Course Introduction

About me

Rensselaer

- My name is Radoslav Ivanov
  - Call me Rado
- Undergrad degree in CS and ECON from Colgate in 2011
- Got my PhD in CS from UPenn in 2017
- My research is on safe and secure autonomous systems
  - Verification of neural networks
  - Attack-resilient sensor fusion
  - Context-aware detection and estimation
- Started at RPI in Jan. 2022

Impressive Progress in Autonomy

Rensselaer

Control | Perception | Learning

Boston Dynamics

YOLO v. 3

Kormushev, Calinon, Caldwell, IROS'10

JPL-Caltech, DARPA Robotics Challenge

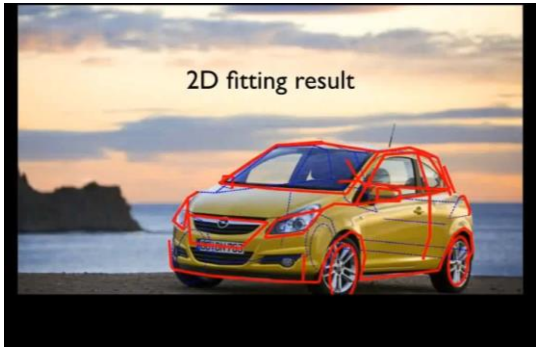

Zhu, Zhou, Daniilidis, ICCV'15

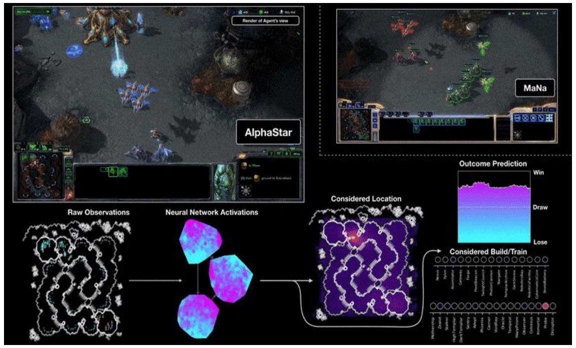

DeepMind

But we're not there yet…

Rensselaer

Tesla Driver Was on Autopilot Eating a Bagel When He Smashed into a Fire Truck

A National Transportation Safety Board investigation found the driver had hands off the wheel and ignored warnings in the 2018 crash.

By Clifford Atiyeh SEP 4, 2019

Culver City Firefighters via Instagram

- The driver of a 2014 Tesla Model S that ran into the back of a fire engine in California in 2018 was using Autopilot at the time, according to a National Transportation Safety Board (NTSB) report this week.
- The agency's investigators reported that the driver was having breakfast while he let Autopilot take over the driving; his hands were not on the steering wheel, and he did not brake prior to the crash.

Uber self-driving car involved in fatal crash couldn't detect jaywalkers

The system had several serious software flaws, the NTSB said.

Steve Dent, @stevetdent 11.06.19 in Transportation — 32 Comments — 2136 Shares

ум: Australian open 2020 - | Дискусия на живо: Аустрелийн | Live Scores | Australian Open

320936/waymo-self-driving-car-crash-arizona

GOOGLE \ TECH \ TRANSPORTATION

Waymo self-driving minivan involved in crash in Arizona

Minor injuries reported

By Andrew J. Hawkins | @andyjayhawk | May 4, 2018, 4:52pm EDT

SHARE

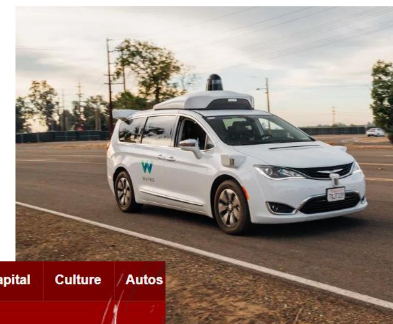

Boeing 737 Max Lion Air crash caused by series of failures

25 October 2019

Share

Lion Air plane crash

BBC NEWS TECHNOLOGY

News | Sport | Weather | Capital | Culture | Autos

Home | US & Canada | Latin America | UK | Africa | Asia | Europe | Mid-East | Business | Health | Sci/Environment

25 July 2013 Last updated at 19:04 ET

Share

Car hackers use laptop to control standard car

By Zoe Kleinman
Technology reporter, BBC News

The Washington Post | PostTV | Politics | Opinions | Local | Sports | Natio

National Security

Home > Collections > Surveillance

Iran says it downed U.S. stealth drone; I acknowledges aircraft downing

By Greg Jaffe and Thomas Erdbrink, December 04, 2011

A secret U.S. surveillance drone that went missing last week in western Afgha have crashed in Iran, in what may be the first case of such an aircraft ending u an adversary.

