Homeland Security
Science and Technology

Homeland Security Science and Technology Advisory Committee (HSSTAC): Quadrennial Homeland Security Review Subcommittee

Artificial Intelligence White Paper

March 10, 2017

Homeland Security

Science and Technology

This publication is presented on behalf of the Homeland Security Science and Technology Advisory Committee, Quadrennial Homeland Security Review Subcommittee, Artificial Intelligence, chaired by Dr. Jim Hendler with contributions from Mr. Byron Collie, Dr. Mark Dean, Mr. Dan Dubno, Dr. Steven Flynn, Mr. Mark Maybury, Mr. Gary Schenkel, Dr. Christina Williams, Dr. Ted Willke as part of recommendations to the Department of Homeland Security, Under Secretary for Science and Technology, Robert Griffin (*Acting*).

<Signature on File>

_____

James A. Hendler
Tetherless World Professor of Computer, Web and Cognitive Sciences
Director, Rensselear Institute for Data Exploration and Applicaions
Rensselear Polytechnic Institute

HSSTAC Staff: Michel Kareis, HSSTAC Executive Director/DFO and Gretchen Cullenberg, QHSR Subcommittee support.

Homeland Security
Science and Technology

ARTIFICIAL INTELLIGENCE AND HOMELAND SECURITY

Dr. Jim Hendler, Subcommittee Chair; Mr. Byron Collie, Dr. Mark Dean, Mr. Dan Dubno, Dr. Steven Flynn, Mr. Mark Maybury, Mr. Gary Schenkel, Dr. Christina Williams, Dr. Ted Willke

White Paper for HSSTAC Quadrennial Homeland Security Review (QHSR) Subcommittee in support of the 2018 QHSR

Introduction and Problem Statement:

Artifical Intelligence (AI), the appearance of logical, sometimes human appearing, analysis and decision making in hardware, software and robotic mechanisms, has been a science fiction staple as both a benefactor and an antagonist since the 1950's. The reality of AI is more complex with advanced, digital logic systems able to learn and perform specific functions in ways exceeding human abilities. These technologies can both support missions to protect the United States, and conversely be employed by our adversaries to identify vulnerabilities and opportunities for malicious actions.

What is Artificial Intelligence?

The term "artificial intelligence" was coined in the mid-1950s to describe the effort to use computers to do tasks that are generally thought of as "intelligent" when humans do them. This defined a very wide range of activities ranging from playing games like chess, understanding human language, and planning activities for the military and many others.

In a February 2017 video[1], John Launchbury, the Director of the Defense Advanced Research Projects Agency's Information Innovation Office, described AI as having three waves: the first, realized in the 1980s and used in many products today, focused on "handcrafted knowledge" – that is the development of rules and procedures which, when applied, were able to apply step-by-step knowledge to a situation and "reason" about the outcomes. These systems, which first came into prominence in the 1980s, were used in medical diagnosis, military planning systems, and many other applications. This technology has now become a "mainstream" technique, and systems have been widely deployed and commercially successful (Launchbury uses the example of TurboTax, one of the most widely used tax preparation systems deploying "handcrafted knowledge"). However, the narrow fields in which these systems are most successfully deployed, and the significant effort in developing applications, has limited their use.

A second wave of AI technology turned from handcrafted rules to techniques in which AI systems used statistical mathematics of various kinds to solve problems based on machine-learning techniques applied to large amounts of data. One of the first major successes of this approach occurred in 1997, when the world's chess champion, Gary Kasparov, was beaten by an IBM system called Deep Blue. This led to new directions in computational approaches, one of the most obvious being significant advances in language processing technologies leading to major improvements in search engines, speech recognition, and many applications that processed large quantities of data to essentially "predict" future situations based on past occurrences. For example, military planners were able to use information about the past deployment of

[1] https://www.youtube.com/watch?v=-O01G3tSYpU

IEDs in the Iraq and Afghanistan theatres to better predict where attacks might occur and to avoid these locations.[2]

One of the most important advances in this latter area has occurred over the past few years, and come to be known as "Deep Learning." Using significant computational power, neural-network based computing has been able to make major advances. For example, it was assumed as recently as 2015, that it would be computationally impossible for computers to be able to beat humans at the game of Go in the foreseeable future due not just to the very large number of possible combinations of game positions, but also due to the wide-range of emerging patterns that needed to be recognized to do well at the game. By March of 2016, a combination of deep learning with other game-playing techniques led to the defeat of one of the best human players by the Google-funded program, AlphaGo. These systems have also been used to significantly improve computational processing of texts, images, sensors and videos. One notable use of these capabilities has been in the significant improvement of autonomous systems, particularly systems for automated vehicles. (The specific threats and capabilities posed by autonomous technology are addressed in a separate HSSTAC white paper, Autonomous Technology, and thus will not be addressed in this paper).

