When Robots Attack: Examining Artificial Intelligence, Autonomy, and Unmanned Threats

2018 OSAC Annual Briefing

Produced by the Research & Information Support Center

Executive Summary

From spam filters on our email to smart home devices, artificial intelligence (AI) has become a ubiquitous component of daily life that often improves the efficiency and effectiveness of routine tasks. In many cases, advances in the field have allowed machines to complete complex tasks once assigned to humans while removing the opportunity for human error. Despite the many promises of AI, there is a potential dark side to machine learning and increased autonomy. AI and increased machine autonomy continue to expand across areas such as robotics and weapons systems. As the technology proliferates and the barriers to use decrease, it will likely fall into the hands of nefarious actors and influence the threat landscape in a number of meaningful ways. This report will examine several future risks posed by these developments, specifically as they relate to the physical security[1] of U.S. private-sector organizations. Ultimately, this report and the subsequent case studies will address what security managers in the private sector should be doing now in order to get ahead of the AI-enabled autonomous threats of tomorrow.

Table of Contents

What is AI? 1
What is Autonomy? 2
Commercial and Military Applications of AI 2
Promise and Peril of AI and Machine Autonomy 3
Preparing for the AI Revolution 5

Driverless Dilemma: Security Implications of Self-Driving Cars on Physical Security 10
The Basics of Driverless Cars 10
Understanding Autonomy within the Automotive Industry 11
Driverless Vehicles and Terrorism 11
Impact of Driverless Cars on U.S Private Sector Physical Security 13

A Mind of Their Own: The Impact of Drone Autonomy on Physical Security 15
Examining the Drone Threat 15
Increasing Autonomy and the Coming of Drone Swarms 16
Impact of Drone Autonomy on U.S. Private Sector Physical Security 17

What is AI?

AI, or artificial intelligence, is the ability of machines to learn and perform tasks similarly to humans. This learning is accomplished through a variety of technologies that assist machines in identifying patterns within large datasets. The process of providing computers access to large amounts of data and allowing them to learn on their own is often referred to as "machine learning."

[1] Artificial intelligence will undoubtedly shift the threat landscape in other areas, such as cyber security. However, this analysis will focus solely on its impact to physical security.

The contents of this (U) report in no way represent the policies, views, or attitudes of the United States Department of State, or the United States Government, except as otherwise noted (e.g., travel advisories, public statements).
This report was compiled from various open sources and (U) embassy reporting.
Please note that all OSAC products are for internal U.S. private-sector security purposes only. Publishing or otherwise distributing OSAC-derived information in a manner inconsistent with this policy may result in the discontinuation of OSAC support.

It was this concept of allowing computers to teach themselves, along with the arrival of the internet and the large increase in available data that have propelled AI to its current prominence. A subset of machine learning called deep learning powers the most human-like artificial intelligence available today. Rather than teaching computers everything they need to know in order to complete a task, it is now possible to code them to "think" like humans via artificial neural networks, and simply provide them with the data necessary to do their jobs.

AI can be broken down into two categories: general (broad) and applied (narrow). General AI is applied to systems or machines that are designed to be adaptable and carry out a variety of different tasks. This is the type of AI that most commonly shows up in movies like The Terminator; while it does not currently exist in a Terminator-like capacity, general AI is where a significant portion of the current cutting-edge research is concentrated. Applied AI, on the other hand, exists in abundance all around us, enabling computers to learn to carry out very specific tasks without being given instructions. These narrowly focused applications of AI can be widely observed, from automated social media feeds to targeted online advertising and bank fraud detection.

What is autonomy?

Many prominent figures, including Stephen Hawking and Bill Gates, have spoken about AI security concerns in recent years, with Hawking going as far as to say, "The development of full artificial intelligence could spell the end of the human race." The key to these objections often hinges on the level of autonomy that computers are given, a factor determined by the machine's design. In the context of AI, autonomy refers to the ability of a machine to perform a function on its own. Many opponents of increased autonomy within AI contend that limiting autonomy through some level of human control is essential to ensuring safety, especially with machines like vehicles or weapons where lives are at stake. Opposition to fully autonomous lethal weapons has also led to the formulation of vocal opposition groups, such as the Campaign to Stop Killer Robots.

Autonomy has three dimensions: the type of task the machine is performing; the relationship of the human and machine in performing the task; and the sophistication of the machine's decision making abilities when performing the task. These dimensions are independent from one another, and a machine becomes more "autonomous" if autonomy is increased along any of the three. As machines become more complex (in tasks and decision making) and require less human involvement, they become harder to predict. This may lead to unintended consequences.

Commercial and Military Applications of AI

In recent years, AI and machine learning have made rapid advances, enabling humans to achieve breakthroughs across a number of sectors from farming to aviation. In the healthcare field, AI is being used to augment the work of radiologists in diagnosing early signs of cancer or strokes. Meanwhile, the automotive industry is using AI to make vehicles safer and more reliable for consumers. In these cases, AI is speeding up processing of information and removing human error in diagnosing discrepancies or abnormalities. These qualities will also make AI a useful tool for many tasks within the security industry currently completed by humans.

