11-695: AI Engineering
Introduction

LTI/SCS

Spring 2020

Logistics

Carnegie Mellon

- Course staff:
  - Instructor: Hai Pham, htpham@andrew.cmu.edu
  - TAs:
    - Jianlin Du, jianlind@andrew.cmu.edu
    - Yiyang Ge, yiyangge@andrew.cmu.edu
- Class website: `http://www.cs.cmu.edu/~11695-s20`
  - Syllabus
  - Lecture slides
  - Assignments
- Piazza: `http://piazza.com/cmu/spring2020/11695`
  - Announcements
  - Additional Materials/References
  - Discussions

What is this class about?

Deep Learning.

What is this class about?

Carnegie Mellon

Deep Learning.

And how to implement Deep Learning.

What is this class about?

Carnegie Mellon

Deep Learning.

And how to implement Deep Learning.

(in Tensorflow/Pytorch, for some essential models)

What is this class about?

Carnegie Mellon

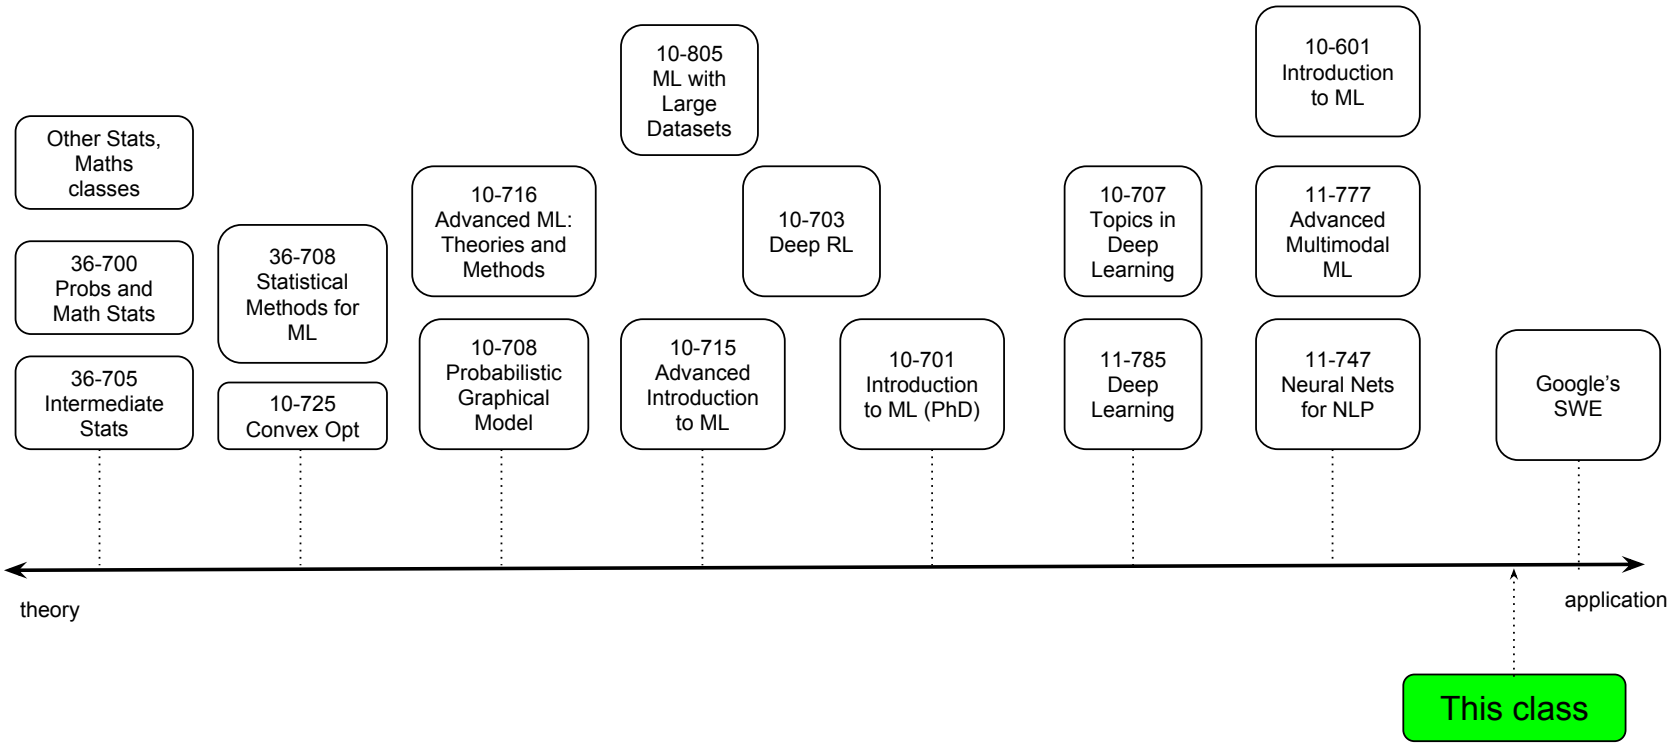

What is this class about?

