

Grand Challenge: Applying Artificial Intelligence and Machine Learning to Cybersecurity

Kirk Bresniker, Hewlett Packard Labs

Ada Gavrilovska, Georgia Institute of Technology

James Holt, Laboratory for Physical Sciences

Dejan Milojevic, Hewlett Packard Labs

Trung Tran, Ampl.io

Cybersecurity is a key risk for any business as the number of attacks are increasing. Growing attacks on cybersecurity are threatening our existence. Artificial Intelligence (AI) and machine learning (ML) can help detect threats and provide recommendations to cyber analysts. Advancing the adoption of AI/ML applied to cybersecurity requires partnering of industry, academia, and government on a global scale.

Cybersecurity is a key concern for nearly every business today.¹ The number of cybersecurity threats is increasing on a daily basis and compromising our private, professional, and national

existence.²⁻⁵ The time required to address threats is increasing, and human capital is under-resourced.⁶ Cybersecurity Ventures estimates that there will be 3.5 million open cybersecurity jobs by 2021, up from 1 million openings last year. People in these jobs are hard to train, and turnover is high. The Enterprise Security Group and the Information Systems Security Association estimate that the average

Digital Object Identifier 10.1109/MC.2019.2942584
Date of current version: 22 November 2019

chief information security officer works for a company only 18 months before moving on to another job.

The growing attack surface includes amateur threats, such as phishing, sophisticated distributed denial of service attacks, and skilled nation-state actors. Prevention is nearly impossible. Given enough time, attackers will get in:

THE BIGGEST PROBLEM IN CYBERSECURITY IS NOT BETTER ENDPOINT DETECTION BUT HOW TO ENABLE THE ANALYST TO KEEP PACE WITH THE SHEER VOLUME OF ALERTS BEING GENERATED.

the cost of attack is low, and automated probing will eventually find a weakness. Advanced persistent threats show that hackers are patient. Defense depends on security analysts who are rare, lack adequate training needed for the job, and have high turnover rates. Artificial intelligence and machine learning (AI/ML) can help with the detection of threats across the enterprise and with recommendations to security analysts. AI/ML can drive down response times from hundreds of hours to seconds and scale analyst effectiveness from one or two incidents to thousands daily. With an adequate knowledge base, it can preserve corporate knowledge and use that knowledge to automate tasks and train new analysts.

AI has been used before in cybersecurity applications for creating pattern-matching tools that alert analysts to security issues in their network. Depending upon the analysis, the tools respond to the events in real time. The

problem is that these tools generate output at rates that quickly outpace the analyst's ability to respond. The biggest problem in cybersecurity is not better endpoint detection but how to enable the analyst to keep pace with the sheer volume of alerts being generated. Humans simply cannot keep up. This leaves us with no other option

but to automate as much as possible⁷ by capturing cybersecurity analysts' existing behavior and thought processes. This can only be done by observing cybersecurity analysts in action. However, this has proven to be challenging due to the proprietary nature of business. One promising way to capture and understand the actions of cybersecurity analysts is by using grand challenges.^{8,9}

Automation also requires a clear definition of goals. This has proven to be very difficult due to the widely varying missions and priorities of computer systems being protected and the multiple simultaneous objectives that often result in tradeoffs. Grand challenges require clear goals, and recurring challenges of increasing complexity can advance the state of the art in representing and measuring complex goals. By gamifying the analysts' actions, we can take advantage of the emerging work from Google's Deepmind and the Open AI team, which

are trying to teach machines how to understand and react to strategies and moves of human actors in such games as *Go* and *StarCraft*.

PREDICTION

For cybersecurity to advance, AI and machine learning (ML) must be used to automate mundane tasks, thus effectively enabling cybersecurity analysts to scale and respond to more events in real time. Unfortunately, cybersecurity is not a static game but one that evolves constantly over time, which means that AI must continuously capture analyst behavior, strategies, successes, and failures to learn new tactics and techniques as they are invented. AI and ML will first assist and then eventually automate lengthy efforts to identify threats and act upon them.

For this to happen, training sets must be created to enable cybersecurity research so that new AI tools can be developed to enhance the effectiveness of current cybersecurity analysts. Operating a public challenge (arena/rodeo) on a recurring basis (at least annually) will ensure that the training sets are constantly capturing the innovations of players in the arena, thus fostering the evolution of AI tools as they are upgraded to keep pace with those innovations.

