2020 CACR AI/ML Lessons Learned

Exploring Machine Learning for Cybersecurity

Ryan Kiser, Emily K. Adams, Austin Cushenberry, Ishan Abhinit, Kelli Shute

Acknowledgements

This work was funded by the Indiana University Center for Applied Cybersecurity Research, the Indiana University Vice President for IT, and the Indiana University Vice president for Research. We would also like to thank the NSF Cybersecurity Center of Excellence, Trusted CI.[1]

The work summarized in this document would not have been possible without the substantial efforts of many people. We would like to thank the following individuals for their contributions and time:

Dr. S. Jay Yang, Rochester Institute of Technology
Dr. Ahmet Okutan, Rochester Institute of Technology
Simon Kirkwood, Rochester Institute of Technology
Gordon Werner, Rochester Institute of Technology
Ayush Goel, Rochester Institute of Technology
Ren Chauret, Rochester Institute of Technology
Steven Su, Rochester Institute of Technology
Tom Davis, Associate Vice President for Information Security, Indiana University
Andrew Korty, Chief Information Security Officer, Indiana University
Keith Lehigh, University Information Security Officer, Indiana University
Scott Orr, SOC Operations Manager, OmniSOC
Rob Carlsen, Lead Security Engineer, OmniSOC
CJ Kloote, Platform Engineering Lead, OmniSOC
Dr. Xiaojing Liao, Indiana University Luddy School of Informatics, Computing, and Engineering
Matt Link, Associate Vice President for Research Technologies, Indiana University
Rob Henschel, Director, Research Software and Solutions, Indiana University
Dr. Scott Michael, Manager, Research Applications and Deep Learning, Indiana University
Dr. Scott Teige, Research Applications and Deep Learning, Indiana University

[1] NSF Award 1920430

Executive Summary

Since Fall of 2019, the Indiana University Center for Applied Cybersecurity Research (CACR) has been exploring the application of machine learning to cybersecurity workflows with the intent of developing the applicable expertise necessary to maintain a commanding lead in the cybersecurity domain where machine learning solutions are expected to increasingly become the norm. In order to serve the objectives laid out in the project charter, CACR primarily worked in partnership with OmniSOC and researchers at Rochester Institute of Technology to explore the application of the ASSERT research prototype to SOC analyst workflows.[2] The intent of this effort was to better understand both the utility of the ASSERT prototype and the challenges associated with the implementation of machine learning approaches to cybersecurity workflows more broadly.

At a fundamental level "machine learning" is a term which encompasses a powerful set of statistical techniques for analysis of data with significant applications to cybersecurity. Security practitioners cannot afford to ignore the capabilities available through use of these techniques. The application of machine learning techniques to cybersecurity comes with new considerations and challenges which require specialized expertise both in machine learning and the domain in which the solution is to be applied. We recommend that machine learning be treated as a toolkit which can greatly amplify the effectiveness of security practitioners and caution against treating these techniques as a panacea for the longstanding personnel shortages in the cybersecurity field.

Project Activities

Initiation & Planning

Artificial intelligence and machine learning (AI/ML) technologies are on track to becoming key to cybersecurity operations grant competitiveness. UITS Research Technologies is investing heavily in high-performance computing hardware, supporting software infrastructure, and staff expertise to support AI/ML research at Indiana University.[3] [4] IU is well positioned to explore the application of machine learning to the cybersecurity domain through the capabilities established within organizations such as OmniSOC, ResearchSOC, and REN-ISAC. In order to position IU to better take advantage of these opportunities, CACR chartered a project with the goal of developing expertise in how AI/ML can be used to solve cybersecurity challenges.

CACR focused initial efforts on a survey of machine learning techniques with limited explorations into some ongoing research topics. These early efforts included review of published materials such as research papers, books, training materials, and third-party interviews with subject matter experts. These initial explorations provided key context of the history and current state of AI/ML which we used to frame potential use cases for machine learning in cybersecurity. This also provided CACR staff with the

[2] Okutan, A., Yang, S.J. ASSERT: attack synthesis and separation with entropy redistribution towards predictive cyber defense. Cybersecur 2, 15 (2019). https://doi.org/10.1186/s42400-019-0032-0
[3] https://itnews.iu.edu/articles/2020/IU-unveils-supercomputer-Big-Red-200%20-.php
[4] https://news.iu.edu/stories/2020/06/iub/releases/01-jetstream-cloud-computing-awarded-nsf-grant.html

basic understanding needed to conduct meaningful discussions with cybersecurity researchers actively investigating machine learning approaches.