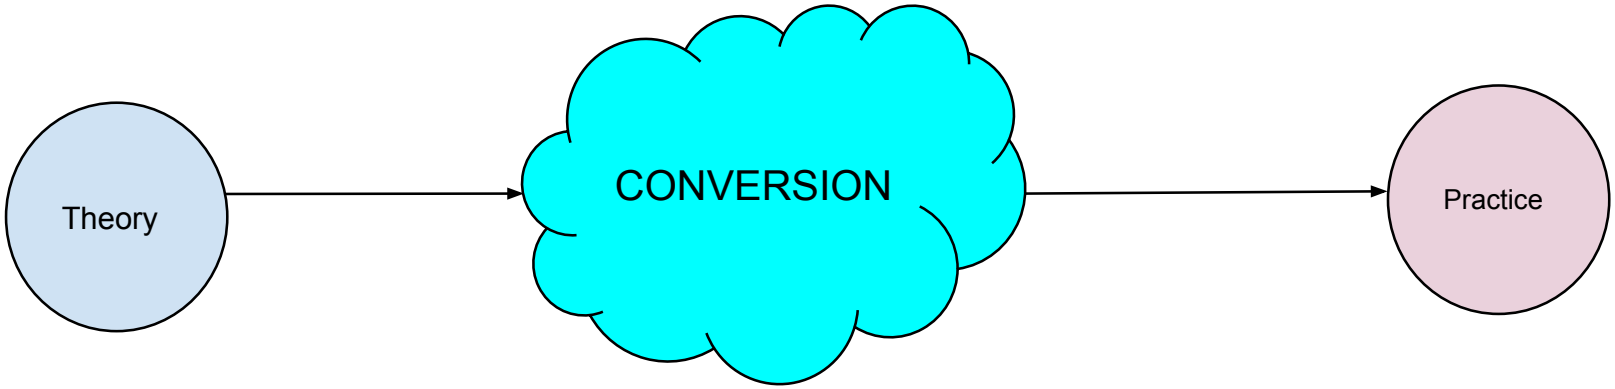

What is this class about?

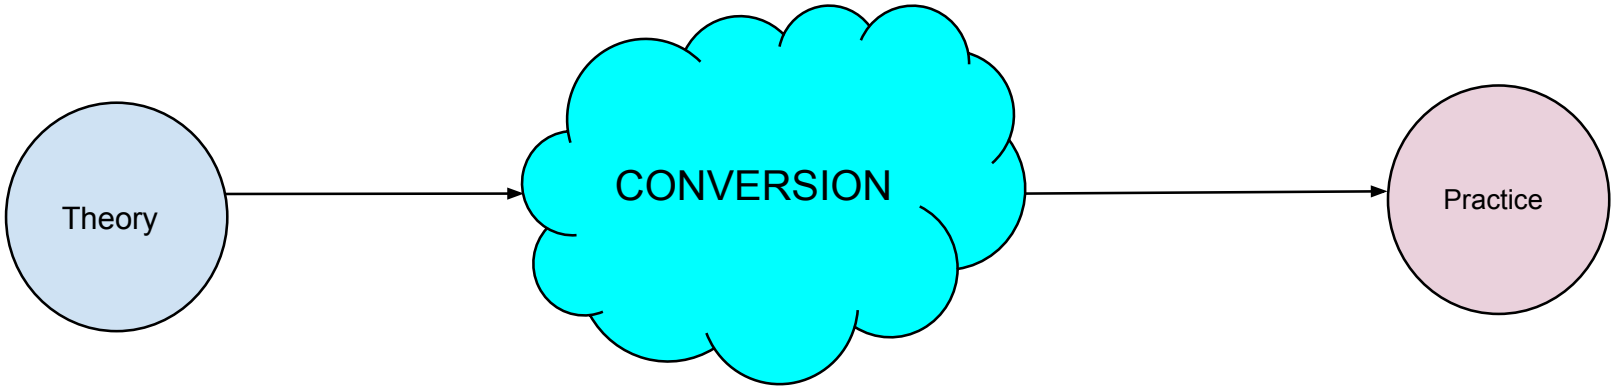

- Focus on **implementation**
  - With the help of some intuitions
  - And some best practices
- Boost your understanding of **(essential) deep learning models and algorithms**