In addition to its many commercial applications, AI also has military applications that are changing the battlefield and warfare landscape in significant ways. A 2017 Department of Defense study on the future operating environment described AI as "the most disruptive technology of our time." One key shift will likely be the ability of AI to increase the pace of conflicts as intelligence is processed faster and computers learn to make battlefield recommendations faster than their human counterparts. AI will also revolutionize military training by enhancing simulations and improving the cognitive skills of top military leaders. Of the many military applications of AI, none is more controversial that the development of autonomous, warfighting robots and weapons systems capable of identifying targets and making life or death decisions. The proliferation of these weapons could present a significant security threat to the general public in the hands of terrorists or other nefarious actors. This possibility will be explored further in a subsequent section.

World leaders have taken note of these advances in AI. Indian Prime Minister Narendra Modi has opined that artificial intelligence will one day "drive the human race." Likewise, Russian President Vladimir Putin has predicted that whoever establishes themselves as a leader in AI will "become the ruler of the world." These remarks reflect the views of many experts working to develop and understand this technology. Many analysts believe that AI represents the next industrial revolution, and will significantly disrupt the global economy. Others have described AI as being as transformative as electricity. If past industrial revolutions are any indication, this shift could also cause significant disruptions in the global balance of power, international competition, and international conflict.

Promise and Peril of AI and Machine Autonomy

AI and machine autonomy will impact private-sector physical security in two meaningful ways. First, it will provide promising opportunities for computers and robots to make organizations more secure. By outsourcing tasks like video monitoring and perimeter security to AI-enabled machines, organizations will be able to increase threat recognition and reduce the burden on human personnel, who are potentially more prone to error. With this opportunity comes increasing concern regarding how to secure the technology against hacking and other cyber-attacks.

Second, AI and machine autonomy will present new opportunities for terrorists and other nefarious actors to conduct scalable, coordinated attacks remotely, involving less personnel and decreasing the risk to attackers. In the past, threat actors have shown both the capabilities and intent to adapt widely available consumer technologies to conduct attacks. It is very likely that they will also incorporate artificial intelligence into their arsenals as the technology proliferates.

AI Enhanced Physical Security

Despite the many concerns surrounding lethal applications of autonomy, AI has a number of positive physical security applications. For example, though closed-circuit television (CCTV) systems are common, many are passive and serve as a deterrent rather than an active security measure. Those that are regularly monitored by security personnel present a situation where threats can go unnoticed due to limited capacity of operators to monitor multiple live feeds. To overcome this hurdle, AI is providing "brains" to the "eyes" of CCTV, enabling analysis of live video footage by computers without human involvement.

One example of this is the Japanese AI-augmented security camera called "AI Guardman" that helps store owners identify shoplifters by matching body poses to predefined suspicious behavior. When combined with other technology, like facial recognition, AI-enabled security cameras are able to provide real-time situational awareness to security personnel. Despite its many promises, views on these capabilities are mixed (especially in cases of government use) due to privacy concerns. Also, the technology is not infallible.

Robots fueled by AI are also expected to augment security personnel in conducting perimeter security operations. Through increased AI-enabled automation, unmanned ground security robots have the capability to patrol perimeters, detect security breaches or other potential risks (e.g. damaged fences), and alert appropriate security personnel. Aerial drones are also being used to carry out autonomous patrol routes without human intervention. In addition to increasing an organization's ability to detect intrusive and risky activity, these machines can also take on tedious, remote, or high-risk tasks, leaving the more strategic security duties to human personnel.

AI Enabled Attacks

Malicious actors have taken note of unmanned technology, with some seeking to exploit advances in machine learning and increased autonomy to conduct physical attacks. As has been noted in previous OSAC analysis (see: Drone Operations and Threats Abroad) these actors have been successful in employing unmanned machines, like commercially available drones, to conduct physical attacks on targets. Many experts believe it is only a matter of time before terrorists and other nefarious actors will leverage the capabilities of AI in order to conduct attacks. Some go as far as to argue that rogue states and terrorists will be the biggest winners of the coming AI arms race. In an open letter to the United Nations, 114 leading AI and robotics industry experts, including Tesla's Elon Musk, described the frightening possibility that autonomous machines could become "weapons that despots and terrorists use against innocent populations."

The unmanned threat posed by drones and other robotic machines currently exists independently of AI, and generally relies on remote piloting by a human operator. However, one can imagine how the threat would evolve with increased autonomy. What if it were possible to remove the human operator and have machines act independently based on collected environmental data? What if it were possible for one operator to direct a scalable attack involving multiple autonomous machines? What if those machines had the power to make life or death decisions independently?

Though the prospect of killer robots may seem like dystopian delusion or futuristic farce, the technological capabilities exist, and the intent of nefarious actors to employ new, cutting-edge methods has not diminished. Seemingly harmless technology like driverless cars, which promise to revolutionize global transportation, has already been coopted by terrorist groups like ISIS for the purpose of conducting unmanned vehicle-borne improvised explosive device (VBIED) attacks. And as the recent drone attack against Venezuelan President Nicolas Maduro illustrates, nefarious actors worldwide have discovered the operational benefits of current robotics and unmanned technology, and continue to leverage that which is commercially available. The impact

of autonomy on both drone and vehicle-borne threats will be addressed as case studies in an addendum to this report.

Looking back to the military use of AI and autonomous weapons, more than 30 countries have defensive weapons meant to respond in situations that require a quicker responses than humans are able to provide. Several countries also employ military-grade drones controlled by humans. However, Israel has crossed into the realm of full autonomy with its Harpy drone, which can seek out enemy radar and conduct a "suicide" attack – self-detonating at a target -- without human instruction. Though these weapons seem like they belong in futuristic science fiction films, the technology behind them is very real, and is rapidly advancing and proliferating around the world.