This will be accomplished by capturing host and network data in a standardized way that allows for the creation of structured and labeled data sets. It will also give us the ability to analyze workflows and the decision processes used by players in the game. Thus, a common repository will be created, containing cybersecurity data and thought processes/strategies of red/blue teams required to train AI algorithms for today and tomorrow.

Achieving this goal requires the following five elements:

1. The behavior of cybersecurity analysts must be captured in a realistic yet controlled cybersecurity setting in a standardized way to record and store behaviors, threats, strategies, and host and network states.
2. Gamification must be introduced to enable experiments (competitions) to be conducted in a constrained and instrumented environment so that the actions of cybersecurity analysts can be tracked in a standardized way.
3. Researchers must consider new types of attacks targeting distributed enterprise and communication infrastructures, including edge computing and 5G.
4. Standardized methods and associated metrics must be developed for representing complex goals of specific challenges or exercises.
5. Global collaboration must be encouraged so that knowledge, strategies, and data can be archived and disseminated worldwide and continuously updated to keep up with state of the art in cybersecurity.

These activities entail research and development, experimental/game design, standardization, global outreach, lobbying, and education. We now explore these five elements in more detail.

Capturing the behavior of cybersecurity analysts

Today, cybersecurity focuses mainly on creating tools that can provide detailed assessments of activities on a network. Very little work has been done to understand the strategies and

moves of the cybersecurity actors themselves. This makes attribution hard and forces defenders into a reactionary mode as they wait for the next move to occur. Waiting to be attacked is not a successful strategy. Mitre's TAXI/STIX ATT&CK frameworks are good ways to do forensic analysis of events after the fact. But nothing currently exists that actually details, in a standardized way, what analysts look for in alerts, logs, and other data and how they should respond to those observations. These form the moves of the game.

Current ML techniques have a powerful ability to learn from examples.

ONE PROMISING WAY TO CAPTURE AND UNDERSTAND THE ACTIONS OF CYBERSECURITY ANALYSTS IS BY USING GRAND CHALLENGES.

However, this requires the existence and availability of a set of examples representing the behavior to be learned. Creating this set by capturing analyst behavior is an essential prerequisite to realizing the potential of ML for automating tasks for cybersecurity analysts.

In existing security operations center (SOC) environments, this task is notoriously difficult. The cybersecurity analyst must have both a deep technological and an intuitive understanding of the behavior of adversaries. Such an analyst might be unable to explain why a particular step was taken to detect or remediate an attack or identify exactly what signal tipped them off that an attack was ongoing.

The first step is to create a standardized knowledge model or ontology that defines a common semantic for describing the activities and decision trees that the analyst should follow for any given set of observed data. The next step is to create standard data-reporting formats (see the section "Standardized Representation of Complex Goals and Associated Metrics"), so that the data being provided by the analyst are consistent and capture a steady workflow. Both steps require an intensive standardization effort because each organization in cybersecurity tends to develop its own knowledge model and reporting formats, or it

lets the analysts develop them on their own. Without a standard, ad hoc and unstructured approaches to reporting actions make it impossible to train an AI system.

Standards should be required for endpoint equipment and other sensors to report events in the format decided upon. Often the raw data collected from these sensors do not report the same values for required fields. In some cases, even the definition of data changes from one vendor to another. The goal is to make the sensor data simple and clear like moves in a chess game, where letters and numbers are sufficient to describe the entire situation on a board without visual cues. These data points will provide ideal

labels for AI and a standard platform for training new analysts.

The intent of these efforts is to create a large corpus of data that can be used to train AI systems. A recent example of how this was done is DeepMind's AlphaGo effort. DeepMind's first step was to take pictures of Go games. The pictures captured the placement of pieces in a standard way so that deep neural networks (DNNs) could be used to classify the players' moves in a way that was simple and easy to understand. Over 30 million moves were captured, and they were placed in a sequence that mirrored the actual moves of the game. Each game was graded based on whether it was successful. DeepMind then built a reinforcement-learning algorithm using Markov tree searches to select what the optimal next move is based on the prior placement of the stones. The same process that allowed DeepMind to capture the moves of the players can be used in cybersecurity to capture the moves of the analyst. But it requires a standard method of capturing actions in cybersecurity, grading those actions based on outcomes, and creating enough examples that it can capture the state space of likely next moves.

