had a Deadly Goal. Experts Fear Another Try," *The New York Times*, March 15, 2018, https://www.nytimes.com/2018/03/15/technology/saudi-arabia-hacks-cyberattacks.html; Dragos, Inc., "TRISIS Malware Analysis of Safety System Targeted Malware," December 13, 2017, https://www.dragos.com/wp-content/uploads/TRISIS-01.pdf.

[8] Chris Foreman, Intel Corp., "Cyber-Security in Industrial Control Systems," Perdue University, accessed June 2, 2020, https://engineering.purdue.edu/VAAMI/ICS-modules.pdf.

[9] CISA, U.S. Industrial Control Systems Cyber Emergency Response Team, "ICS Cybersecurity (301V)," Training Available Through CISA, CISA Department of Homeland Security, 2019, https://us-cert.cisa.gov/ics/Training-Available-Through-ICS-CERT#virtual.

[10] "HMISTO705 - 4.3," Schneider Electric, accessed July 22, 2020, https://www.se.com/ww/en/product/HMISTO705/4.3"-wide-screen-touch-panel,-rs-232-terminal-block/.

[11] Robert M. Lee and Michael J. Assante, "The Industrial Control System Cyber Kill Chain," *SANS Institute Information Security Reading Room*, October 2015, 1–24.

[12] "EKANS Ransomware and ICS Operations | Dragos," Dragos | Industrial (ICS/OT) Cyber Security, February 3, 2020, https://www.dragos.com/blog/industry-news/ekans-ransomware-and-ics-operations/.

[13] Lee and Assante, "The Industrial Control System Cyber Kill Chain"; Eric M. Hutchins, Michael J. Cloppert, and Rohan M. Amin, "Intelligence-Driven Computer Network Defense Informed by Analysis of Adversary Campaigns and Intrusion Kill Chains," *Lockheed Martin*, April 3, 2014.

[14] Benjamin Green, Marina Krotofil, and Ali Abbasi, "On the Significance of Process Comprehension for Conducting Targeted ICS Attacks," in *Proceedings of the 2017 Workshop on Cyber-Physical Systems Security and Privacy*, CPS '17 (New York, NY, USA: Association for Computing Machinery, 2017), 57–67.

[15] Benjamin Green, Marina Krotofil, and Ali Abbasi, "On the Significance of Process Comprehension for Conducting Targeted ICS Attacks," in *Proceedings of the 2017 Workshop on Cyber-Physical Systems Security and Privacy*, CPS '17 (New York, NY, USA: Association for Computing Machinery, 2017), 57–67. This graphic has been edited for spelling.

[16] Zetter, "An Unprecedented Look at Stuxnet, the World's First Digital Weapon."

[17] Green, Krotofil, and Abbasi, "On the Significance of Process Comprehension for Conducting Targeted ICS Attacks," 2017.

[18] Dragos, Inc., "CRASHOVERRIDE: Analysis of the Threat to Electric Grid Operations," June 13, 2017, https://www.dragos.com/wp-content/uploads/CrashOverride-01.pdf?hsCtaTracking=ec040772-a407-4187-bd1d-e750cb8e1d99%7Ced3e20a3-1321-4fad-b91b-d19eb25b2350.

[19] Kevin E. Hemsley and Ronald E. Fisher, "History of Industrial Control System Cyber Incidents," 2018, https://doi.org/10.2172/1505628.

[20] Chris Bing, "Trisis Malware Has the Security World Spooked, Stumped and Searching for Answers," CyberScoop (CyberScoop, January 16, 2018), https://www.cyberscoop.com/trisis-ics-malware-saudi-arabia/.

[21] Blake Johnson et al., "Attackers Deploy New ICS Attack Framework 'TRITON' and Cause Operational Disruption to Critical Infrastructure," *Threat Research Blog* 14 (2017), https://www.fireeye.com/blog/threat-research/2017/12/attackers-deploy-new-ics-attack-framework-triton.html.

[22] Marina Krotofil, "Deep Dive: Likely, Real and Unlikely Cyber-Physical Threats to ICS" (September 27-29, 2017), https://www.youtube.com/watch?v=fn7eF9xKWYg.

[23] Green, Krotofil, and Abbasi, "On the Significance of Process Comprehension for Conducting Targeted ICS Attacks," 2017.

[24] Kim Zetter et al., "Inside the Cunning, Unprecedented Hack of Ukraine's Power Grid," *Wired* (WIRED, March 3, 2016), https://www.wired.com/2016/03/inside-cunning-unprecedented-hack-ukraines-power-grid/.: An approximation derived from the timeline between phishing emails being received and the date of the attack.