As AI capabilities develop and the autonomy of these machines increases, nefarious actors will likely realize the benefits of using smarter technologies. For the most part, the unmanned machines that pose a threat today require remote operation by a human attacker that is within operating range of the device's command and control signals. However, AI control would allow these machines a greater degree of autonomy, enabling larger-scale attacks that require less human input. Though the software required to bring these attacks to fruition is fairly sophisticated, it can easily be found in open source. For example, the algorithms for facial detection, navigation and planning, and multi-agent swarming can easily be found online and employed in an attack.

Beyond decreasing the manpower and proximity barriers to an attack, AI-enabled unmanned vehicles also increase the number of spaces that are vulnerable to automated physical attacks. Unmanned vehicles tend to be smaller, lighter, faster, and more maneuverable than traditional human-inhabited vehicles, which provides strategic advantages for both soldiers and terrorists. They are not inhibited by the same endurance limitations as humans, and often offer enhanced perception or increased physical capacities. Finally, and most notably, they also offer the ability to conduct dangerous missions without risking human lives.

Preparing for the AI Revolution

Thus far it has been established that: 1) AI is drastically enhancing the ability of existing unmanned machines to act autonomously; 2) autonomous weapons are actively being developed by governments around the world for military use; and 3) terrorists and other nefarious actors will likely both seek out these weapons and attempt to leverage AI capabilities to enhance the lethality and operability of other current attack platforms, like over-the-counter drones. In that case, what should security managers be doing now in order to get ahead of the threats of tomorrow?

There are three key steps to this process:

1. Identify the full scope of the threat and risk.
2. Develop and employ near-term mitigation measures.
3. Begin evaluating long-term measures that account for AI and increased autonomy.

Identify the full scope of the threat and risk

To fully recognize the physical security threat posed by AI and unmanned, autonomous machines, security managers should conduct risk and vulnerability assessments that inform mitigation measures. *Threat* often refers to the capability of a terrorist or other nefarious actor to conduct an attack against an organization, whereas *risk* generally involves both the probability that the organization may be involved in an attack (either deliberately or by chance) and the harm that such involvement would cause. Simply put, *threat = capability x intent*, and *risk = probability x harm*. With that in mind, risk assessments for unmanned, autonomous threats should address the following questions:

- Who are the individuals or organizations (threat actors) that may seek to damage your interests (personnel, facilities, reputation, etc.)?
- Have threat actors exhibited the capability to conduct unmanned attacks or attacks that require similar sophistication (technological expertise, finances, manpower, etc.)?
- Have the threat actors demonstrated specific intent to attack your organization or other similar entities? If so, have they demonstrated interest/expertise in unmanned attacks? Have they previously employed attack methods that seek to limit loss of personnel?
- What is the likelihood that the threat actor could muster the resources necessary to make an unmanned attack successful? How capable are local security personnel in disrupting and disarming threat actors before attacks take place?
- Based on the known capabilities and intentions of the threat actor, what would the impact to the organization be in the event of a successful unmanned attack?
- Finally, how would the risk profile shift if the threat actors were able to conduct unmanned attacks via multiple machines that act independently from human operators?

Take the example of ISIS, a group that has previously exhibited both substantial capabilities and intent to carry out unmanned attacks against Western targets in Iraq and Syria. Both intent and probability are likely impacted by preference toward military targets. However, the capabilities and potential for significant harm are certainly there. Also, through both success and effective propaganda, ISIS attacks have inspired lone wolf attackers to leverage unmanned platforms when planning attacks, a factor which extends the threat well beyond the organization's physical control.

Other groups have also proven capabilities and intent to use unmanned attack platforms, including Hezbollah and Hamas. Additionally, as the number of countries employing armed unmanned systems in combat increased, the chances for the technology falling into the wrong hands also increases. Rogue states (e.g. Iran) present a particularly worrying scenario in regard to unmanned weapons systems, like kamikaze drones or boats, as they are often more willing to provide them to non-state actors (e.g. Houthi rebels) whose actions benefit their strategic aims. This presents an opportunity for unpredictable and potentially nefarious use thereafter.

Develop and employ near-term mitigation measures

As has been established previously, the technology required for terrorists and other nefarious actors to conduct unmanned attacks already exists. Though they often require remote operation, attacks using unmanned machines have been planned and successfully enacted in multiple regions of the world. They also transcend multiple domains, as attacks are being orchestrated from the air, land, and sea. Rather than waiting for the threat to evolve with the coming wave of AI and autonomy, security managers should be developing protocols and employing mitigation measures now to meet the existing threat head on.

Establishing key security vulnerabilities is critical to this process. Because of their previously described tactical advantages, many unmanned systems might be difficult to detect or, if encountered, might not arouse suspicion. Any evaluation of how AI and unmanned machines might impact an organization's security profile in the near term must also be accompanied by a vulnerability assessment that takes into account their tactical advantages. Consider the organization's key security vulnerabilities absent this threat, and then evaluate how those vulnerabilities might be further exploited via unmanned systems (regardless of the operator).