Despite the significant advances in AI coming from the new learning technologies, the contexts in which they can be deployed is still fairly narrow. For example, at the time of this writing, vision systems can be trained to do facial recognition or emotion detection from faces, but not to do the combination of the two at the same time. A 2016 report by the Office of Science and Technology Policy[3] entitled "Preparing for the Future of Artificial Intelligence" categorizes this kind of system as "narrow AI" – able to convincingly handle problems in a specific context, but lacking the capability to replace a human's abilities to apply knowledge learned in one context to another.

The third wave of AI identified by Launchbury will be the ability to do the "contextual adaptation" that is needed to apply what is learned in one area to another, and also to project AI learning over much longer time spans. These technologies are still in the early research stage, and while not likely to have major impact in the next four to seven years, may well have significance to the DHS mission in the next two decades.

In the remainder of this white paper, we look at the threats which current and emerging AI technologies pose to DHS missions and also the opportunities these technologies provide if properly applied to DHS needs.

DHS Mission	Threat to DHS missions from AI	Enhancement of DHS capabilities using AI
Prevent terrorism and enhance security	MT	ST
Secure and manage our borders	MT	ST
Enforce and administer our immigration laws	LT	MT
Safeguard and secure cyberspace	ST	MT

[2] https://www.sciencedaily.com/releases/2009/12/091210153655.htm

[3] https://obamawhitehouse.archives.gov/sites/default/files/whitehouse_files/microsites/ostp/NSTC/preparing_for_the_future_of_ai.pdf

DHS Mission	Threat to DHS missions from AI	Enhancement of DHS capabilities using AI
Strengthen national preparedness and resiliency	ST	MT

ST = within 4 years, MT = 4-7 years, LT = longer time horizon

Table 1: Summary "roadmap" of time to impact of AI technologies on critical aspects of DHS Missions.

Threats from AI to DHS missions

AI, like most powerful technologies, can be applied both in the defense of the homeland or in attacks on homeland infrastructures and systems. Often similar techniques can be used in enforcement (for example, better biometrics) and in frustrating that enforcement (finding the minimal changes needed to prevent biometric identification) when applied by an adversary. However, it is clear that adversarial use of the current and emerging AI technologies could have the greatest impact in two of the key DHS missions: *strengthening national preparedness and resiliency* and *safeguarding and securing cyberspace*. We discuss each of these in turn.

Threats to strengthening national preparedness and resiliency

Probably the greatest threats from current and near-future AI systems to the DHS mission space is in attacks to national resiliency. While AI also has potential to enhance resiliency, as we will discuss in the next section, the threats from AI come from three key areas: social disruption, attacks on national infrastructure and use of intelligent design technologies in the CBRN arena.

Social Disruption: As outlined in the December 2016 report from the Executive Office of the President entitled "Artificial Intelligence, Automation and the Economy,"[4] it is difficult to predict the exact magnitude of job loss and social disruption that can be caused by the deployment of AI technologies across the industrial and service base of our economy, it is likely to be quite large. The job loss from automated vehicles alone could be in the 2-3 million job range, and the overall impacts could be much higher. For pre-high school education, they cite a figure of as many as 44% of the jobs could be highly automated, with 19% of high school jobs, and 8% of trade school positions. Further, the increases in productivity of knowledge workers with AI-based systems could cause significant job loss in more highly educated and highly skilled professions. In other countries, and in past generations, social unrest, mass protests and acts of violence have been endemic in times of major work force disruption especially in urban environments. We particularly note that first responder training for urban unrest may have to focus on significantly different parts of the nation and different infrastructure components that may be effected by a rise of "neoluddism" as workers try to disrupt the automated systems that have put them out of work. Thus, **DHS must consider the impacts on the nation that could be caused by significant changes in the work force leading to protests and riots, increases in homelessness and crime, and violent acts by displaced workers.** The OSTP report does describe steps to mitigate the disruption via education and training, but note that without major investment, the likelihood is that the people most likely to lose their jobs, or unable to be employed, are not likely to be those who will be able to occupy the new occupations that arise. Howeber, on the positive side, we note, however, that using AI in helping to educate displaced workers and in conjunction with service delivery, and esp. the Smart City /

[4] https://obamawhitehouse.archives.gov/sites/whitehouse.gov/files/documents/Artificial-Intelligence-Automation-Economy.PDF

IOT concepts of delivery of services, may also help mitigate some of the more drastic results.)