Gamification of cybersecurity by running cybersecurity challenges in an annual arena/rodeo setting

A fully instrumented arena will enable the comprehensive and high-resolution capture of analyst actions, including steps to discover, observe, investigate, understand, and mitigate adversary actions (Figure 1). Analysts' choices of tools and data used in each step and their methods for gathering and synthesizing disparate pieces of information would be tracked.

These analyst actions would then be evaluated and understood within the




FIGURE 1. The use of an arena to capture standardized actions of a cybersecurity analyst and store them in a global public database.

context of the scenario or exercise in which they occurred. In the arena, the goals and metrics of success for each scenario are clearly defined. Since the goals are known, analyst actions can be more fully understood as attempts to achieve the goals in the given scenarios. Metrics from the scenario, captured with the analyst actions, provide information on how effective those actions were in achieving the defined goals. All of this can inform learning algorithms on what to do and not do under similar conditions.

In creating an arena/rodeo environment as a fully instrumented digital twin of both the resources to be protected as well as the SOC in which the cybersecurity analysts operate, we establish the environment needed to capture the in vivo behavior of the analyst as well as add external expert analyst commentary. This platform and the standardization of logging form the basis for the training, simulation, and analysis of strategies analogous to those related to strategic competitive games, such as chess and poker. As the global library of real-life threats are captured, modeled, and added, the environment becomes both an epidemiological laboratory and an analyst training tool.

The goals of the arena could be adjusted to the current needs, or they

can vary from year to year. Challenges could be against one opponent or multiple ones, and they could be assisted by AI or against AI (eventually). Possible goals include the following:

- › *Data protection (steal the data; capture the flag).* This challenge consists of gaining access to protected data. The winner is the first to gain access to those data. This challenge can be combined with a variety of other defensive and offensive strategies.
- › *Controlling network (takeover).* This goal includes gaining full access to network and controlling it.
- › *Intrusion detection (penetrate the firewall; hide and seek).* This challenge incorporates penetrating a network without being noticed.
- › *Zero day patching (find, patch, report; bug bounty).* This challenge resembles typical bug bounty, where a bug needs to be discovered, patched in a timely manner, and reported.
- › *Overall system running (not binary, partial recovery).* The goal here is to keep the system running. The system can be partially down, but as long as some essential parts are running, the goal is met.

The arenas for competitions should mirror both the assets to be protected and the tools and interfaces available to the analyst in the SOC. In the case of public or private cloud infrastructures, this will mean the development-suitable containerized environments for both the assets and the SOC. This has the dual advantage of being deployable at varying scales and having the same infrastructure for the same arenas as would be used for the real assets

to be protected. The scalability of containerized workloads allows them to be deployed on any infrastructure from individual laptops and PCs up to a full-scale hybrid cloud. This is also a low barrier to entry, which should encourage individual practitioners and academics to participate. However the arena is built, whether entirely in the public cloud or by a stack of open-source tools to be installed on the premises, it will be made public and easily reproducible. That will enable academic researchers, training companies, government agencies, commercial tool developers, and others to conduct tests related to their own needs in support of their research, development, and operations.

Systems that require a combination of cyberphysical or Internet of Things (IoT) edge devices can either be represented by the physical systems or digital twins of the physical devices. To be realistic and provide true examples of potential threats, arenas should be equipped with actual devices. However, such devices are challenging to maintain in a verifiable state after exposure in the arena. Physical resources are also difficult to provide at scale to individuals and academic environments. Given the challenges of dedicated physical equipment, IT and operations technology equipment manufacturers should be encouraged to provide containerized digital twins of their infrastructure products and perhaps offer bug bounties for discovering possible weaknesses.

Consideration of new tiers in the technology landscape

The number of client and IoT devices is growing and rapidly expanding the digital footprint that can serve as a launchpad for automated attack botnets. This

is evident from such events as the recent Mirai attack, which created a flood of packets at more than 600 Gb/s, generated by an army of 200,000–300,000 compromised IoT devices.¹³ These types of attacks are only going to grow in scope and frequency. The enterprise infrastructure and services deployed closer to the devices, at the perimeters of the enterprise networks, comprise a clear first line of defense to deal with the scale and distribution for these attacks.¹⁴ The edge infrastructure tier presents new opportunities for supporting cybersecurity strategies because

- › It is closer and therefore can see, detect, and react to events more quickly.
- › It presents a new security perimeter and therefore can contain attacks within a smaller footprint.