Current unmanned physical security threats generally involve the remote deployment of a vehicle to deliver a malicious payload to a nearby target. The delivery vehicles range from full-size automobiles to small drones, both of which present detection issues. That said, many common risk mitigation measures used to protect against vehicular attacks, explosives, surveillance, etc. can be equally effective against various unmanned threats. For example, effective perimeter security design that includes adequate standoff distance and physical barriers can also be effective at protecting against a VBIED via a driverless car. Similarly, having a shelter-in-place plan or simple surveillance prevention procedures like closing blinds could be sufficient enough to protect personnel and information in the case of an unauthorized unmanned aerial vehicle (UAV) on the premises.

It is not prudent simply to assume that current measures will be broadly applicable and adequate to meet current and future security needs. Security managers should integrate the possibility of an unmanned attack into security planning and protocols, to include conducting drills that involve unmanned threats. Any new or revised protocols and security measures should also take into account the key differences between unmanned and human threats. For example, it is much easier to decipher the capabilities and intentions of a suspicious human actor than it is to do so for a machine, like a UAV. Unmanned vehicles also have a novelty factor that often influences security personnel or bystanders to move toward the device rather than away. Terrorists have recognized this and preyed on this instinct to conducts attacks in the past.

Security managers should also consider both passive and active countermeasures for detecting and interdicting unmanned threats. In the past several years, there have been over 200 products developed in the counter-UAV market alone that aim to detect and disable these devices. As the threat exists currently, the primary solution to stopping unmanned vehicles is to jam the fragile links the device depends on to communicate with its human operator. These same (often radio frequency-based) communication links are also used to identify possible threats alongside

methods like identifying visual, acoustic, and heat signatures. Security managers with operations where unmanned threats present a significant risk should be exploring all available countermeasures. It is worth noting though that many active countermeasure options are still largely untested and can be limited by effectiveness and, in some cases, legality.

Begin evaluating long-term measures that account for AI and increased autonomy

To date, much of the conversation regarding avoiding AI misuse and subsequent threats is very broad and high-level. For example, the often cited Malicious Use of Artificial Intelligence Report makes a number of key recommendations regarding how policymakers, researchers, engineers, etc. can work together to develop AI that is safer and much less susceptible to malicious use. However, broad assessments generally do not provide specific recommendations regarding physical protection against these technologies as they evolve.

Security managers should identify the areas in which AI and autonomy are most likely to impact their security operations. This includes considering physical security, which has been the primary focus of this report, as well as other security areas, such as cyber or political. For example, in the cyber realm, AI can be used to help hackers find identify data to breach and discover vulnerabilities within a system. In the political arena, AI can be leveraged to create personalized disinformation campaigns and even so-called "deepfakes," where a computer can doctor images or video so well that the manipulated files are indistinguishable from the originals. It easy to imagine how using this technology against governments could potentially result in unrest.

As for physical security, consider the ways in which machine autonomy will alter the previously discussed unmanned threats. For example, how does the threat posed by drones (land, air, or sea) change if there is no remote operator nearby and the device has the ability to recognize faces and identify targets? Maybe plans and protocols are in place for dealing with a singular unmanned threat, but how would security personnel respond to a swarm of unmanned machines? Also, as was previously noted, one current solution to stopping unmanned threats is to jam the communication links and eliminate contact with the human operator. However, as the level of autonomy increases, these communication links are no longer necessary and this solution becomes less viable. How will security personnel respond to these threats as they become more difficult to disable electronically?

After considering the ways that AI could be disruptive or threatening to your organization's operations, consider the many positive applications that could revolutionize your security operations. As was noted previously (and will be further evaluated in the case studies) developments in AI, autonomy, and robotics will also present a number of exciting opportunities to boost security via technology. These should not be overshadowed due to concerns regarding threats posed by nefarious uses of the technology. In some cases, AI may actually be the solution to detecting and/or deterring AI-based autonomous threats.

Conclusion

Preparing for the coming AI revolution is going to require security managers to be imaginative in considering the number of ways that artificial intelligence and autonomy might impact security

operations abroad. The components required to automate and weaponize a number of machines already exist, and terrorists and other nefarious actors have already taken note. If history is any indication, threat actors will innovate and implement AI and autonomous technologies as they become more advanced and more widely available. AI and autonomous machines are not simply a security threat of tomorrow; their predecessors are here today, and the security community must adapt quickly in order to address their potential impact to security operations abroad.

Additional Resources

Army of None – Paul Scharre
The Malicious Use of Artificial Intelligence: Forecasting, Prevention, and Mitigation
Artificial Intelligence and National Security – The Belfer Center
Artificial Intelligence and International Security – Center for New American Security (CNAS)
Strategic Competition in an Era of Artificial Intelligence – CNAS
Between a Roomba and a Terminator: What is Autonomy? – War on the Rocks
Department of Defense Unmanned Systems Integrated Roadmap 2017-2042
An AI Glossary – The New York Times

Driverless Dilemma: Security Implications of Self-Driving Cars on Physical Security

Overview

Driverless cars represent a transformative technology that promises to revolutionize the transportation industry by making roads safer and more efficient. The industry has made significant advances in recent years, with some companies expecting success with high levels of automation as soon as 2020. Given the preference of many terrorist groups to conduct vehicle-borne improvised explosive device (VBIED) and vehicle-ramming attacks, it's easy to imagine how the growth of driverless cars might also revolutionize global terrorism. We are certainly years away from fully autonomous vehicles being broadly employed on public roads. However, the potential to weaponize autonomous driving technology already exists, and threat actors have taken note. This report will focus on artificial intelligence (AI) and autonomy within the automotive sector, specifically focusing on how this technology will affect the ability of terrorists and other nefarious actors to conduct vehicle-based attacks that may impact the U.S. private sector.