[25] Zetter et al.

[26] William J. Broad, John Markoff, and David E. Sanger, "Israeli Test on Worm Called Crucial in Iran Nuclear Delay," *The New York Times*, January 15, 2011, https://www.nytimes.com/2011/01/16/world/middleeast/16stuxnet.html.

[27] Robert Lee, Joe Slowik, Ben Miller, Anton Cherepanov, Robert Lipovsky, "Industroyer/Crashoverride: Zero Things Cool About a Threat Group Targeting the Power Grid" (July 26-27, 2017), https://www.youtube.com/watch?v=TH17hSH1PGQ.

[28] Robert Lee, Joe Slowik, Ben Miller, Anton Cherepanov, Robert Lipovsky.

[29] Yuqi Chen et al., "Learning-Guided Network Fuzzing for Testing Cyber-Physical System Defences," *ArXiv [Cs.CR]* (September 12, 2019), arXiv, http://arxiv.org/abs/1909.05410.

[30] CISA, U.S. Industrial Control Systems Cyber Emergency Response Team, "ICS Cybersecurity (301V)."

[31] Cheng Feng et al., "A Deep Learning-Based Framework for Conducting Stealthy Attacks in Industrial Control Systems," *ArXiv [Cs.CR]* (September 19, 2017), arXiv, http://arxiv.org/abs/1709.06397.

[32] Cheng Feng et al., "A Deep Learning-Based Framework for Conducting Stealthy Attacks in Industrial Control Systems," *ArXiv [Cs.CR]* (September 19, 2017), arXiv, http://arxiv.org/abs/1709.06397.

[33] Feng et al., "A Deep Learning-Based Framework for Conducting Stealthy Attacks in Industrial Control Systems," September 19, 2017. Researchers acknowledge this limitation.

[34] Bergman and Halbfinger, "Israel Hack of Iran Port Is Latest Salvo in Exchange of Cyberattacks."

[35] M. Elnour et al., "A Dual-Isolation-Forests-Based Attack Detection Framework for Industrial Control Systems," *IEEE Access* 8 (2020): 36639–51; Cheng Feng et al., "A Systematic Framework to Generate Invariants for Anomaly Detection in Industrial Control Systems," in *Proceedings 2019 Network and Distributed System Security Symposium* (Network and Distributed System Security Symposium, Reston, VA: Internet Society, 2019), https://doi.org/10.14722/ndss.2019.23265; Yoshiyuki Harada et al., "Log-Based Anomaly Detection of CPS Using a Statistical Method," *ArXiv [Cs.SE]* (January 12, 2017),

arXiv, http://arxiv.org/abs/1701.03249; J. Goh et al., "Anomaly Detection in Cyber Physical Systems Using Recurrent Neural Networks," in *2017 IEEE 18th International Symposium on High Assurance Systems Engineering (HASE)*, 2017, 140–45; J. Inoue et al., "Anomaly Detection for a Water Treatment System Using Unsupervised Machine Learning," in *2017 IEEE International Conference on Data Mining Workshops (ICDMW)*, 2017, 1058–65; Moshe Kravchik and Asaf Shabtai, "Detecting Cyber Attacks in Industrial Control Systems Using Convolutional Neural Networks," in *Proceedings of the 2018 Workshop on Cyber-Physical Systems Security and PrivaCy*, CPS-SPC '18 (New York, NY, USA: Association for Computing Machinery, 2018), 72–83; Moshe Kravchik and Asaf Shabtai, "Efficient Cyber Attacks Detection in Industrial Control Systems Using Lightweight Neural Networks and PCA," *ArXiv [Cs.CR]* (July 2, 2019), arXiv, http://arxiv.org/abs/1907.01216; Chuadhry Mujeeb Ahmed, Jianying Zhou, and Aditya P. Mathur, "Noise Matters: Using Sensor and Process Noise Fingerprint to Detect Stealthy Cyber Attacks and Authenticate Sensors in CPS," in *Proceedings of the 34th Annual Computer Security Applications Conference*, ACSAC '18 (New York, NY, USA: Association for Computing Machinery, 2018), 566–81; Wissam Aoudi, Mikel Iturbe, and Magnus Almgren, "Truth Will Out: Departure-Based Process-Level Detection of Stealthy Attacks on Control Systems," in *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, CCS '18 (New York, NY, USA: Association for Computing Machinery, 2018), 817–31.