Given these trends, the arena will provide capabilities to evaluate and analyze the role that the edge-based cybersecurity strategies play in defending against these types of distributed attacks as well as others. The creation of these arenas can be structured in a way that leverages the approach taken in the development of the EdgeNet testbed.¹⁵

The key characteristics of edge computing devices make it clear that these devices will not be able to support the same type of cybersecurity functionality supported in enterprise data centers and clouds. One difference is that the edge has a limited scale compared to the cloud, which is elastic. Another difference is that the functionality deployed on the edge operates in a localized context. By comparison, in data centers and the cloud, functionality is centralized to support a global scope.

These differences impact the interplay of AI/ML and security in several ways. First, concerning AI-based cybersecurity methods, the type of AI that can be executed at an edge component will need to operate within the reduced resource footprint available at the edge, while still providing the timeliness and data-reduction benefits that the edge promises to deliver. Second, the mismatch of the capabilities at the edge versus the data center will lead security strategists to devise and deploy different techniques at the edge than in data centers. For instance, an edge-based detector may serve as a more lightweight estimator of complex classification logic for generating an alert of a potential impending attack and enabling measures to be taken to prevent or mitigate such an attack.¹⁴ Strategists may orchestrate the aggregation of data from edge-based detectors in an application/enterprise-specific manner that combines a customized understanding of data-mining or information theory techniques. These techniques, when applied to their context, compensate for the reduced scope at which individual edge locations operate. For more effective and more scalable cybersecurity, detailed study of these strategies will help determine which can provide the basis for automation in transforming back-end/data center-based cybersecurity defenses into their edge-based counterparts or spokes.

Standardized representation of complex goals and associated metrics

AI and ML algorithms are excellent at optimizing, but to optimize they need a signal, such as a score, indicating how well they are performing. To apply these powerful tools to cybersecurity problems, it is necessary to have training

data and scenario outcomes associated with scores. The performance of the models trained on this data will be very sensitive to how these scores are calculated, so it is important to make sure the scores truly represent progress towards achieving the desired outcome.

To generate these scores, it is necessary to have definitions of the goals for any given scenario, challenge, or exercise. As the purpose and mission of computer systems vary widely, the goals vary from scenario to scenario in the

The answers will be different depending on the environment. A representation of goals for this environment might describe a collection of goals and assign weights to each of them, expressing their relative importance.

For each goal that is defined, it is necessary to have a way to measure it or at least approximate it. Many things can be measured in a networked system of systems. Determining what to measure and what each measure reveals about the defined goals is an important

policies for AI and ML configuration, AI strategies, the execution environment, and the captured knowledge of security analysts. It is important that this effort be globally supported by such organizations as the IEEE. The IEEE can provide a home to standardization and be a trusted and neutral (politically and organizationally) global body to run the arenas, store the knowledge base, disseminate results using its *Xplore/Computer Society Digital Library* database, and pursue education using online and in-person courses. A successful outcome can only be obtained if the global aspect is accepted by worldwide cybersecurity experts.



THE APPLICATION OF AI AND ML TO CYBERSECURITY HAS MANY BENEFITS.

IMPACT

The application of AI and ML to cybersecurity has many benefits. At the entry level AI and ML will eliminate some threats, such as phishing, at the noise level. But as AI and ML are applied to a broader set of more complex cases, their value in the field of cybersecurity will become more apparent. They will enable system owners to detect earlier and deter more efficiently some of the distributed denial-of-service attacks and prevent data leakage and network penetration. The real benefits of this approach are in sharing the results globally so that attacks can be prevented closer to the source, and we can create a community effort to regulate the current Wild West of cybersecurity.

By gathering data for AI and ML training, industry, academia, and governments of the world will be able to establish a starting point for additional work that may or may not be shared. Sharing, in some cases, may be limited. Some companies may not want to share data for competitive reasons. Governments may wish to keep data confidential to protect national security. However, all initial results will be globally shared.

arena and from environment to environment in real operational networks.

Generally, system owners have multiple goals that compete with each other, leading to tradeoffs. For example, a system owner may want a data store to be available, authorized users to be able to access it, users to be able to communicate with each other and with people outside the network, and the data flow coming into the network to be uninterrupted. But the owner also wants to maintain current updates/patches on his or her systems to minimize vulnerability, prevent unauthorized access to accounts and data, prevent system intrusions via spearphishing, detect any unauthorized data exfiltration, and prevent attackers from controlling any of their systems. Those are many different goals. The key question is how important are they relative to each other? Is it more important to allow users to communicate outside the network or prevent data exfiltration?

task that will require a sustained effort of the cybersecurity community.