The Basics of Driverless Cars

Driverless cars combine a complex set of systems and software that allow the machine to control and navigate itself. They are fitted with cameras, sensors, and communication systems to enable the vehicle to generate large amounts of data that, when combined with the capabilities of artificial intelligence (AI), enables the vehicle to make decisions similarly to human drivers. Though this technology is still in the early stages, it is becoming more common. While there are no legally operating, fully autonomous vehicles on the road today, the increasing impacts of autonomy are apparent through brake and lane assistance technology and the number of successful self-driving prototypes that are being produced.

So how does a vehicle learn to drive itself? Though each model is different, most self-driving vehicles create and maintain an internal map of their surroundings, based on data collected by sensory equipment. Software then processes that incoming data, plots a path, and sends instructions to the vehicle's actuators, which control acceleration, braking, and steering. The computer systems that operate this software consist of three components. The first is **perception**, which takes information from the previously mentioned sensors and identifies relevant objects nearby. The data from these sensors is combined to build a model of the environment; AI systems then identify nearby cars, bicycles, pedestrians, etc. The second component is **prediction**, which forecasts how each of those objects will behave over the next few seconds. Finally, the third component establishes the driving **policy** by using the predictions to determine vehicle response.

Though current models of driverless cars are highly successful at identifying familiar objects and following rules, their ability to generalize and avoid accident-prone situations is still under development. Unpredictable situations present complications for a narrowly-focused, applied AI system that thrives on predictability. This has resulted in a few significant driverless car incidents, including the death of a pedestrian following a collision with a self-driving vehicle in Tempe, Arizona. For this reason, some semi-autonomous products (e.g. Tesla's Autopilot) require human intervention in unpredictable situations.

Understanding Autonomy within the Automotive Industry

As is the case more broadly with AI, much of the current debate regarding the risks posed by driverless cars focuses on the issue autonomy. The Society of Automotive Engineers, a professional association that sets industry guidelines, has developed a level-by-level automation guide (see graphic below) that has become the industry standard. The most significant pivot occurs between Levels 2 and 3, when responsibility for monitoring the driving environment shifts from the driver to the system. The goal for the driverless car industry is full autonomy (Level 5). However, getting there will require several semi-autonomous phases that leverage AI and other technologies to automate key driving tasks and move toward vehicles that mimic human drivers.

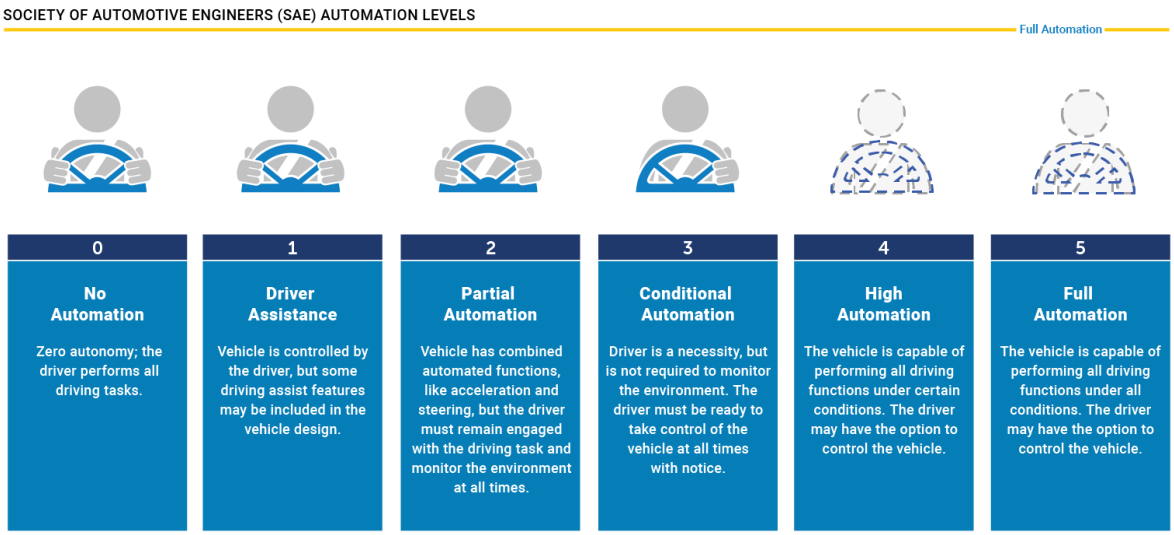

Source: Society of Automotive Engineers (SAE)

Driverless Vehicles and Terrorism

Driverless vehicles may present a number of complications for the security industry in the coming years. At the forefront of these concerns will likely be if/how organizations should adopt the technology for their own transportation purposes. In this case, understanding the technology behind driverless car incidents will be incredibly important. Though this will be a serious concern in the future, we are many miles (and years) away from successful and safe deployment of fully autonomous vehicles on public roads. If that is the case, what risk is posed by this technology today, and how is that risk likely to shift as autonomy increases in the future?