[36] Alessandro Erba et al., "Real-Time Evasion Attacks with Physical Constraints on Deep Learning-Based Anomaly Detectors in Industrial Control Systems," *ArXiv [Cs.CR]* (July 17, 2019), arXiv, http://arxiv.org/abs/1907.07487.

[37] Ling Huang et al., "Adversarial Machine Learning," in *Proceedings of the 4th ACM Workshop on Security and Artificial Intelligence*, AISec '11 (New York, NY, USA: Association for Computing Machinery, 2011), 43–58.

[38] Alexander Stiehler and Sundeep Gantori, "Longer Term Investments: Automation and Robotics" (UBS Chief Investment Office GWM, February 26, 2020), https://www.ubs.com/content/dam/WealthManagementAmericas/documents/automation-and-robotics-lti-report.pdf.

[39] International Business Machines, "Maximo APM - Overview," International Business Machines, accessed July 10, 2020, https://www.ibm.com/products/ibm-maximo-asset-performance-management.

[40] Statista, "Industrial Internet of Things Market Size Worldwide from 2017 to 2025* (in Billion U.S. Dollars)," Chart (Statista, March 6, 2020), https://www-statista-com.proxy.library.georgetown.edu/statistics/611004/global-industrial-internet-of-things-market-size/. International Business Machines, "Maximo APM - Overview."

[41] International Business Machines, "Maximo APM - Overview."

[42] Erba et al., "Real-Time Evasion Attacks with Physical Constraints on Deep Learning-Based Anomaly Detectors in Industrial Control Systems."

[43] Erba et al.

[44] U.S. Nuclear Regulatory Commission, “NRC: Backgrounder on the Three Mile Island Accident,” US Nuclear Regulatory Commission, 2014, https://www.nrc.gov/reading-rm/doc-collections/fact-sheets/3mile-isle.html.

[45] Bernard Marr, “What Is Industry 4.0? Here’s A Super Easy Explanation For Anyone,” *Forbes Magazine*, September 2, 2018, https://www.forbes.com/sites/bernardmarr/2018/09/02/what-is-industry-4-0-heres-a-super-easy-explanation-for-anyone/.

[46] U.S. National Security Agency and U.S. Cybersecurity and Infrastructure Security Agency, “NSA and CISA Recommend Immediate Actions to Reduce Exposure Across All Operational Technologies and Control Systems,” July 2020, https://media.defense.gov/2020/Jul/23/2002462846/-1/-1/1/OT_ADVISORY-DUAL-OFFICIAL-20200722.PDF; Cybersecurity and Infrastructure Security Agency, “Securing Industrial Control Systems: A Unified Initiative” (Department of Homeland Security, July 7, 2020).

[47] Eric Cole, “Threat Hunting: Open Season on the Adversary” (SANS Institute Information Security Reading Room, April 2016), https://www.sans.org/reading-room/whitepapers/analyst/threat-hunting-open-season-adversary-36882.

[48] U.S. Department of Homeland Security. n.d. “NCCIC ICS.” National Cybersecurity and Communications Integration Center. https://us-cert.cisa.gov/sites/default/files/FactSheets/NCCIC%20ICS_FactSheet_NCCIC%20ICS_S508C.pdf.

[49] Glenn A. Fink and Yana Shulga, “Helping IT and OT Defenders Collaborate,” *2018 IEEE International Conference on Industrial Internet (ICII)*, 2018, https://doi.org/10.1109/icii.2018.00036.

[50] Joseph Slowik, “Evolution of ICS Attacks and the Prospects for Future Disruptive Events” (Dragos, Inc, February 25, 2019), https://www.dragos.com/wp-content/uploads/Evolution-of-ICS-Attacks-and-the-Prospects-for-Future-Disruptive-Events-Joseph-Slowik-1.pdf.

[51] Fink and Shulga, “Helping IT and OT Defenders Collaborate.”

[52] NSACyber, “GRASSMARLIN,” GitHub, June 27, 2020, https://github.com/nsacyber/GRASSMARLIN.

[53] U.S. National Counterintelligence and Security Center, “National Counterintelligence Strategy of the United States of America 2020-2022”; Cybersecurity and Infrastructure Security Agency, “Securing Industrial Control Systems: A Unified Initiative.”

[54] U.S. National Security Agency and U.S. Cybersecurity and Infrastructure Security Agency, "NSA and CISA Recommend Immediate Actions to Reduce Exposure Across All Operational Technologies and Control Systems."