CS 6604: Applied Machine Learning in Security

Gang Wang

Fall 2017

VirginiaTech

Invent the Future

About This Class

- Tuesday/Thursday 3:30 PM - 4:45 PM in McBryde Hall 226

- Instructor
  - **Gang Wang**, Assistant Professor of Computer Science
  - Office hour: after class, by appointment (gangwang@vt.edu)
  - Office location: 2202 Kraft Drive, Knowledge Works II, RM 2223

- If any student needs special accommodations because of a disability, please contact me in the first week of classes

No Textbook

- Recent papers from security, networking, and ML conferences

Computer Security

USENIX Security, CCS, SP, NDSS

Machine Learning, Data Mining

ICML, NIPS, CVPR, ICCV, KDD

Networking

SIGCOMM, IMC

Class Website

- Course site: http://people.cs.vt.edu/~gangwang/class/cs6604/index.html

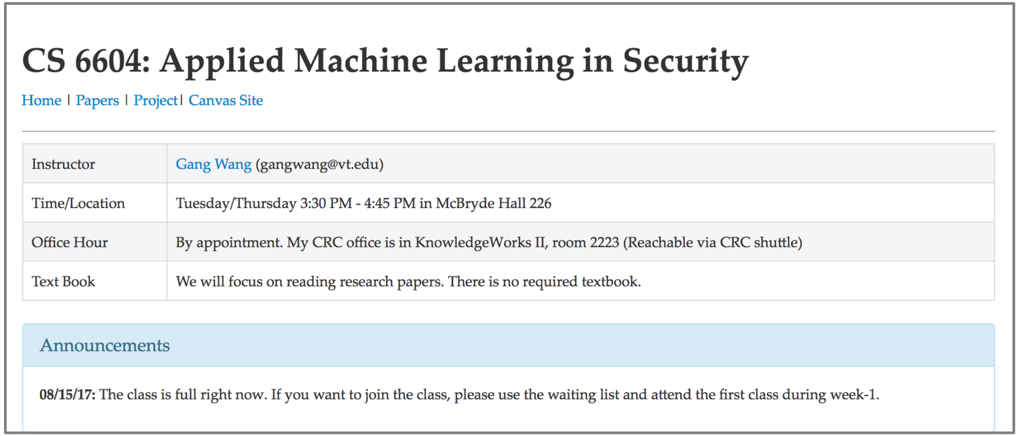

CS 6604: Applied Machine Learning in Security

Home | Papers | Project| Canvas Site

Instructor	Gang Wang (gangwang@vt.edu)
Time/Location	Tuesday/Thursday 3:30 PM - 4:45 PM in McBryde Hall 226
Office Hour	By appointment. My CRC office is in KnowledgeWorks II, room 2223 (Reachable via CRC shuttle)
Text Book	We will focus on reading research papers. There is no required textbook.

Announcements

08/15/17: The class is full right now. If you want to join the class, please use the waiting list and attend the first class during week-1.

- Use Canvas to post paper discussions and submit your project report
  - https://canvas.vt.edu/courses/57607/

High-Level Topics

- Apply Machine Learning for Security Attacks
  - De-anonymization, information leakage, side-channel attacks …
- Apply Machine Learning to Build Security Solutions
  - Phishing detection, malware detection, intrusion detection, authentication …
- Adversarial Attacks Against Machine Learning Models
  - Evasion attack, poisoning attack …
- Improving the Security of Machine Learning Models
  - Poisoning detection, concept drifting detection …

Expected Work

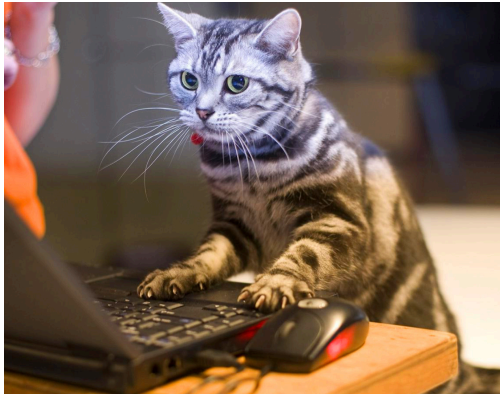

- **Attend all the lectures**
  - Let me know ahead of time if you cannot attend
- **Paper reading**
  - Post your comments online **before** each class
  - Be original, don’t repeat what the authors/your classmates said
  - 200-500 words
- **Paper presentation in the class**
  - Build your own slides, don’t use the authors’ slides stack (30-40 mins)
  - Discuss what you lean, the good and the bad about the paper, how to improve
  - Start your talk with **“attacks and defense of the week”** (5-10 mins)

Expected Work (Cont.)