> "Autonomy will improve driving safety and make mobility more efficient, but will also open up greater possibilities for dual use applications and ways for a car to be more of a potential lethal weapon than it is today."
>
> -Federal Bureau of Investigation (FBI)

Though full autonomy is years away, most of the technological components needed to provide automation to certain driving functions already exist. Full autonomy is also not required in order to make a vehicle lethal. The mere ability of terrorists and other nefarious actors to control a vehicle remotely presents a concerning scenario in which VBIED and ramming attacks do not require an operator's direct presence. Remote operation also presents significant complications for law enforcement personnel trying to identify the perpetrators of attacks. Evidence suggests that ISIS has been successful in developing a remotely-operated, driverless vehicle. Intercepted videos show ISIS instructors training fighters in Syria on how to construct remotely-controlled vehicles using sophisticated radios and electronic components from scrap cars. Additionally, an ISIS-inspired plot to conduct a driverless car attack in the United Kingdom was thwarted in December 2017.

This shift is consistent with previous evolution of terrorist tactics. Stationary roadside bombs have been a key component for guerilla and terrorist groups for decades. The Irish Republican Army began employing them in Northern Ireland in the 1970s, and they also became a key tool in the Lebanese civil war. However, this static approach limited the ability of groups to conduct attacks directly on a target in real time. In 1981 a suicide bomber conducted a VBIED attack on the Iraqi Embassy in Beirut, overcoming the previous limitation.

Over the past few decades, VBIED suicide attacks have become the most lethal form of terrorist attack, killing 4.5 times more people than any other form of terrorist attack between 2001 and 2013. Outside of the Middle East, other vehicle-borne attacks, such as vehicle rammings, have also increased rapidly in recent years. In 2014, ISIS spokesman Abu Mohammed al-Adnani encouraged followers to run over Westerners with their cars, and several extremist publications have provided guidance on the subject. These attacks, like the 2016 Bastille Day truck ramming in Nice, also tend to be suicide-based attacks, where the drivers aim to achieve as many casualties as possible before being killed by law enforcement.

In shifting from stationary roadside bombs to suicide bombers, terrorists simply added a guidance system that leveraged the power of human intelligence to conduct an attack in real time. In shifting from vehicle bombings to vehicle rammings, terrorists were able to remove barriers to obtaining and assembling explosives, a key targeting point for law enforcement personnel. Driverless cars would remove a key recruiting barrier by no longer requiring terrorists to recruit young fighters willing to commit suicide missions for their cause.

Remotely controlling a vehicle presents numerous challenges. The threat actor must be able to transmit a command signal, control the vehicle based on very limited visual feedback, and overcome latency issues between visual feedback and the ability to direct operational change. However, these challenges can be overcome through a combination of both machine-based automation and human control, similar to that which we have already seen in the unmanned aerial

vehicle (UAV) industry. Many UAVs currently on the market have the ability to adjust to feedback from onboard sensors autonomously despite commands from a human operator. Terrorists may seek to overcome current barriers using a similar autonomous-remote hybrid approach and relying on existing technologies and existing vehicles.

Impact of Driverless Cars on U.S. Private Sector Physical Security

Vehicle ramming and VBIED attacks existed absent of vehicle automation, and will continue to exist moving forward. Increased automation and full autonomy will simply serve to reduce some of the barriers to conducting a successful attack, including real-time, targeted delivery and driver recruitment. As we move along the spectrum toward full automation, and the sensors and software necessary to automate driving become more widely available, technological barriers will decrease, increasing the likelihood that threat actors will leverage driverless cars as an attack platform. Decreased barriers could lead to decreased interdiction points for law enforcement personnel, which may increase the likelihood of successful attacks.

Most countermeasures designed to harden facilities against manned vehicular attacks will equally effective against an unmanned attack. In areas at risk for vehicle-based attacks, consider installing security barriers, such as bollards, designed to prevent vehicles from penetrating a secure perimeter. VBIED-prone locations should also include setback or stand-off distance in facility design. There should be as great a distance as possible between the facility and the likely positioning of a VBIED to mitigate the effects of a possible explosion. Given the volume of explosives used in some past VBIED attacks, achieving optimum setback is rarely possible. However, combining this countermeasure with other security-focused design elements can significantly decrease the impact of attacks. For example, positioning security guards at a vehicle control point can also be instrumental in detecting and deterring attacks.

That said, driverless cars do present unique security challenges. Detecting that a vehicle does not have a driver and then disabling that vehicle are not easy tasks. For remote-controlled and semiautonomous vehicles, jamming the fragile command-and-control communication links that connect the machine to the operator may be a viable solution. However, as automation improves, these communication links may not be necessary, and this solution becomes less viable. Though higher-caliber rifles and other weapons systems can disable a vehicle's engine, they often cannot be accessed quickly and cannot be used in many operating environments. Security managers will need to look both to low-tech solutions, like barriers, and high-tech vehicle disabling solutions, like electronic jamming via high-powered microwaves, in order to secure against driverless vehicle threats. Much of the technological innovation in this area is happening in the military.

Threat actors will also continue look for new, innovative ways to evade security countermeasures designed to prevent against driverless car attacks. For example, some secure facilities use scanners that detect heat signatures from humans inside of a vehicle. In the case of a driverless car, no human-like heat signature would be detected, which would alert security personnel to a potential threat. However, ISIS has sought to overcome this limitation by outfitting their remote-controlled vehicles with mannequins equipped with self-regulating thermostats to produce a heat

signature similar to humans. This illustrates the need for security managers to maintain situational awareness of up-to-date tactics and countermeasure avoidance efforts.