Iran's news agencies asserted that the nation's defense forces brought down t Iranian reports said was an RQ-170 stealth aircraft. It is designed to penetrate defenses that could see and possibly shoot down less-sophisticated Predator a

A stealthy RQ-170 drone played a critical role in surveilling the compound in Osama bin Laden was hiding in the months before the raid in which he was ki SEALs in May.

U.S. officials acknowledged Sunday that a drone had been lost near the Irania declined to say what kind of aircraft was missing.

The Iranian government has not released any pictures of the recovered aircra was downed by defense forces after it flew across the border and into the cou unnamed Iranian defense official said in one report that a cyberattack caused

U.S. officials cast doubt on th Iranian assertions. "We have no indication that it was brought down by hostil senior Pentagon official, speaking on the condition of anonymity to discuss se activity.

RESEARCHERS HACK GPS, $80M YACHT VEERS OFF COURSE

Brian Donohue | Follow @TheBrianDonohue | July 30, 2013 , 3:26 pm

A 213-foot luxury yacht veered off course while cruising in the Mediterranean Sea this

Neural Network (NN) Vulnerabilities

- Neural networks increasingly used in safety-critical systems

Perception (autonomous cars)

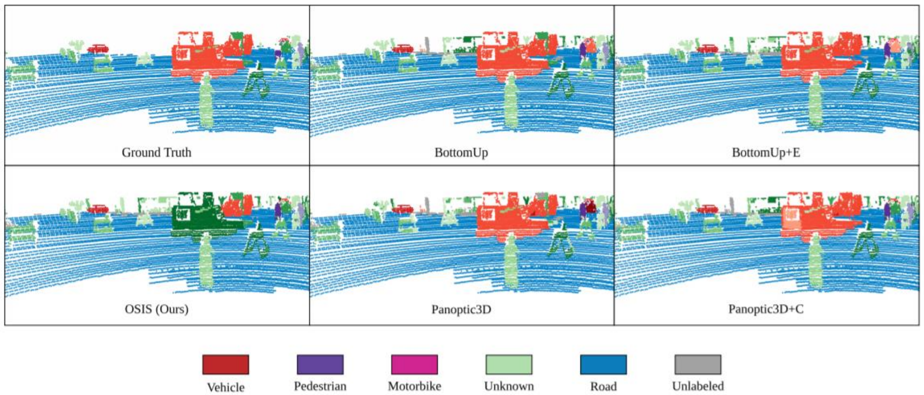

Wong et al., CoRL'19

Control (air traffic avoidance)