- No Exam

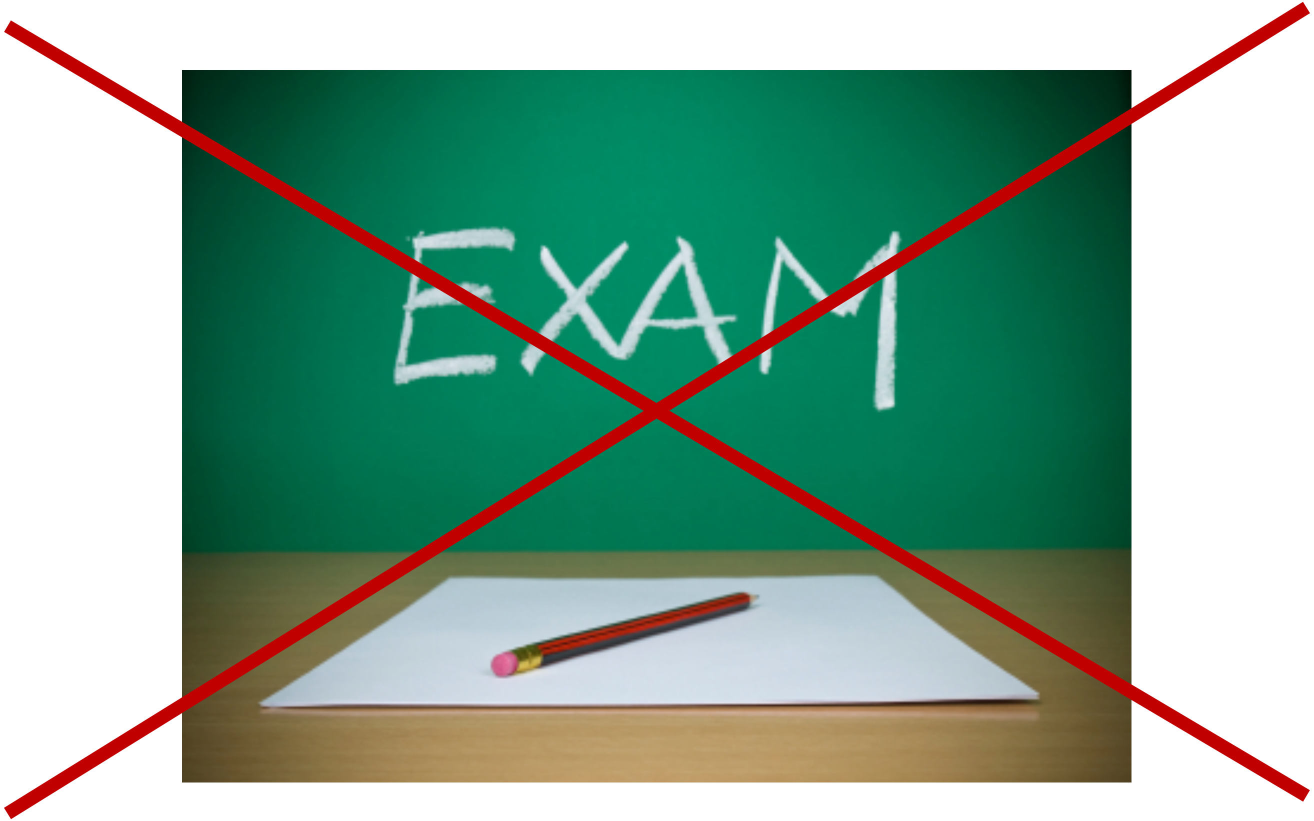

Expected Work: Class Project

- **Proposal**
  - 1-2 pages, describe the problem, background, plan and timeline
  - Publishable idea, talk to me before starting writing
  - Submission (in 2 weeks) → Feedback → Final version (3rd week)
- **Midterm presentation**
  - The research problem, your progress, and next steps
- **Final presentation**
  - Presenting your idea, your approaches, key findings and results
- **Project report**
  - 8-page workshop style paper (excluding references and appendixes)

