Security Analytics

Spring 2019

1 Course description

The cyber threat landscape is diverse, and includes principal threats such as malware, botnets, spam, compromised accounts, fake accounts, and phishing, to name a few. These threats are constantly evolving and can negatively impact how we interact on the web, with other people, with our personal devices, and even threaten our safety at home (with the proliferation of IoT devices). How can we better understand such threats, and build protective measures? Given the diverse nature of these threats, where do we even start from? An effective approach is to start from *data* — most systems leave vast traces of data when they operate, e.g., logs of user activity, machine activity, and communication. In this class, we will explore how such *data* combined with appropriate *algorithms* can provide powerful tools to analyze security threats. We will start by covering the threat landscape from a data-driven perspective, by following research that takes a *measurement and analysis* approach to understand real world threats. This will help us understand incentives for attackers today, their attack strategies and how attacks evolve over time. Next, we will learn to apply techniques from machine learning, graph analysis, and natural language processing schemes in a security context. This includes understanding the strengths and limitations of different family of algorithms, and how certain combinations of data and algorithms may strengthen or weaken the "arms race" between attackers and defenders. Finally, we will cover the emerging space of *data-driven attacks*, where we consider malicious adversaries capable of leveraging data and machine learning (especially deep learning) to launch powerful attacks.

2 Topics covered

While it is hard to provide a solution to every security threat you may encounter, the topics below should equip you to use data and algorithms to better approach security problems.

- **Understanding common threats:** Measurement studies of various threats: e.g., malicious crowdsourcing services, large-scale botnets, fake news, spam, and phishing campaigns, reputation manipulation on social media, fake accounts, click fraud, and denial of service attacks.
- **Threats against machine learning systems:** Topics on adversarial machine learning, e.g., adversarial samples to fool classifiers, model poisoning attacks.
- **Data and algorithms for better security:** Application of machine learning (including deep learning), graph-based approaches, and NLP schemes to build robust defenses.
- **Data and algorithms for evil:** Misuse of deep learning for attacks, e.g., deep learning to bypass existing defenses, deep learning to generate fake online content to mislead users.

3 Reference materials

Most reading material will be drawn from research papers published at venues such as IEEE S&P, Usenix Security, CCS, NDSS, IMC, SIGCOMM, NSDI, CoNeXT, WWW, WSDM, and KDD. An

optional textbook for reference is "Machine Learning and Security: Protecting Systems with Data and Algorithms", by Clarence Chio, and David Freeman, O'Reilly Media, 2018.

4 Prerequisites

Prerequisites for the course include undergraduate courses on information systems, and high level programming languages. Students are expected to have a basic understanding of graph theory, algorithms, networks and distributed systems, and also be ready to learn concepts from machine learning, NLP and information retrieval. Knowledge of a scripting language such as Python or Perl would greatly aid you in your work. Students who enroll for the course are expected to be highly motivated to learn and work hard and be ready to make up for any prerequisite deficiencies they may have.

5 Grading

Final grade will be based on the following components:

- Paper summaries
- Paper presentation
- Homework assignments
- Project