UCF DEPARTMENT OF COMPUTER SCIENCE

Faculty Candidate Seminar

SECURE LEARNING IN ADVERSARIAL ENVIRONMENTS

Friday March 24, 2017 ● 12:00PM to 1:00PM ● Teaching Academy Room 117

Advances in machine learning have led to rapid and widespread deployment of software-based inference and decision making, resulting in various applications such as data analytics, autonomous systems, and security diagnostics. Current machine learning systems, however, assume that training and test data follow the same, or similar, distributions, and do not consider active adversaries manipulating either distribution. Recent work has demonstrated that motivated adversaries can circumvent anomaly detection or classification models at test time through evasion attacks, or can inject well-crafted malicious instances into training data to induce errors in classification through poisoning attacks. In addition, by undermining the integrity of learning systems, the privacy of users' data can also be compromised.

In this talk, I will describe my recent research addressing evasion attacks, poisoning attacks, and privacy problems for machine learning systems in adversarial environments. The key approach is to utilize game theoretic analysis and model the interactions between an intelligent adversary and a machine learning system as a Stackelberg game, allowing us to design robust learning strategies which explicitly account for an adversary's optimal response. I'll briefly discuss human subject experiments that support the results of mathematical models, and I will also introduce a real world malware detection system deployed based on adversarial machine learning analysis.

Dr. Bo Li

University of Michigan

Dr. Bo Li is a postdoctoral research fellow in the department of Electrical Engineering and Computer Science at University of Michigan, and is a recipient of the Symantec Research Labs Graduate Fellowship in 2015. Her research focuses on both theoretical and practical aspects of machine learning, security, privacy, game theory, social networks, and adversarial deep learning. She has designed several robust learning algorithms, a scalable framework for achieving robustness for a range of learning methods, and a privacy preserving data publishing system. She is also active in adversarial deep learning research for training generative adversarial networks (GAN) and designing robust deep neural networks against adversarial examples. Her website is http://www.crystal-boli.com/home.html

Hosted by: Faculty Cluster Initiative, Cyber Security and Privacy and Department of Computer Science

UCF

4328 Scorpius Street Room 346 Orlando, FL 32816

WWW.CS.UCF.EDU