Late in August 2019, we met with Dr. Jay Yang to discuss some of his research applying machine learning to cybersecurity. Dr. Yang is a researcher from Rochester Institute of Technology who was an inaugural member of the Trusted CI Fellows program.[5] He has several machine learning research prototypes and publications which were of interest to the CACR team. Of these, ASSERT seemed most relevant for potential applications to cybersecurity operations at IU through IU's role in the operation of OmniSOC and ResearchSOC.

ASSERT uses information theory based measures to learn from intrusion alerts across intrusion alerting platforms. Over time, ASSERT learns to synthesize empirical "attack models", each of which abstracts a unique cyberattack behavior derived by the observed malicious activities. By analyzing these attack models instead of viewing and sorting through the large number of intrusion alerts, analysts may discover critical relations between intrusion events and thus spend time more effectively to reveal the high impact tactics and patterns and focus on attacks that target critical assets. ASSERT attack models may also be leveraged to assess the aggregate of data sources from multiple organizations served by OmniSOC to provide predictive awareness or intelligence on attack activities.

Project Timeline

CACR's initial efforts were focused on understanding how the ASSERT prototype could improve the ability of a SOC analyst to mitigate cybersecurity threats. We established a plan to foster a collaboration between OmniSOC and RIT researchers with the goal of developing an understanding of how the ASSERT prototype could be applied to effectively analyze IU network intrusion alert data.

To pursue this, we developed a data usage proposal for Suricata[6] alert data and began work on the necessary test infrastructure to pipeline an anonymized subset of the IU Suricata data provided to OmniSOC to a prototype deployed on a host within IU's Intelligent Infrastructure. The intent was to provide results to OmniSOC analysts to determine if the method correctly identified attack characteristics beneficial to their threat hunting workflows. The architectural details of the deployment are documented in Appendix A.

Once the test environment was complete and the appropriate parties had approved the usage of the data, OmniSOC began to send data to the prototype. CACR staff and the ASSERT developers were then able to begin to refine the testbed configuration based on the results produced by the prototype and evolving operational needs for it as refinements were made to the software.

In January 2020, CACR developed a plan for activities during the first half of 2020. Due to unforeseen circumstances resulting in the ASSERT visualization specialist being unavailable, we implemented a contingency plan for CACR staff to attempt to pilot an alternate way to interact with results generated by ASSERT which resembled existing workflows at OmniSOC. Given that Elastic Stack[7] provides a key

[5] https://www.trustedci.org/fellows/about
[6] https://suricata-ids.org/about/
[7] https://www.elastic.co/what-is/elk-stack

suite of tools used by OmniSOC's analysts, we determined that the core focus of this effort would be to attempt to integrate ASSERT's outputs with Elasticsearch[8] and Kibana.[9] This would also provide CACR staff with the opportunity to explore tools and workflows used by OmniSOC to better understand how ASSERT could be used in a production environment.

This presentation of ASSERT results was also an opportune basis for Austin Cushenberry, an undergraduate student at the Luddy School of Informatics, Computing, and Engineering, to develop a capstone project for the B.S. in Informatics with a Human Centered Computing cognate. The confluence of these factors provided the direction for continuing work. At the conclusion of this initial effort we began work early in 2020 with the intent to develop methods to interact with ASSERT results in Kibana.

To maximize the benefits of the project, we conducted a series of activities to document and disseminate our process and what we learned. In addition to this report, we have created guidance which summarizes our recommendations to further mature the ASSERT prototype as well as offer specific recommendations regarding its application to a broader set of IU data. This report is being provided to OmniSOC and the developers of the ASSERT prototype. We have submitted a case study proposal to an applied cybersecurity conference[10] to describe some of what we learned using ASSERT to analyze alerts as well as a proposal for a training session to the NSF Cybersecurity Summit[11] on using machine learning approaches in cybersecurity operations.