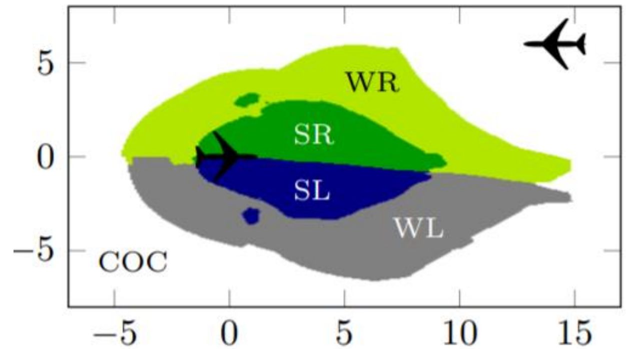

Katz et al., CAV '17

- Safety concerns discovered in both domains

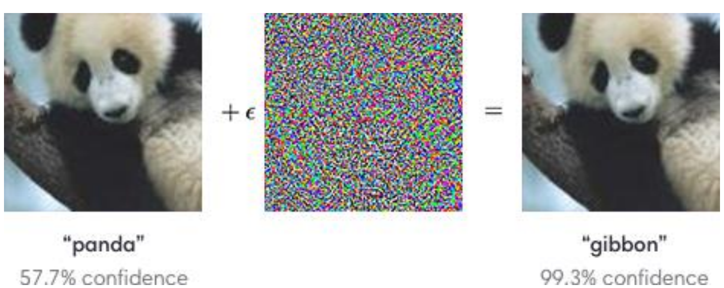

Goofellow et al., ICRL'15

Table 2: Verifying properties of the ACAS Xu networks.

	Networks	Result	Time	Stack	Splits
ϕ_1	41	UNSAT	394517	47	1522384
	4	TIMEOUT			
ϕ_2	1	UNSAT	463	55	88388
	35	SAT	82419	44	284515
ϕ_3	42	UNSAT	28156	22	52080
ϕ_4	42	UNSAT	12475	21	23940

Cyber-Physical Systems (CPS)

Rensselaer

Tight coupling between **communication**, **computation** and interaction with the **physical world**

Aircraft

Autonomous Cars

Medical CPS

Military

Smart Grids

Robotics

A standard CPS design

F1/10 Autonomous Racing Competition, ES Week 2016

Problem: How do we know car won't crash?

- How do we build safe algorithms?
- How do we analyze algorithms?
- What about "black-box" components such as neural networks?
- How do we convince other people car is safe (assurance argument)?

CPS Autonomy: Problem Landscape and Complexity

Rensselaer

Information Processing and Acquisition

- Perception, prediction
- Active info acquisition
- State estimation w.r.t. the environment

Interaction with the Environment

- Human machine interaction
- Privacy, trust
- Distributed control and computation

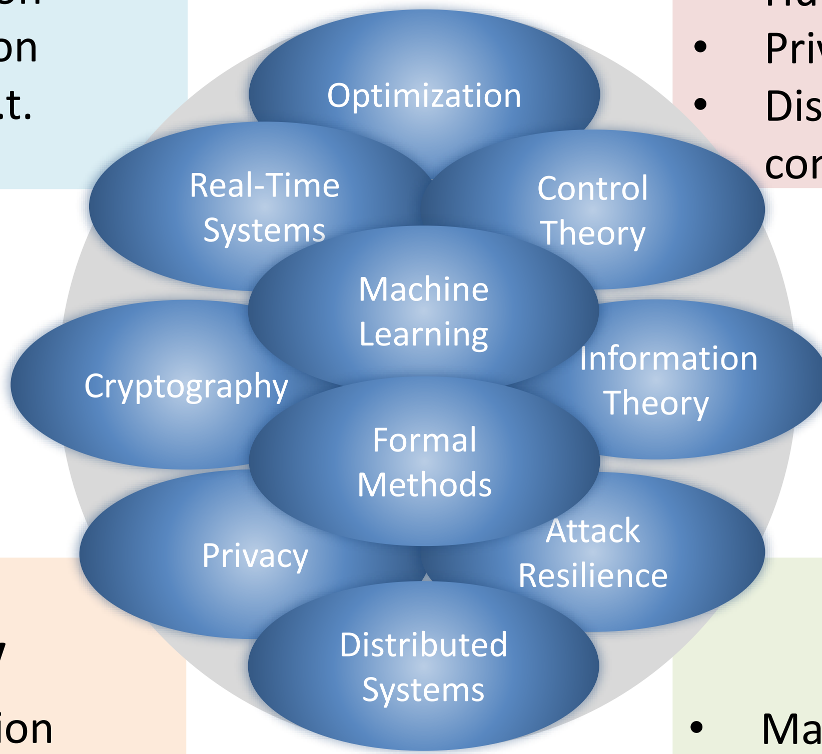

System Security

- Secure communication
- Secure computation
- Attack detection, recovery

System Safety

- Machine learning verification
- Anomaly detection, recovery
- Assurance cases

Why is safe autonomy so hard?

Even navigating hallways is not easy!