Developing standards for expressing complex goals like those described here is a critical step. Developing standards for cyber metrics, and for ways to connect those metrics to goals, is also critical. These interconnected standards, in combination, will make it possible to define complex scenarios with complex goals and produce scores indicating how successful attackers and defenders are during a scenario. This, in turn, will enable the development and evolution of increasingly complex and realistic challenges, scenarios, exercises, and training modules.

Global collaboration

In the coming years, the proposed grand challenge will result in a repository of training data, shared data logs, and attack details. The resulting knowledge base will contain captured network traffic,

ABOUT THE AUTHORS

KIRK BRESNIKER is the chief architect of Hewlett Packard Labs and a Hewlett Packard Enterprise (HPE) fellow and vice president. Bresniker received a B.S.E.E. from Santa Clara University, California. He is a Senior Member of the IEEE. Contact him at kirk.bresniker@hpe.com.

ADA GAVRILOVSKA is an associate professor at the College of Computing and the Center for Experimental Research in Computer Systems at the Georgia Institute of Technology, Atlanta. Her interests include operating systems, virtualization, systems software for heterogeneous many-core platforms, emerging nonvolatile memories, large-scale data-center and cloud systems, and high-performance communication technologies. Gavrilovska received a Ph.D. in computer science from the Georgia Institute of Technology. She currently serves as an associate editor for *IEEE Transactions on Cloud Computing* and as a technical program cochair for the USENIX Annual Technical Conference. She is a Member of the IEEE, Association for Computing Machinery, and USENIX. Contact her at ada@cc.gatech.edu.

JAMES HOLT is a researcher at the Laboratory for Physical Sciences, where he focuses on applying AI and ML techniques to address cybersecurity problems. His research interests include malware detection, autonomous cyberdefenses, processing streaming data, and anomaly detection. Holt received a B.A. in computer science and mathematics from East Carolina University, Greenville, North Carolina. He is a member of the American Association for Artificial Intelligence. Contact him at holt@lps.umd.edu.

DEJAN MILOJICIC is a distinguished technologist at Hewlett Packard Labs. His research interests include operating systems, distributed systems, and systems management. Milojicic received a Ph.D. from the University of Kaiserslautern, Germany. He received multiple best paper awards. He is a member of the editorial boards of *IEEE Internet Computing* and *IEEE Transactions on Cloud Computing*. He is general cochair of the IEEE Infrastructure and Association for Computing Machinery (ACM) Middleware (Industry track) conferences. He is a Fellow of the IEEE, an ACM distinguished technologist, and a member of Eta Kappa Nu and USENIX. Contact him at dejan.milojicic@hpe.com.

TRUNG TRAN is an entrepreneur focused on developing second- and third-generation AI systems. His research focus is on finding new computing paradigms and unique approaches to understanding how machines can help us think and see the world in new ways. He received an M.B.A. from the Wharton School of the University of Pennsylvania. He is a member of the American Association for Artificial Intelligence and Tau Beta Phi.

Finally, outcomes could be used for training human cybersecurity analysts by providing access to the behavior of recorded actions during the arena competitions and enabling them to play the defender in arena scenarios and receive feedback in the form of a score. The knowledge database can be mined and leveraged for creating courses.

Most importantly, through improved tools to increase automation and support the analyst, AI and ML will dramatically increase the speed and scale of our cybersecurity defenses and the speed and effectiveness of our human analysts. Through improved training capabilities, AI and ML will improve the training of new cybersecurity experts and enhance the quality of cybersecurity analysts available for hire.

TECHNOLOGY CHALLENGES

Three main actions are needed to address technology challenges.

1. Research is needed on AI and ML for cybersecurity primarily focused on deep-learning algorithms. Many other approaches are available, but only DNNs have achieved wide adoption. Other techniques for automating cybersecurity analyst behavior may need to be explored.
2. The ability to effectively standardize vertical data from the network traffic, including everything from attack signatures to cybersecurity analyst strategies, needs to be developed. Correlating data captured across the stack will absolutely be required.
3. Methods need to be devised to distribute AI and ML across the attack surface with both centralized SOC-like deployment and deployment at the edge.


Maintaining consistency in behavior and real-time distribution of models will require new research and development.