Grading

- Sum up the points: x out of 100
- Convert x to letter grade:
  - [0-60] F, [60-62] D-, [63-66] D, [67-69] D+, [70-72] C-, [73-76] C, [77-79] C+, [80-82] B-, [83-86] B, [87-89] B+, [90-92] A-, [93-100] A.
- I do not curve the grades

Class attendance and participation	10%
Paper discussion online	20%
Paper presentation in class	15%
Project: proposal	10%
Project: midterm presentation	10%
Project: final presentation	20%

Policies

- **Late policy**
  - All the deadlines are hard deadlines
  - Late submission score = 0.5*(your raw score), if delay < 72 hours (3 days)
  - Late submission score = 0, if delay >= 72 hours (3 days)

- **Virginia Tech honor code**
  - Make proper citations for other people's ideas, tools, code, datasets
  - Write your own slides and project reports

A big of background

MACHINE LEARNING IN SECURITY

What is Machine Learning?

The complexity in traditional computer programming is in the code (programs that people write). In machine learning, algorithms (programs) are in principle simple and the complexity (structure) is in the data. Is there a way that we can automatically learn that structure? That is what is at the heart of machine learning.

-- Andrew Ng

Machine learning is the about the construction of systems that can learn from data.

What is Machine Learning?

Traditional Programming

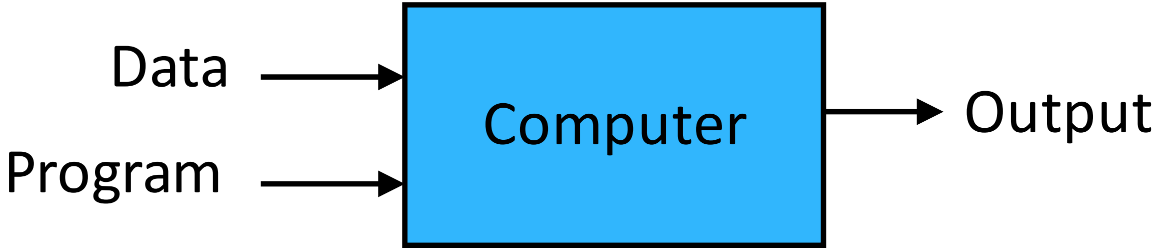

Machine Learning

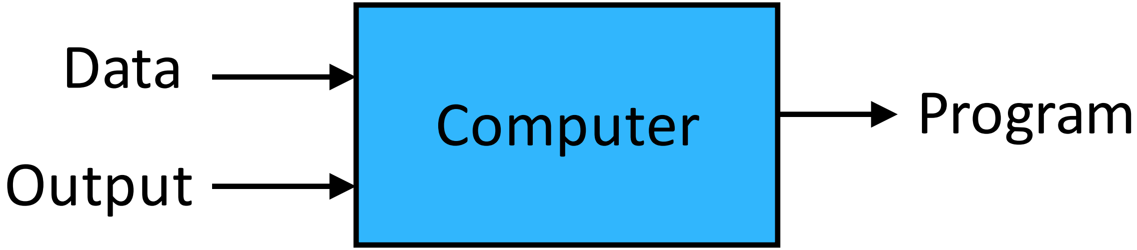

Slides credit to David Meyer

When Would We Use Machine Learning?

- **When patterns exists in our data**
  - Even if we don't know what they are

- **We can not obtain the functional relationships mathematically**
  - Else we would just code up the algorithm

- **When we have lots of (unlabeled) data**
  - Labeled training sets harder to come by
  - Data is of high-dimension
  - Want to discover lower-dimension representations

Why Machine Learning Is Hard

What is a "2"?

