DATA-DRIVEN SECURITY, PRIVACY, AND FAIRNESS

Special Topics in Computer Science: CS4390/CS5390

Instructor: Saeid Tizpaz Niari (saeid@utep.edu)

The objective of this course is to familiarize students with the state-of-the-art machine learning techniques including deep neural network (DNN), and the applications of these models to address challenges in cybersecurity such as finding bugs and debugging. In addition, the course will cover state-of-the-art techniques to enhance security, robustness, privacy, and fairness of machine learning models when used in adversarial and social-critical settings. Finally, the course highlights techniques beyond traditional machine learning such as causality.

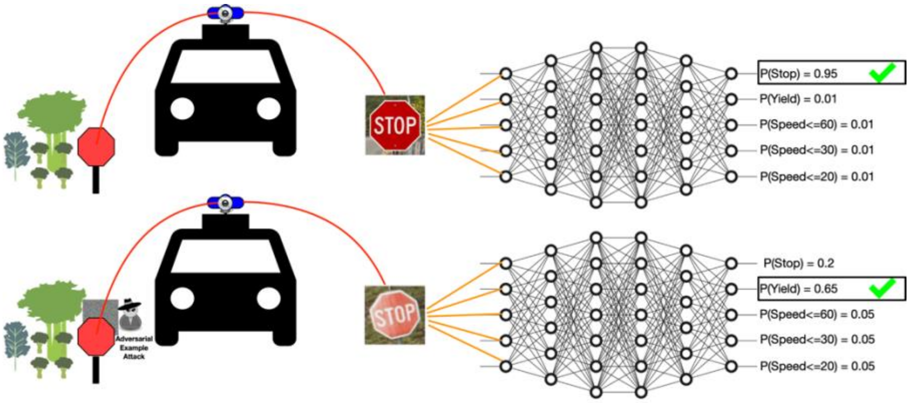

Course Topics Include:

- Overview of machine learning (ML): KNN, Linear Classification, Optimization, and Neural Networks.
- Attacks against ML: Adversarial Machine Learning and Data poisoning.
- Challenges in Defensing and Detecting Adversarial Example and Data Poisoning Attacks.
- Challenges and Opportunities in Deploying ML-based Software Systems.
- Data Privacy and Membership attacks.
- Differential Privacy.
- Fairness in Data-Driven Applications.
- Causality, Intervention, and Counterfactuals.
- Traditional and Modern Software Testing and Debugging.
- White-box and Gray-Box Testing for Machine Learning Systems.
- The Application of ML for Security Fuzzing and Software Testing.
- The Application of ML for Debugging and Fault Localizations.

Prerequisite:

This course requires no prior experience in security and privacy but assumes the willingness to seek out and read background material as needed. Although it is not a requirement, knowledge in core topics of machine learning and familiarity with Python and Numpy is a significant advantage.

Course Structure:

This is a research-oriented and discussion-based course, which also includes hands-on exercises and programming assignments. The students are required to write a review for assigned papers prior to the class so that they can participate in class discussions. Every student needs to present a major paper listed in course syllabus and lead discussions. Students will also work on a major project in group of 1, 2, or 3 and deliver write-ups, code, and presentations in phases.