Disruption of national infrastructure: Machine-learning techniques being developed in AI can be used to identify and categorize the components and interactions in complex systems. Given much of the US infrastructure in areas such as the electrical grid, water infrastructure, and financial systems are complex systems of systems, often under separate ownership and/or control, the ability of AI systems to identify weaknesses and areas of vulnerability becomes manifest. **Countering this ability requires DHS and its partners to develop similar technologies and deploy them in identifying vulnerable infrastructure components that need to be reinforced.**

A longer-term, but very real threat can be caused by "intelligent agent" technologies deployed against financial systems (as well as cyber attacks on the banking system and/or intelligent phishing attacks for identity theft). The 2010 "flash crash" of the US stock market, a 9% drop in the Dow Jones Index in a period of minutes, was attributed to a number of factors including both system problems and trader behaviors. The banning of a number of activities (spoofing, layering and front-running) resulted and e-trading is now monitored to prevent these behaviors. However, **more complex interactions could be enabled by AI algorithms that created complex disruptions in the markets, varying their behaviors to avoid the known detection routines.** Similar disruption techniques have been postulated to be possible by systems which spoofed electrical or water monitoring systems causing cascading effects on national infrastructure. **During the coming four years, DHS S&T efforts should include exploring the magnitude of the potential intentionally introduced threats and identifying mitigation techniques.**

Another area of concern, especially in the longer term, is the use of AI techniques to "intelligently design" both biological threats (via CRISPR genome editing, for example) and physical materials that could evade detection and/or be resilient to counterattack. The separate HSSTAC white paper on CBRN threats discusses emerging threats and counters, and the role of AI in those areas is one that also needs to be explored.

Threats to safeguarding and securing cyberspace

As detailed in the HSSTAC white paper on cybersecurity, the overall safeguarding and securing of cyberspace is a complex and multifaceted problem. However, one of the key problems faced in the cyber world is the speed with which attacks can occur versus the time taken for appropriate counter measures to be deployed. **One potential threat from currently available AI technology is that large scale attacks could be "directed" by the kind of pattern-identification techniques that enabled AI to win at Go.** For example, current Distributed Denial of Service (DDOS) attacks are launched massively, but generally with a prepared attack scenario in which the locus of attack is either singular or "moves" in a predictable way. **An attacker using an AI-based system could be able to detect where defenders are having impact by identifying the patterns of machines where packets are being blocked, and dynamically shift attacks to other IP addresses.** This sort of "directed" attack could be even more significant as the Internet of Things is further realized, and attacks can be directed against not just servers, but individual devices in a massive way. **Countering such attacks needs to be a focus of securing cyberspace.**

An the emergence of potential cyber threats become possible as deep learning AI systems have been shown to be able to debug and even to generate computer code. This work is still primarily in a research stage, but significant commercial investment is being put into this technology as the costs of

program development and the need for the maintenance of complex systems increases. **This technology is currently being used by commercial firms to try to automate the identification of software vulnerabilities to build more secure code, but the same techniques can be used to probe for attack vectors and develop zero-day attacks.** DHS needs to engage the commercial sector and explore how these techniques, often proprietary to the individual software providers, can be understood and better countered.

Other DHS Missions

In the longer term, it is possible to see AI increasingly used by adversaries to "game" US systems and to identify weaknesses beyond those of physical systems to systems with humans involved. As knowledge technologies come on board, dedicated adversaries, if undeterred, will be able to track DHS assets, to identify weaknesses in DHS protections and preparedness and to plan how best to threaten US interests.

Potential use of AI to enhance DHS missions

While there are threats that are posed by the use of AI technologies, as described above, an even greater threat might be not using appropriate AI technologies. Making appropriate S&T investments in AI, to be able to take advantage of the improved capabilities AI systems, could offer significant advantages to the DHS mission areas. We outline some of the key areas below where current and near-term AI technologies could have impacts across the DHS missions. We highlight the emerging technologies and how they could be used to enhance DHS capabilities.