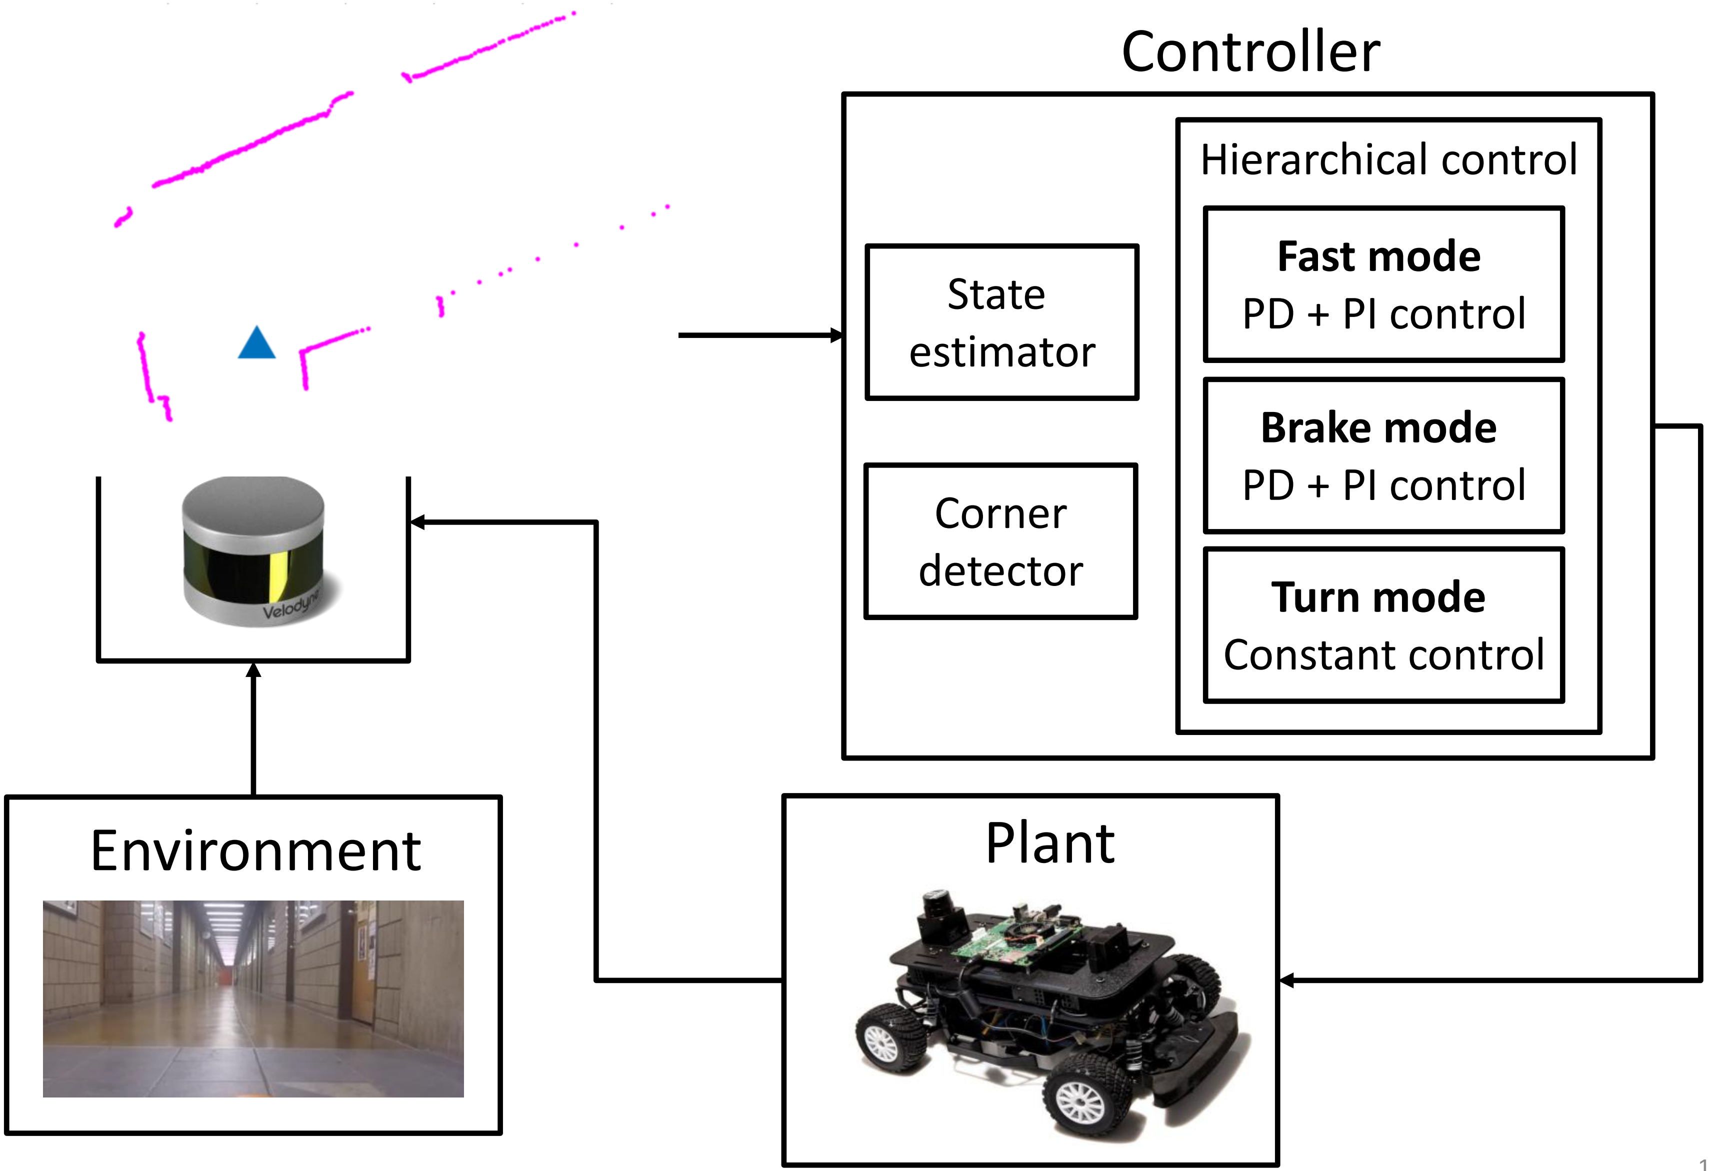

Building blocks of autonomous systems

- Model measurements
- Often high-dimensional
- Model is certainly inaccurate!

