9th International Workshop on Artificial Intelligence and Security (AISec 2016)

David Mandell Freeman
LinkedIn Corporation, USA
dfreeman@linkedin.com

Katerina Mitrokotsa
Chalmers University of Technology, Sweden
aikmitr@chalmers.se

Arunesh Sinha
University of Michigan, USA
aruneshsinha@gmail.com

Background

Artificial Intelligence (AI) and Machine Learning (ML) provide a set of useful analytic and decision-making techniques that are being leveraged by an ever-growing community of practitioners, including many whose applications have security-sensitive elements. However, while security researchers often utilize such techniques to address problems and AI/ML researchers develop techniques for Big Data analytics applications, neither community devotes much attention to the other. Within security research, AI/ML components are usually regarded as black-box solvers. Conversely, the learning community seldom considers the security/privacy implications entailed in the application of their algorithms when they are designing them. While these two communities generally focus on different directions, where these two fields do meet, interesting problems appear. Researchers working in this intersection have raised many novel questions for both communities and created a new branch of research known as secure learning. The AISec workshop has become the primary venue for this unique fusion of research.

In recent years, there has been an increase of activity within the AISec/secure learning community. There are several reasons for this surge. Firstly, machine learning, data mining, and other artificial intelligence technologies play a key role in extracting knowledge, situational awareness, and security intelligence from Big Data. Secondly, companies like Google, Facebook, Amazon, and Splunk are increasingly exploring and deploying learning technologies to address Big Data problems for their customers. Finally, these trends are increasingly exposing companies and their customers/users to intelligent technologies. As a result, these learning technologies are being explored by researchers both as potential solutions to security/privacy problems and also as a potential source of new privacy/security vulnerabilities that need to be addressed. The AISec Workshop meets this need and serves as the sole long-running venue for this topic.

AISec, having been annually co-located with CCS for nine consecutive years, is the premier meeting place for researchers interested in the junction of security, privacy, AI, and machine learning. Its role as a venue has been to merge practical security problems with advances in AI and machine learning. In doing so, researchers also have been developing theory and analytics unique to this domain and have explored diverse topics such as learning in game-theoretic adversarial environments, privacy-preserving learning, and applications to spam and intrusion detection.

AISec 2016

The ninth annual event in this series, AISec 2016 drew a record 38 submissions, of which approximately ten were selected for publication and presentation. Submissions arrived from researchers in 16 countries, from a wide variety of institutions both academic and corporate. Paper topics included the following:

- Theoretical topics related to security: adversarial learning, robust statistics, learning in games, economics of security, differential privacy.
- Security applications: computer forensics, spam detection, phishing detection and prevention, botnet detection, intrusion detection and response, malware identification, authorship identification.
- Security-related AI problems: distributed inference and decision making for security, privacy-preserving data mining, adaptive side-channel attacks, design and analysis of captchas, AI approaches to trust and reputation, vulnerability testing, security policy management & access control, anomalous behavior detection.
- Machine learning and security at scale: high-throughput abuse detection systems, large-scale active learning, big data analytics for security, techniques dealing with well-resourced adversaries.

The keynote address was given by Elie Bursztein of Google, Inc., who discussed challenges in the reproducibility of scientific results from machine learning algorithms and what we can do about it. Dr. Bursztein's talk touched on issues arising from proprietary hardware, dataset availability, adversarial machine learning, and the ethics of data. He also considered several privacy questions related to machine learning models.

CCS'16 October 24-28, 2016, Vienna, Austria

ACM ISBN 978-1-4503-4139-4/16/10.

DOI: http://dx.doi.org/10.1145/2976749.2990479