Prerequisites

Carnegie Mellon

- Proficiency in Python

Prerequisites

Carnegie Mellon

- Proficiency in Python
- Statistics, Optimization, Linear Algebra, Calculus

Prerequisites

- Proficiency in Python
- Statistics, Optimization, Linear Algebra, Calculus
- But we will help review fundamental knowledges if required

Toolkits: Tensorflow/Pytorch

Carnegie Mellon

- Leading deep learning toolkits
  - Strong GPU (Tensorflow: +TPU) support, seamless distributed computing, ...
  - Active online community, from academia to industry
- They are most popular toolkits
  - But there are also other options: DyNet, MxNet, ...
  - In theory, any of such frameworks can help you implement your solutions

Roadmap: Fundamentals

Carnegie Mellon

- Supervised Learning Fundamentals
  - Regression, Classification
  - Optimization, Loss
  - Training and Testing, Overfitting, Regularization
- Neural Networks Fundamentals
  - Feed-forward NN
  - Convolutional NN
  - Recurrent NN
- Distributed Applications
  - Distributed Optimization
  - Data-Parallelism

Roadmap: Advanced Topics

Carnegie Mellon

- Unsupervised Learning and Generative Models
  - AutoEncoder (AE), Denoising AutoEncoder
  - Variational AutoEncoder (VAE)
  - Generative Adversarial Networks (GAN)
  - Normalizing Flow Models
  - Style Transfer and Conditional Generation
- Reinforcement Learning and Related Topics
  - MDP, POMDP and Model-based RL
  - Model-free RL
  - Transfer Learning and Meta Learning
- Some other potential topics: Other Probabilistic Models, Applications, Guest Lectures

Gradings

- All Quizzes and Individual Coding Assignments
  - Quizzes: *every* single lecture 25%
  - 3-5 Coding Assignments: 75%
- Grace days: 5
  - You can use all at once, or
  - You can distribute however you want
  - Excuse: with *valid* evidence
- No late submission is accepted. Likewise, Zero is given.

Gradings

Carnegie Mellon

- All Quizzes and Individual Coding Assignments
  - Quizzes: *every* single lecture 25%
  - 3-5 Coding Assignments: 75%
- Grace days: 5
  - You can use all at once, or
  - You can distribute however you want
  - Excuse: with *valid* evidence
- No late submission is accepted. Likewise, Zero is given.
- Please start early! Don't ever underestimate the debugging requirement for a deep learning model (even the easiest one).

Weekly Quizzes

Carnegie Mellon

- Test knowledge *from the most recent lecture*
  - Multiple choices (mostly); Fill in the blank; Short answers.
  - 10-15 minutes, at the end of class.
- Sample quizzes:

1. What is the Regularization effect towards bias?
   A. Increase B. Decrease C. No effect.
2. In a feed-forward neural network, layer 5 has 128 units, and layer 6 has 256 units. What is the size of $\mathbf{W}^{5,6}$? (assuming no bias included)
   A. 5*128 B.6*256 C.128*256 D. 256*256
3. If bias is included in question 2., what is the answer instead?
4. Find $f'(x)$ where $f(x) = x \cdot \text{sigmoid}(\beta x)$.

Programming Assignments

Carnegie Mellon

- Python: `NumPy` and `TensorFlow/Pytorch`
- Download starter code, download datasets and submit assignment via Canvas.
- Assignments are **not** hard, but they require reasonable effort:
  - We believe it's fun and helpful for honing your skills in Deep Learning
  - You will have enough time if you **start early**
  - You will be provided enough resources
- Make use of Piazza discussions, you will learn a lot from them.

Policy on Cheating

Carnegie Mellon

- We will be around to answer your questions during your quizzes
  - So please keep yourselves honorable
  - If you get caught, things are out of your control.
- Plagiarism check is run *for each assignment*
  - Violations of the Academic Integrity Policy are taken very seriously
  - Instructors MUST report any violations as provided by policy