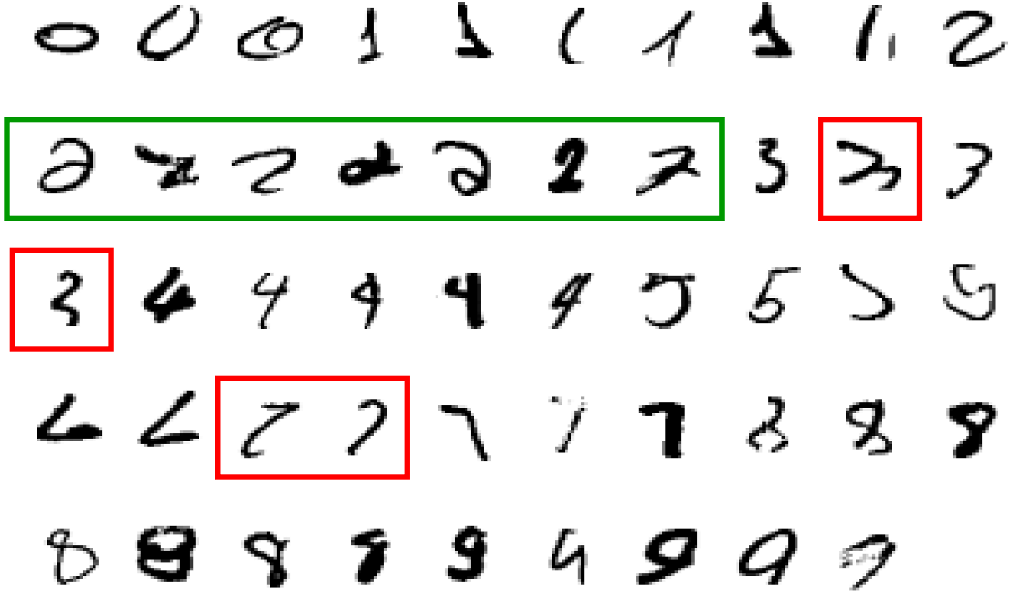

Examples of Machine Learning Problems

- Pattern recognition
  - Facial identities or facial expressions
  - Handwritten or spoken words (e.g., Siri)
  - Medical images
  - Network traffic

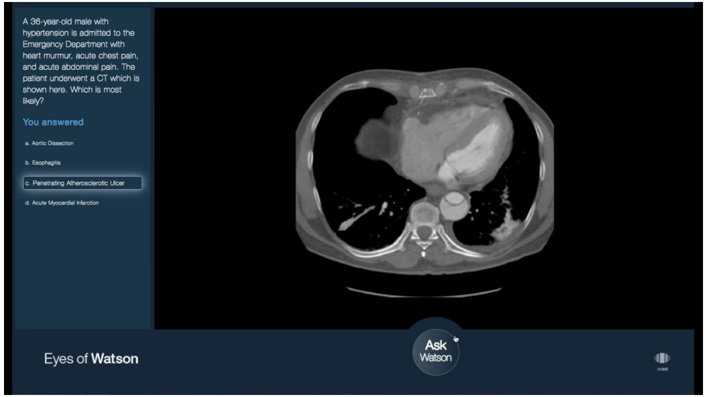

A 36-year-old male with hypertension is admitted to the Emergency Department with heart murmur, acute chest pain, and acute abdominal pain. The patient underwent a CT which is shown here. Which is most likely?

You answered

a. Aortic Dissection

b. Esophagitis

[✓ c. Penetrating Atherosclerotic Ulcer]

d. Acute Myocardial Infarction

- Pattern generation
  - Generating images
  - Motion sequences

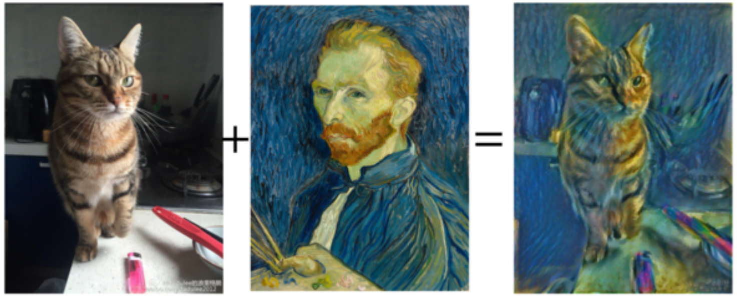

Examples of Machine Learning Problems

- Anomaly detection
  - Unusual patterns in data center networks
  - Unusual sequences of credit card transactions
  - Unusual patterns of sensor data from a nuclear power plant

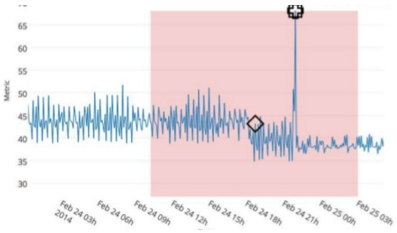

- Prediction
  - Future stock prices or currency exchange rates
  - Network events