- Perception – recognize objects
- Control/planning – decide on actions

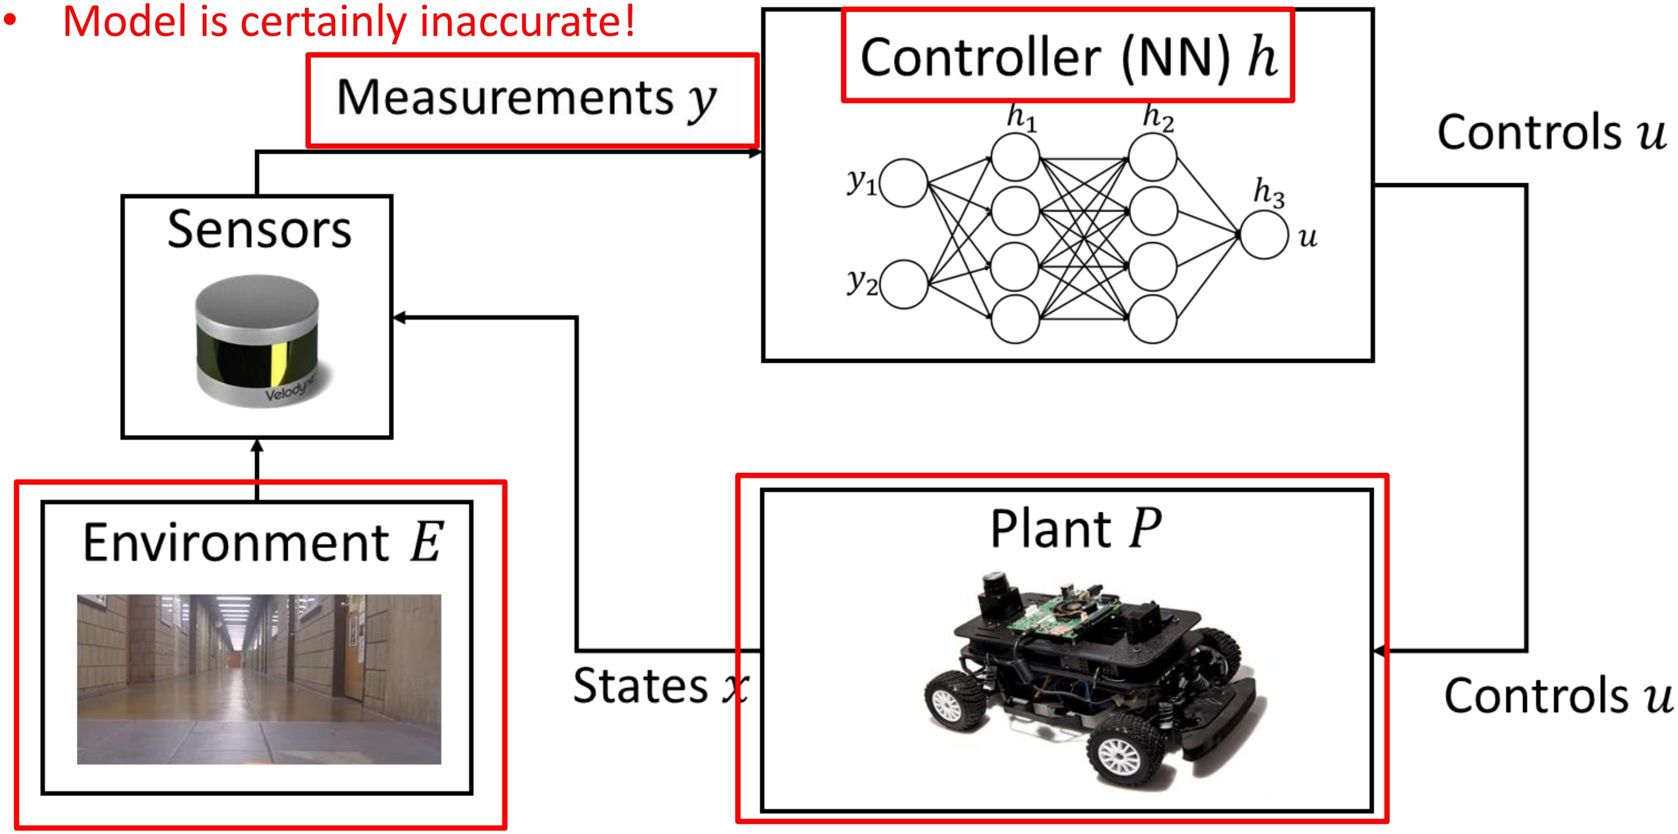

- Model environment agents
- Perceive agents, predict movement
- Essentially the singularity problem!

- Model plant dynamics
- Analyze dynamics properties
- What if model is inaccurate?

Topics

1. Supervised machine learning
   – Linear regression and classification
   – Generalization
   – Deep learning
2. Reinforcement learning
   – Modeling, relation to standard control theory
   – Markov chains, Markov reward/decision processes
   – Policy/Value iteration
   – Q learning
   – Policy gradients
   – Actor-critic methods

Course Mechanics

- Meeting time: TF 10-noon
  - Each lecture will be split up into two 50-minute sessions, with a 10-minute break in between
- We will meet in Sage 3101
- Office hours: T 1-2pm, W 2-3pm, Th 11am-noon
  - Lally 309
  - Office hours will be in person unless noted otherwise

Course Mechanics, cont'd

- TA: Thomas Waite
  - PhD student
  - Email: waitet@rpi.edu
- Mentor: Anthony Shaw
  - Email: shawa9@rpi.edu
- TA/Mentors will be monitoring Piazza and will be helping with grading/marking

Course Mechanics: Piazza

- We will be using Piazza for questions and discussions
- Sign-up link: https://piazza.com/rpi/fall2024/csci41606963ecse49656965
- Access code:
- Please let me know if you're having issues enrolling