Conclusion

Widespread use of driverless vehicles is on the horizon, and the threat of attacks via unmanned vehicles already exists. Security managers operating in, or responsible for, areas at risk for vehicle-borne attacks should consider conducting risk and vulnerability assessments that account for unmanned technology. As previously mentioned, driverless technology will reduce the barriers to conducting a successful attack, thus raising the overall risk for many organizations. Future risk and vulnerability assessments should account for this increased risk, while considering how increasing automation will change the operating environment and overall risk profile. Finally, security managers should also stay abreast of new autonomy and AI developments in the automotive industry. The shift to widespread employment of driverless cars on public roads will undoubtedly be gradual. As this occurs, security managers should maintain awareness of new developments and any future lethal driverless vehicle incidents in order to develop a risk profile that can inform a cost-benefit analysis for future business use.

Additional Resources

Hostile Vehicle Ramming: U.S Department of State Resources
FBI Analytic Report: Autonomous Cars
Environment Net Assessment: Autonomous Vehicles – Homeland Security
The Wired Guide to Self-Driving Cars

A Mind of Their Own: The Impact of Drone Autonomy on Physical Security

Overview

In recent years, unmanned vehicles (UVs), more commonly referred to as drones, have received worldwide attention both as a high-impact military technology, as well as immensely popular toys for hobbyists. Advances in robotics and drone technology have also elevated commercially-available UVs to prominence as powerful business tools. At the same time, terrorists and other nefarious actors have also realized the benefits of unmanned technology and have deployed drones to conduct attacks. With the emergence of artificial intelligence (AI) and machine learning, drones will likely soon incorporate more autonomy and require less human guidance. As machines become more programmable, it may become possible for one individual to unleash multiple machines that carry out a pre-programmed mission on their own. AI may even offer these machines the ability to think like humans and make decisions, including ones with lethal results. This is not science fiction, but rather a technological reality that security managers are likely to face in the coming years. This report will focus on the near and long-term impacts that AI and autonomy will have on drone threats to the U.S. private sector overseas.

Examining the Drone Threat

> "As the Venezuela incident demonstrates, hobbyist, commercial UAVs are readily available and demonstrably capable of conducting limited attacks, not just on the battlefields of Iraq and Syria, but in a public assassination attempts."
>
> -Kathryn Dura, *National Interest* (2018)

On August 4, two commercially-available aerial drones, each equipped with a kilogram of powerful plastic explosives, were used in an attempted assassination against Venezuelan President Nicolas Maduro. Though the machines were unsuccessful in reaching their target, the attack raised the alarm and brought drone threats to the forefront of security conversations. Drone-based attacks from the air, land, and sea are not a new phenomenon. Terrorist and other nefarious actors have shown both the capability and intent to adjust attack platforms to use modern technology and increase the overall impact or lethality of attacks. Recent technological advances and proliferation of commercially-available drone technology have simply made UVs a high-tech tool in the arsenal for these threat actors.

Drone technology can be categorized in a number of ways, but two key distinctions are most common. First, drones can be classified by their availability. The key division is between commercially-available, over-the-counter products and military-grade products. There are also different types of drone platforms suited for different operating environments. While the term "drone" is most often used to refer to unmanned aerial vehicles (UAVs), there are also drones designed to operate on land, the sea surface, and underwater, known as unmanned ground vehicles (UGVs), unmanned surface vehicles (USVs), and unmanned underwater vehicles (UUVs) respectively.

Most of the innovation and focus in the drone industry surrounds UAV development. Likewise, UAVs are the means by which we see the most examples of nefarious use, both as reconnaissance and attack platforms. UAV innovation has been fueled in large part by the smartphone revolution. Many smartphone components – gyroscopes, accelerometers, GPS,

processors, and cameras – have been integrated into the onboard technology for modern, commercially-available UAVs. Key technological advances in the UAV industry have included changes to the sizes of aerial drones, their payload capacity, and their flight longevity. However, the most promising (and concerning) advancements in UAV technology, and the drone industry as a whole, are in the area of autonomy.

As over-the-counter drone technology has evolved, non-state actors have sought to leverage the technology to carry out surveillance and attacks. Several groups, including ISIS, have been particularly successful in this regard. In January 2017, ISIS announced the establishment of a drone unit known as "Unmanned Aircraft of the Mujahideen," which organizes the group's unmanned aircraft campaigns on the battlefield. According to uncovered ISIS documents, the program is not a series of one-off incidents, but rather an institutionalized, bureaucratic unit that has been planning for drone weaponization since 2015. ISIS has also released propaganda footage that illustrates the group's ability to drop munitions onto crowds and to hit stationary vehicles and tanks.

Military applications for drones have also dramatically increased in recent years. In the 1990s, the U.S. military spent around $300 million per year on drones. Annual military spending on drones has since increased to over 23 times pre-9/11 levels. This includes large drones like the MQ-9 Reaper down to hand-launched models like RQ-11 Raven. For Fiscal Year (FY) 2019, the Department of Defense has requested approximately $9.39 billion for unmanned systems and associated technologies. The proposal includes funding for the procurement of 3,447 new air, ground, and sea drones, a significant increase over FY 2018 numbers. Though the personal and commercial market for drones primarily focused on UAVs, military usage continues to rise across land and sea platforms as well.

These military developments are significant because non-state actors have also proven an ability to acquire and utilize multiple types of military drones in operations against both military and civilian targets. For example, Houthi rebels in Yemen unveiled their drone capabilities in January 2017 not in the air, but at sea. Using an armed, unmanned maritime craft, the group was able to strike a Saudi warship in the Red Sea, killing two sailors and injuring three others. Around the same time, Houthi rebels also displayed their UAV capabilities. Both technologies are believed to have been provided by Iran. Similarly, Hamas and Hezbollah have also employed military-grade, Iranian-made UAVs.

