The AAAI-19 Workshop on

Artificial Intelligence for Cyber Security (AICS)

January 27 or 28, 2019
Honolulu, Hawaii, USA

Part of the Association for the Advancement of Artificial Intelligence 2019 Conference https://aaai.org/Conferences/AAAI-19/

CALL FOR PARTICIPATION

AICS Workshop Description

This year the AICS workshop emphasis will be on adversarial learning. The workshop will address technologies and their applications, such as machine learning, game theory, natural language processing, knowledge representation, automated and assistive reasoning, and human machine interactions. The workshop will emphasize research and applications of techniques to attack and defend machine learning systems, especially in the context of cyber security.

Machine learning capabilities have recently been shown to offer astounding ability to automatically analyze and classify large amounts of data in complex scenarios, in many cases matching or surpassing human capabilities. However, it has also been widely shown that these same algorithms are vulnerable to attacks, known as adversarial learning attacks, which can cause the algorithms to misbehave or reveal information about their inner workings. In general, attacks take three forms: 1) data poisoning attacks inject incorrectly or maliciously labels data points into the training set so that the algorithm learns the wrong mapping, 2) evasion attacks perturb correctly classified input samples just enough to cause errors in classification, and 3) inference attacks that repeatedly test the trained algorithm with edge-case inputs in order to reveal the previously hidden decision boundaries. As machine learning-based AI capabilities become incorporated into facets of everyday life, including protecting cyber assets, the need to understand adversarial learning and address it becomes clear.

The above discussion of adversarial learning highlights one of the main challenges in applying AI to problems in cyber security. Poisoning attacks that inject incorrectly labeled malicious traffic or data can be leveraged by the adversary to enable their attacks to go undetected, while data evasion attacks can be used to cause false classification of benign traffic as malicious thereby eliciting a defense response. If AI is to succeed in helping cyber security, it must be secure and robust to attacks itself.

This year we are asking the AI for cyber security community to submit solutions to a challenge problem. The challenge problem (http://www-personal.umich.edu/~arunesh/AICS2019/challenge.html) is focused on solving an adversarial attack scenario based on redacted data.

Understanding and addressing challenges associated with adversarial learning requires collaboration between several different research and development communities, including the artificial intelligence, cyber security, game theory, machine learning, as well as the formal reasoning communities. This workshop is structured to encourage a lively exchange of ideas between researchers in these communities from the academic, public, and commercial sectors.

Workshop Topics

- Machine learning approaches to make cyber systems secure and resilient
  - Natural language processing techniques
  - Anomaly/threat detection techniques
  - Big Data noise reduction techniques
  - Adversarial learning
  - Human behavioral modeling
- Formal reasoning, with focus on human element, in cyber systems
- Game theoretic reasoning in cyber security
- Robust AI metrics
- Multi-agent interaction/agent-based modeling in cyber systems
- Decision making under uncertainty in cyber systems
- Modeling and simulation of cyber systems and system components
- Automation of data labeling and ML techniques that learn to learn
- Quantitative human behavior models with application to cyber security
- Operational and commercial applications of AI
- Adversarial planning

Challenge Problem*

For information on this year's AICS challenge problem:
http://www-personal.umich.edu/~arunesh/AICS2019/challenge.html

Workshop Format

Invited speakers, presentations, panel and group discussions

Submission Requirements

One of two submissions is solicited:

- Full-length papers (up to 8 pages in AAAI format)
- Challenge problem papers (up to 8 pages in AAAI format)

Submissions are not anonymized. Please submit PDF via AICS Workshop website. **The deadline to submit papers has been extended to November 12, 2018.**

Workshop URL

http://www-personal.umich.edu/~arunesh/AICS2019/

Publication

Accepted papers will be published in the AICS Workshop proceedings on arXiv after the event.

Workshop Co-Chairs

- William W. Streilein, MIT Lincoln Laboratory, MA, USA
- David R. Martinez, MIT Lincoln Laboratory, MA, USA
- Jason Matterer, MIT Lincoln Laboratory, MA, USA
- Howard Shrobe, MIT/CSAIL, MA, USA
- Arunesh Sinha, University of Michigan, MI, USA

Workshop Program Committee

- George Cybenko, Dartmouth College
- Robert Goldman, Smart Information Flow Technologies (SIFT)
- Gal Kaminka, Bar Ilan University, Israel
- Christopher Kiekintveld, University of Texas at El Paso
- Sven Krasser, Crowdstrike
- Robert Laddaga, Vanderbilt University
- Brett Meyer, Crowdstrike
- Ranjeev Mittu, Naval Research Laboratory
- Una-May O'Reilly, MIT/CSAIL
- Benjamin Rubinstein, University of Melbourne, Australia
- Salvatore Stolfo, Columbia University
- Milind Tambe, University of Southern California
- Robert Templeman, Navy Surface Warfare Center, Crane Division

Administrative Contact

Brent Cassella
MIT Lincoln Laboratory
Email: brent.cassella@ll.mit.edu
Voice: 781-981-7580
Fax: 781-981-8166

CROWDSTRIKE FOUNDATION

*Challenge Problem sponsored by Crowdstrike

Approved for public release. Distribution is unlimited. This material is based upon work supported by the United States Air Force under Air Force Contract No. FA8702-15-D-0001. Any opinions, findings, conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the United States Air Force.