PennState

Institute for Network and Security Research

Welcome: ARO Workshop on Adversarial Learning

Palo Alto, CA – September 14, 2017

Patrick McDaniel, Dan Boneh, Kamalika Chaudhuri, David Evans,
Somesh Jha, Dawn Song, Ananthram Swami

Welcome to the ARO Workshop!

PennState

- On behalf of the event organizers and our colleagues at the Army Research Office, we would like to welcome you to this workshop on adversarial machine learning ...

Thanks …

- We wish to
  - Cliff Wang of ARO, who has tirelessly supported this workshop
  - Ruth Harris of Stanford University who assisted Frankie
  - Judy Bowes of Penn State University, who has worked behind the scenes to help coordinate the proposal and initial setup
  - University staff, catering and other folks …
  - Student volunteers …
  - Speakers ..

Todays Agenda Highlights ...

Thursday, September 14th, 2017

9:15	Welcome and overview
9:30-10:15	Ian Goodfellow, Google
10:15-10:45	Jacob Steinhardt, Stanford
10:45-11:00	Break
11:00-11:30	Nicolas Papernot, Penn State
11:30-12:00	Aleksander Madry, MIT
12:00-12:30	Tian Pham, ARL
12:30-14:00	Lunch (provided, in room)
…	
16:30-16:35	Closing

Todays Agenda (cont.) ...

Thursday, September 14th, 2017

…

12:30-14:00 Lunch

14:00-15:00 Breakouts I

15:00-15:30 Breakouts II

15:30-16:00 Dawn Song, UC Berkeley

16:00-16:30 Dave Evans, Virginia

16:30-16:35 Closing

Breakouts (2-3pm)

- 6 different technical sessions (2 sessions of 3) spanning topics of interest collected from the attendees
- Rules
  - Pick and choose what interests you …
  - No presentations … (workshop rules)
  - Active participation ... it is on you
  - Be constructive, listen (self promotion should be limited)
  - Come and go as you like … try more than one if you want
  - Have fun, meet people, learn and get new ideas

- Q: What are the key challenges and research areas?

Breakouts

- 14:00-14:30
  - Privacy (Jha, Papernot)
  - Measuring and achieving resilience (Evans, Grosse)
  - Fairness (Chaudhuri, ???)
- 14:30-15:00
  - Autonomous cyber defense (Swami, Grosse)
  - Adversarial reinforcement learning (Wellman, …)
  - Attacks on training attacks (Boneh, …)

Machine Learning

- Perhaps no area of computer science has had more impact on systems and society in the last 5 years than machine learning
  - Analytics
  - Autonomous systems
  - Vision …

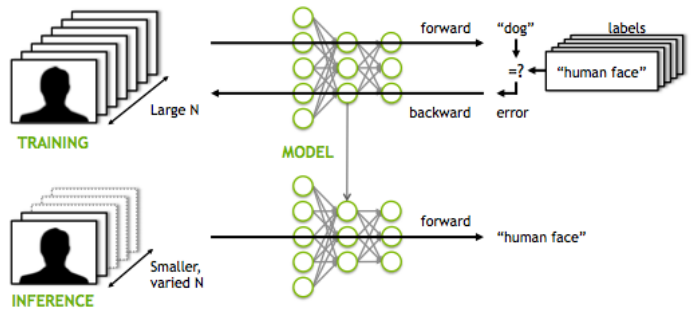

Focus: Security and Privacy in ML

PennState

- Challenge: what are the security and privacy challenges of the use of machine learning in adversarial settings?
  - Fundamental science:
    - What are the limits of machine learning with respect to accuracy and resilience?
    - What vulnerabilities are general vs. those are a consequence of the techniques used?
    - Can the advantages of ML be realized while preserving privacy?
  - Applied science
    - What countermeasures are likely to be effective in practice?
    - What are the domain specific challenges and safeguards for security and privacy?
    - Ethics:
      - Just because a system may be able to understand environment, should it?
      - Can the advantages of ML be realized fairly without discriminating minorities?
    - Education (what and how to integrate security into machine learning/security courses)

PennState

Institute for Network and Security Research

Welcome: ARO Workshop on Adversarial Learning (Closing)

Palo Alto, CA – September 14, 2017

Patrick McDaniel, Dan Boneh, Kamalika Chaudhuri, David Evans, Somesh Jha, Dawn Song, Ananthram Swami

Follow on …

- Readout: over the next few weeks the organizers will try to capture the substance of the conversations with a focus on the breakouts.
- Make a short statement of key areas and challenges faced by the community, as well as identify important application domains.
- We will send an email with a reference to the readout to all participants shortly thereafter.

- Please send any comments, questions to mcdaniel@cse.psu.edu

Thanks …

- We wish to
  - Cliff Wang of ARO, who has tirelessly supported this workshop
  - Ruth Harris of Stanford University who assisted Frankie
  - Judy Bowes of Penn State University, who has worked behind the scenes to help coordinate the proposal and initial setup
  - University staff, catering and other folks …
  - Student volunteers …
  - Speakers ..

- You!