Lessons Learned

1. Machine learning and security

Machine learning encompasses a powerful set of statistical techniques which can be used to accomplish things which would be prohibitively difficult or in some cases impossible to encode into systems otherwise. Machine learning solutions are high effort to implement and require specific expertise both in the methodologies being used and the subject matter the solution is being applied to. Cybersecurity is a broad domain with continuously evolving demands and an asymmetrical adversarial relationship. The capabilities enabled by machine learning will be increasingly adopted out of necessity by organizations with limited resources who have to hire from a limited pool of experienced cybersecurity professionals. These problems are particularly pronounced in research and educational institutions where cybersecurity resources are often limited.

It is our assessment that machine learning cannot be viewed as a panacea to personnel shortages. Applying machine learning techniques effectively requires expert knowledge of the operating environment to interpret results and make important distinctions between valuable data and misleading data. In addition, deep knowledge of the data and statistical approaches being used is required to effectively apply these techniques to solve specific problems. Determining whether a problem is amenable to a machine learning solution, producing a viable machine learning solution to said problem, and using said solution in a viable manner requires both skill sets to varying degrees. Machine learning

[8] https://www.elastic.co/elasticsearch/
[9] https://www.elastic.co/kibana
[10] https://www.acsac.org/
[11] https://www.trustedci.org/2020-nsf-summit

provides a powerful set of tools and techniques which can be leveraged by skilled personnel to enable them to accomplish more. That is, it *augments* rather than *replaces* human expertise and effort.

Due to the asymmetric relationship between attacker and defender and the ever increasing complexity of the challenges faced by security professionals, defenders cannot afford to ignore machine learning techniques as a manner to procure solutions. We expect that applications of these approaches in cybersecurity will by necessity continue to grow over time. Cybersecurity professionals will need to understand the utility, challenges, and common pitfalls of machine learning solutions in order to effectively make use of them.

2. ASSERT Example Use Cases

ASSERT is a machine learning system which automatically categorizes related attacker behaviors derived from alerts and other information into descriptive models which can help analysts and other security practitioners to better understand attacks. Over the course of this work, we developed example use cases for SOC analyst use of the models generated by ASSERT. Some of these include:

1. Sudden changes in models indicate changes in attacker behavior associated with the model. These rapid changes suggest a chain of events which require further investigation. For example, if the traffic flow direction of new events associated with a specific model suddenly changes it may be indicative that a successful compromise occurred and follow-up behavior such as data exfiltration or lateral movement is now being attempted from the compromised host.
2. Certain attacks generate very few alerts or alerts which appear similar to less important alert types. On the other hand, some malicious activities result in a large number of alerts which are duplicative. Models generated by ASSERT can identify activity which is important but also generates only a limited number of alerts or related alerts of a type which may otherwise be considered low-priority. The models generated by ASSERT provide a way to alert analysts to high-importance, low-signal event types that may be drowned out by high-volume, low-priority alerts . This determination can in turn be leveraged to signal an increase in the importance of certain events and more importantly, identify new relationships between events which can be tracked by other methods.
3. The ASSERT methodology has the capability to derive models with high volume and low variability which may indicate that normal and expected behavior within an environment is causing alerting mechanisms to be triggered *i.e.* false positives or low priority malicious activities. This presents an opportunity to further tune IDS system configurations and SOC analyst threat hunting activities. For example, ASSERT has the capability to map various Suricata alerts into an “Information Discovery” category. If an attack model continues to contain a majority of “Information Discovery” alerts, even those that include diverse ports and services, this can inform the analyst of scanning activities that may be categorized as inconsequential network activity.
4. ASSERT utilizes `Attack Intent State’ (AIS) to map intrusion alert descriptions to the likely intended consequences of an observed attack action.[12] The AIS mapping facilitates

[12] https://arxiv.org/pdf/2002.07838.pdf

interpretation of intrusion observables of attacker intent by decomposing the MITRE ATT&CK[13] framework 12 enterprise tactics into ~30 types of attack activity that do not rely on platform-specific techniques. Yang's research group has now developed a machine learning model to automatically map Suricata alert descriptions into AIS, and will extend the model to manage other intrusion detection and network monitoring systems such as Zeek.