Improved Predictive Systems The use of machine learning techniques, and increasingly the power of deep learning, coupled with large amounts of data is creating a major boon to predictive systems across many industrial sectors. DHS can take advantage of improved prediction in a number of ways:

- New technologies for improvements in predicting natural disasters; including severe weather, earthquakes, etc. **As weather patterns grow more erratic, with extreme weather predicted to increase, better weather forecasting and storm tracking can help DHS to issue earlier warnings and give first responders more time to prepare.**
- Improved social media analytics are being developed for predicting human behavior. **These systems can be deployed in the counter-terror and countering violent extremism arenas to provide improvements to DHS capabilities.**
- Acquisition of more data about patterns of border crossings, illegal entry and attempts to circumvent legal immigration procedures, **predictive capabilities will increase analysts' ability to help agents to enforce customs and immigration laws.**

Improved vision systems and biometrics The rapidly improving capabilities of sensor understanding and fusion algorithms, being enabled by Deep Learning, promise to be commercially available in the near future. These systems may include automated sensor detection or systems which enhance human sensing capabilities. DHS can take advantage of these capabilities for:

- **Extending the coverage and accuracy of intelligent border patrol or protection**
- **Using enhanced sensors for cargo inspection, customs interventions etc.**
- **Increasing the ability of sensors to be more accurate, resulting in the reduction of incidences of false alerts and making it more difficult to avoid by those attempting to circumvent border protections.** These sensors, coming on line over the next four to seven years, should lead to DHS being able to differentiate, for example, humans from animals (rather than just motion), to better identify heat signatures in complex images from thermal imaging, etc. In the longer term,

these sensors should greatly improve the accuracy of sensing activities both under land and water surfaces improving counter smuggling operations.

- **Improving facial recognition, biometrics and other identification techniques using AI systems coming available on the market.** These systems will be increasingly available for purchase and, with specialized training, be available for identifying individuals despite efforts to change appearance or other identifying characteristics. These systems will also enhance the ability to better recognize individuals at a distance in images or videos with multiple participants resulting in improving the enforcement of immigration laws and securing and managing borders.

Optimization and resource allocation: AI systems, based on game theory and other approaches are now transitioning from laboratory to practice. These "agent-based" systems make it possible to take information about resources available, targets of opportunity, and changing information about the environment to make better allocation decisions as to where to deploy limited resources including personnel, sensors, detectors, etc. DHS S&T should invest in the improvement of these systems and commercialization for deployment for use in a number of arenas including:

- Resource deployment decisions. Systems based on these approaches have already been demonstrated in areas such as deploying police in protection of infrastructure and preventing crime, assignment of air guards to flights, and increased alert statuses for potentially targeted responders. **As threats increase, without the ability to continually increase manpower, these systems will become more important for enforcement and responder resource allocations.**
- **These same technologies will also be increasingly useful for the better acquisition and distribution of relief supplies for natural and manmade emergencies**, especially when coupled with better predictive technologies as described above. In particular, there will be advantages in resilience by DHS having resource visibility for advance planning and subsequent delivery of necessary resources in time of need. This can range from logistical concerns, (water, food, power), to critical life saving elements, (doctors, facilities, transportation, public safety).

Other AI techniques, or the integration of the ones above, will also be of importance to the DHS mission areas. DHS will need to improve the ability to automate medical diagnosis for first responders or victims leading to decreased loss of life and to have the ability to enhance the detection of insider cyber (and physical) threats using machine-learned, behavior-based models. AI technologies, that are increasingly becoming available, will be important for the ability to process text, spoken speech, and to translate between languages. In addition, as AI is increasingly used in training systems it could help train DHS personnel in interpersonal relations and other aspects of their jobs that could reduce tensions and improve public relations in respect to the DHS border and immigration activities. Finally, there is a lot of ptenital to AI being deployed for service to citizens (call centers, forms processing, fraud detection, etc).

Longer-Term Recommendations re: artificial intelligence and homeland security

AI, as a field, is not unique in that it has been known for periods of concentrated research with little results followed by large spurts of exponential progress. However, what is somewhat unique is the specific fit of the technologies developed to DHS missions due to the sensor-based capabilities inherent in data-enhanced machine learning applications and the specific threats that these increasingly capable systems will have if applied by adversaries. **We therefore recommend that it is important that DHS S&T maintain a constantly evolving roadmap of AI technologies vs. DHS missions, and monitor this technology even when the current wave of applications, and the attendant marketing hype, starts to dissipate.**

Homeland Security

Science and Technology