Increasing Autonomy and the Coming of Drone Swarms

At present, most consumer and military-grade drones are controlled by a human operator. However, new developments in UVs are allowing these machines semi-autonomy through components that automate various controls and allow the machines to interact with their surroundings. Some over-the-counter UAV models now have the ability to avoid obstacles and follow pre-determined paths without direct human control during the flight. As this trend progresses and AI is applied, drones could be given the ability to "think" for themselves and carry out assignments free of human control.

With autonomy comes mounting concerns over drone swarms (the use of multiple, autonomous drones acting in unison to overwhelm a target). To some extent, the ability to do this already exists, as drones swarms can be choreographed in advance or controlled by multiple humans. These capabilities have been demonstrated for both positive and nefarious purposes. Pre-programmed drone swarms have also been employed on the battlefield, as evidenced by an attack on Russian bases in Syria earlier this year. However, AI and autonomy present the opportunity for multiple drones with human-like decision-making ability to be dispatched to carry out attacks absent of human control. When combined with other technology (like explosives or facial recognition software), malicious drones could become a much more significant threat.

Fear regarding drone autonomy is the basis for *Slaughterbots*, a short video from the Future of Life Institute depicting a dystopian future in which militaries build autonomous, explosive micro-drones that fall into the hands of terrorists. Though this video is a sensationalized dramatization that may overlook several key assumptions, the central technical concept is grounded in reality. The technological components depicted in the video are quite real. Furthermore, there is not much that can be done to keep the underlying technology depicted in the video out of the hands of terrorists. The technology required to turn over-the-counter drones into weapons is fairly widespread and already being employed across the globe. As autonomy works its way into the drone market, threat actors will undoubtedly seek to use that technology to their advantage.

Impact of Drone Autonomy on U.S. Private Sector Physical Security

If nefarious actors are certain to obtain and use autonomous drone technology as it becomes available, what can be done now to prepare to counter this threat? For security managers in the private sector, the most effective approach to prepare for offensive UV use would be to focus on defensive measures. The threat from UVs must be integrated into an overall security management strategy. This includes conducting risk and vulnerability assessments, as well as integrating countermeasures into security protocols and policies.

Security managers should consider setting specific rules and guidelines regarding how employees should respond when a encountering a drone. Consider how the threat should be reported, what emergency response measures could be implemented, and how security personnel should respond. It is also crucial to conduct drills involving responses to a drone threat – much like fire drills and active shooter drills – to ensure that personnel know how to respond in a crisis. As was noted previously, drones have a novelty factor that complicates security response. Further suggestions specific to UAV incident response can be found in OSAC's reporting on Drone Operations and Threats Abroad.

In the near term, security managers should consider evaluating current technologies that can detect and disable drones. These technologies have proliferated widely in recent years, with most of the innovation focusing on counter-UAV technologies employable for both commercial and military use. Most drone countermeasures are based on jamming radio frequency and GPS signals that drones use to communicate with their human controller or determine their flight path. The assumption here is that drones need those signals in order to operate effectively. However, as drones become more autonomous, these countermeasures may become less effective.

Legal regulations and compliance regarding drone countermeasure technology are also critical. For example, many offensive countermeasures designed to disable UAV threats are not legal (domestically and in many foreign countries) due to anti-jamming regulations or those meant to protect civil aircraft. The irony is that autonomous machines with no centralized command and control systems present a unique type of threat where brute-force, offensive countermeasures may be the most effective approach. Moving forward, the legal environment for drone countermeasures will be every bit as important to follow as technological developments.

Though AI may create new, frightening opportunities for threat actors, it could also serve as a drone security solution. One key issue with the many drone detection and neutralization platforms available is that they require a high level of technical competence on the part of the operator. Such expertise is hard to find outside of the military. Training the system to "think" like a human could allow security personnel to achieve a high level of detection and limited false alarms without human involvement. AI may also provide a number of security benefits by augmenting humans in other areas of physical security (e.g. conducting perimeter patrols).

Conclusion

Defending against drone threats will require security managers to take both an immediate and long-term approach, as the threat already exists but is likely to evolve in the coming years. UVs of various shapes, sizes, and capabilities are already employed by hobbyists, businesses, militaries, and threat actors. As the capabilities of drone technology increase and the barriers to use (e.g. cost, accessibility, operability, etc.) decrease, terrorists and other nefarious actors, including lone wolf attackers, will become increasingly likely to adopt the technology. AI and autonomy will likely enhance this trajectory of upward use and enable more lethal, coordinated attacks with less human involvement. Though threat actors have not achieved Slaughterbot-esque, "killer drone" swarm capabilities to date, their less autonomous predecessors have already arrived and should be addressed as part of a comprehensive drone countermeasure strategy.

Additional Resources

World of Drones – New America
Flying IEDs: The Next Big Threat? – War on the Rocks
Remotely Piloted Innovation – Combatting Terrorism Center at West Point
Autonomous Military Drones: No Longer Science Fiction – NATO Review
Why You Shouldn't Fear 'Slaughterbots' – IEEE Spectrum
Why You Should Fear 'Slaughterbots'—A Response – IEEE Spectrum
This Is How Drones Work – Time