RISKS TO PREDICTION

The obvious risk is time of adoption. The application of AI and ML to cybersecurity is inevitable, but the question is whether any meaningful use of such an application will emerge in the next two years. State-of-the-art AI and ML are still primarily focused on deep-learning neural networks trained using labeled data sets and applied to limited sets of use cases, such as video recognition for assisted driving and in industrial IoT scenarios. We anticipate that an initially limited adoption of AI and ML will occur using rule-based techniques. Over the next few years, the transition to a broader use of AI and ML will take place. Differences in regulations across the world will complicate deployment because governments and cultures have varying attitudes about automation, ethics, and privacy.

An unwillingness of global constituencies to collaborate will limit both the breadth and speed of adoption. We anticipate that, as the world becomes more connected and the playing field levels, (and everyone is exposed to each other), global constituencies will be more open to collaboration, especially if the benefits of using this approach become substantial (reduction of threats, speed to resolution, training workforce, and so on).

Cyberattacks are becoming an increasing concern for businesses. Adequate cybersecurity requires automation, which, in turn, requires a means of capturing the behavior of cybersecurity analysts. We predict that this will start to happen within next two years

as a global, annual competition event. As a result, increased adoption of AI and ML in cybersecurity will reduce the effectiveness and impact of attacks. 

REFERENCES

1. "Artificial intelligence and machine learning applied to cybersecurity," IEEE, Piscataway, NJ, 2017. [Online]. Available: https://www.ieee.org/content/dam/ieee-org/ieee/web/org/about/industry/ieee_confluence_report.pdf
2. "DHS cyber security initiatives," United States Computer Emergency Readiness Team, Feb. 6, 2013. [Online]. Available: <https://www.us-cert.gov/security-publications/dhs-cyber-security-initiatives>
3. "The big data breach suffered by Equifax has alarming implications," *The Economist*, Sept. 16, 2017. [Online]. Available: <https://www.economist.com/finance-and-economics/2017/09/16/the-big-data-breach-suffered-by-equifax-has-alarming-implications>
4. D. J. Trump, "Executive order on America's cybersecurity workforce," May 2, 2019. [Online]. Available: <https://www.whitehouse.gov/presidential-actions/executive-order-america-cybersecurity-workforce/>
5. A. Greenburg, "The untold story of NotPetya, the most devastating cyberattack in history," *Wired*, Aug. 22, 2018. [Online]. Available: <https://www.wired.com/story/notpetya-cyberattack-ukraine-russia-code-crashed-the-world>
6. K. Corbin, "Cybersecurity pros in high demand, highly paid, and highly selective," *CIO*. Aug. 8, 2013. [Online]. Available: <https://www.cio.com/article/2383451/careers-staffing/cybersecurity-pros-in-high-demand-highly-paid-and-highly-selective.html>
7. M. Braga, "In the future, we'll leave software bug hunting to the machines," *Motherboard*, June 16, 2016. [Online]. Available: https://motherboard.vice.com/en_us/article/mg73a8/cyber-grand-challenge
8. D. Frazee, "Cyber Grand Challenge (CGC)," DARPA. Accessed on: June 2019. [Online]. Available: <https://www.darpa.mil/program/cyber-grand-challenge>
9. "Challenges of challenge." Accessed on: June 2019. [Online]. Available: <https://challenge.gov>
10. "Enabling distributed security in cyberspace," Department of Homeland Security, Oct. 4, 2016. [Online]. Available: <https://www.dhs.gov/enabling-distributed-security-cyberspace>.
11. New York Cyber Task Force, "Building a defensible cyberspace," Columbia Univ. School of International and Public Affairs, New York City, Sept. 28, 2017. [Online]. Available: https://sipa.columbia.edu/sites/default/files/3668_SIPA%20Defensible%20Cyberspace-WEB.PDF
12. OASIS STIX, "OASIS Cyber Threat Intelligence (CTI) TC," 2017. [Online]. Available: <https://www.oasis-open.org/committees/cti/>
13. M. Antonakakis et al., "Understanding the Mirai botnet," in *Proc. 26th USENIX Security Symp.*, 2017, pp. 1093-1110.
14. K. Bhardwaj, J. C. Miranda, and A. Gavrilovska, "Toward IoT-DDoS prevention using edge computing," in *Proc. USENIX HotEdge'18*, July 2018. [Online]. Available: <https://www.usenix.org/node/216764>
15. EdgeNet Project. Accessed on: Oct. 14, 2019. [Online]. Available: www.edge-net.org