Course Mechanics: LMS

- All lecture notes and slides will be posted on the website
  - http://cs.rpi.edu/~ivanor/rl/F24/rl.html
- Homework assignments and submissions will be through LMS
- Please use Piazza for questions and discussion
  - I won’t monitor LMS/Webex that frequently

Course Mechanics, cont'd

- Lectures will be a mix of theory and practice
  - RL is inherently a statistical subject, will cover the basics of statistical learning and probability theory
- Homeworks will also be a mix
  - A few problem sets and a few programming assignments
  - Submit through LMS
  - Please make sure you have access now
  - There will be 10 homeworks total

Course Mechanics, cont'd

- Some homeworks will require significant computation
  - One big deep learning assignment
    - Classify buildings on campus
    - Charles Yu '23 and I have collected a dataset of about ~500 images per building in different weather conditions/time of day
  - One deep reinforcement learning assignment
- We will use CCI for these assignments
  - RPI/IBM's computing cluster
  - You will need a basic understanding of how to use a Unix command line and possibly use an editor over it
  - I also recommend using Ubuntu for the other assignments
    - Deep learning libraries mostly developed for Linux systems

Grading

- Homework (100%)

- Please attend the lectures unless you have a good reason not to
  - Won’t take attendance but participating in class helps you learn and helps me teach
  - It will be very hard to complete some assignments if you miss the lectures

Reading: standard ML

- Hastie, Trevor, et al. The elements of statistical learning: data mining, inference, and prediction. Vol. 2. New York: springer, 2009.
  – A very comprehensive book – we will cover some parts only
  – Available online: https://hastie.su.domains/Papers/ESLII.pdf
- James, Gareth, et al. An introduction to statistical learning. Vol. 112. New York: springer, 2013.
  – An introductory version of the above
  – Available online: https://www.statlearning.com/
- Ian Goodfellow, Yoshua Bengio, and Aaron Courville. Deep learning. MIT press, 2016.
  – Introduction to deep learning
  – Available online: https://www.deeplearningbook.org/

Reading: RL

- The RL portion will follow this book
  - Sutton, Richard S., and Andrew G. Barto. Reinforcement learning: An introduction. MIT press, 2018.
  - Available online: http://incompleteideas.net/book/RLbook2020.pdf
  - Good high-level overview of RL
- We will also cover parts of this very comprehensive book on Markov Decision Processes
  - Puterman, Martin L. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, 2014.
  - Physical copy available in the library
  - Solid theoretical introduction to MDPs

More resources

- Many ML texts out there
  - Many views on which topics need to be covered
  - Some good books are:
    - Mohri, Mehryar, Afshin Rostamizadeh, and Ameet Talwalkar. *Foundations of machine learning*. MIT press, 2018.
    - Kearns, Michael J., and Umesh Vazirani. *An introduction to computational learning theory*. MIT press, 1994.
    - Bishop, Christopher M., and Nasser M. Nasrabadi. *Pattern recognition and machine learning*. Vol. 4. No. 4. New York: springer, 2006.
- Not as many RL resources
- Useful lecture notes by David Silver:
  - https://www.davidsilver.uk/teaching/

Course Requirements

- Modern machine learning makes heavy use of linear algebra and probability theory
- Although this is not a theory course, there will be assignments with problem sets
  - It helps if you have some formal background, e.g., FOCS, algorithms, calculus, analysis
  - We will cover some of the basics but we can't cover all the necessary background material
- We will be using Python for programming assignments
  - If you have never used Python, this course will be very difficult for you
- Talk to me if you're not sure if this is the right course for you

Course Difficulty

- "The course got significantly harder after the drop date, which is not cool"
- Keep in mind that assignments will get harder, especially as we get deeper into RL
- RL is an advanced ML topic
  - It's a combination of control theory, dynamical systems and ML
  - ML is already an advanced topic, a combination of statistics and optimization

Meet and greet

- Introduce yourself
  - What year are you (undergraduate/graduate)?
  - What’s your major/research interest?
  - Why are you taking this course?
  - One fun fact about you