3. Elastic Stack & ASSERT

The CACR proof-of-concept visualizing of ASSERT data through Elasticsearch and Kibana was only partially successful due to limitations in Elastic Stack expertise within the CACR team coupled with the complexity of translating the ASSERT model data into a form readily usable in Elasticsearch. We were ultimately able to restructure the data and construct a series of exploratory visualizations in Kibana based on preliminary ASSERT data successfully imported into CACR's implementation of Elasticsearch. Specific details of the challenges we faced regarding Elastic Stack as well as the results we were able to achieve are captured in the ASSERT Implementation Strategies document delivered to OmniSOC and the ASSERT developers to guide any future efforts of integrating ASSERT model data into SOC operations.

4. Researcher Needs and Cyberinfrastructure Operations

Although researchers at academic institutions such as IU and RIT have specialized knowledge in their domain of expertise, they often lack insights into the operational realities of the environments and organizations where results of their research may have the most significant impact. Bridging the gaps between research and its practical applications is the focus of several programs at funding agencies including the National Science Foundation,[14] Department of Homeland Security,[15] and the Department of Energy.[16] Among higher-ed and research institutions, IU is uniquely well situated to enable this transition to practice in technology domains due to the comprehensive nature of IT operations in place at IU.

In addition to the substantial research cyberinfrastructure and expert guidance made available to researchers by UITS Research Technologies, UITS has a wealth of operational resources which could be made available to serve research needs including operational data, data descriptions, the university's cyberinfrastructure itself, and direct engagement with infrastructure operations professionals. Research Technologies in particular is situated to assist researchers in designing workflows and working to integrate researchers into their operational environments.

5. Risk Management for Operational Data Usage

IU is home to substantial operational cybersecurity and cyberinfrastructure operations such as OmniSOC, ResearchSOC, REN-ISAC, and GlobalNOC which include a wealth of data, practitioner expertise, and resources relevant to cybersecurity research. Research Technologies provides significant

[13] https://attack.mitre.org/
[14] https://www.nsf.gov/funding/pgm_summ.jsp?pims_id=504709
[15] https://www.dhs.gov/science-and-technology/TTP
[16] https://www.energy.gov/technologytransitions/office-technology-transitions

resources for use by machine learning researchers at IU and has developed significant expertise in the domain through engagement with over 100 projects utilizing these resources. These hardware resources and expertise could be leveraged to accelerate cybersecurity research, including that which occurs at the Luddy School of Informatics, Computing, and Engineering, the Kelley School of Business, and other IU schools and research departments.

Leveraging these resources and providing operational data to researchers will require careful consideration of security risks as well as flexibility to address the dynamic needs of research projects. The current operational data approval process provides flexibility, but is opaque to researchers seeking approval and can involve substantial turnaround times and human effort to review and maintain documentation of data usage. Both researchers and those responsible for approval of data usage would be well served by clear communication about the process. These should include a high-level overview of the approval process which describes the different stages of an approval, stepwise instructions for a researcher to follow to seek approval, and a communication channel which can be used by both researchers and approving parties to convey further information as needed.

Both parties should expect to remain in communication throughout the lifecycle of a project and have the capacity to address requests for more information or updates. Researchers should expect to update documentation that describes details such as data retention, system security controls, and data access as their project evolves over time. Approving parties should anticipate the need to clearly communicate about risks and approaches which can be used by researchers to mitigate them. These expectations should be explicitly set at the beginning of the process.

This approval process itself could be further streamlined by identifying specific data sets appropriate for use by researchers, and where there are data sensitivity concerns, documenting appropriate means researchers can use to protect this data. Related work is already being done at CACR to document secure and compliant use cases in the form of a research data "cookbook" via the SecureMyResearch initiative.[17] Identifying data sets appropriate for research and appropriate means researchers can take to protect them in advance would allow common institutional data use cases to be addressed in a more simple manner and should free up resources for handling more complex requests and communicating with researchers who are seeking approvals.

[17] https://cacr.iu.edu/projects/SecureMyResearch/index.html

Appendix A: System Description

This appendix provides a brief description of the data flow of Suricata alert data from OmniSOC to the ASSERT prototype. The following diagram describes the system architecture: