Introduction to Deep Learning

0. Logistics

Fall 2024

Outline

- Introduction
- Who should take this Course?
- Objective and syllabus
- Course logistics
- Homeworks, quizzes, projects, grading, oh my!
- Preps, teamwork and mentoring
  - And cheating…
- Challenges

Outline

- Introduction
- Who Should take the Course?
- Objective and syllabus
- Course logistics
- Homeworks, quizzes, projects, grading, oh my!
- Preps, teamwork and mentoring
  - And cheating…
- Challenges

Neural Networks are taking over!

- What are neural networks in simple words: A neural network is a method in artificial intelligence that teaches computers to **process data in a way that is inspired by the human brain**

- Neural Networks have become one of the major thrust areas recently in various pattern recognition, prediction, and analysis problems

- Established state-of-the-art in many problems, often exceeding the previous benchmarks by large margins

Breakthrough with Neural Networks

- Audio and speech recognition: Siri, Alexa, Google Home, Cortana




Breakthrough with Neural Networks

- **Vision/ Perception**
  - Face Detection, Face Recognition, Object Detection, Semantic Segmentation, Pose Estimation, Activity Recognition, Emotion Detection, Biometrics, and Forensics.

https://github.com/facebookresearch/detectron2

Breakthrough with Neural Networks

- Self-driving cars

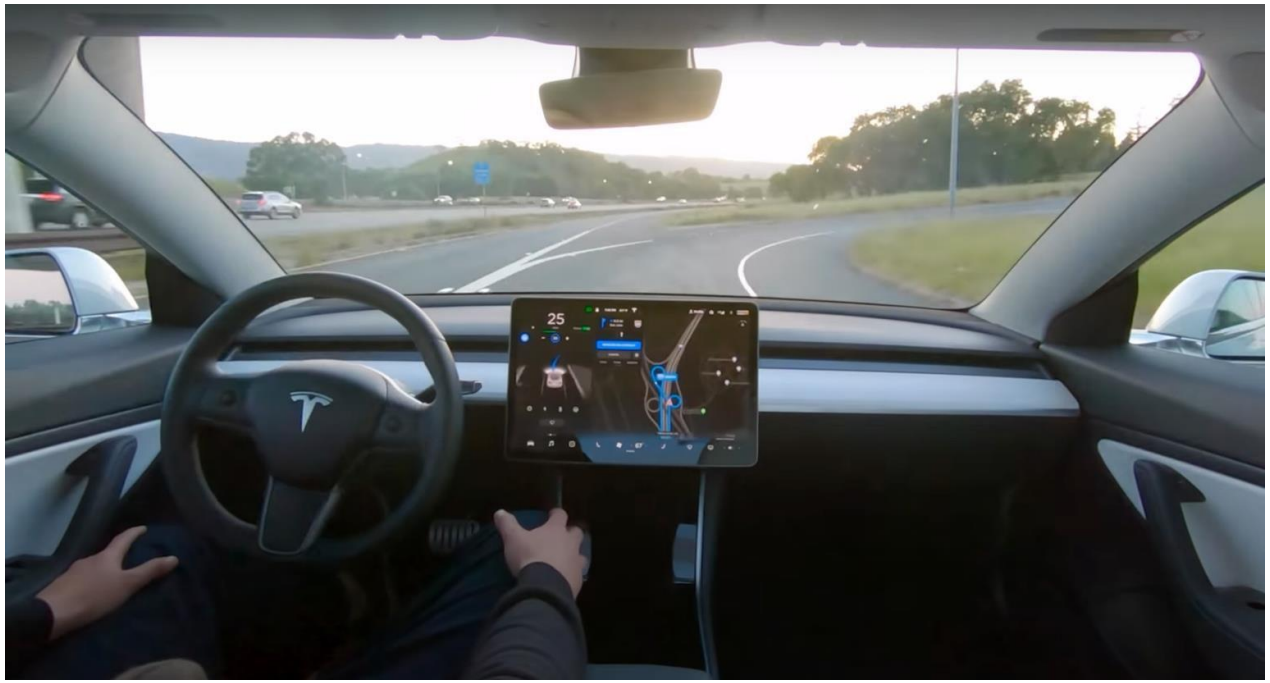

https://youtu.be/tlThdr3O5Qo

Breakthrough with Neural Networks

- Music Generation: Jukebox

https://openai.com/research/jukebox

Most Recent

DALL·E 2

DALL·E 2 is an AI system that can create realistic images and art from a description in natural language.

GPT-4

Input: An armchair in the shape of an avocado

Outline

- Introduction
- Who Should take the Course?
- Objective and syllabus
- Course logistics
- Homeworks, quizzes, projects, grading, oh my!
- Preps, teamwork and mentoring
  - And cheating…
- Challenges

Who is the course designed for?

- Students from any background that want to learn deep learning
- Students who are willing to put in 12-20 hours a week on this course
- Students who give continuous feedback and engage on Piazza
- Students who are mature and want to be challenged
- Students who want to be ready for AI research & engineering roles

Outline

- Introduction
- Who should take the course?
- Objective and syllabus
- Course logistics
- Homeworks, quizzes, projects, grading, oh my!
- Preps, teamwork and mentoring
  - And cheating…
- Challenges

Preparation for the course

- Required Classes
  - Fundamental of Programming (or equivalent) -
  - Python, PyTorch

- Highly Recommended Classes:
  - Calculus and Linear Algebra

- Nice to have:
  - Vector Calculus and Software Engineering

Course Objectives: **High Level**

- Understanding neural networks
- Comprehending the models that do the previously mentioned tasks (and maybe build them)
- Fearlessly design, build and train networks for various tasks

Course Objectives: **Mid Level**

- **Concepts**
  - Some historical perspective
  - Types of neural networks and underlying ideas
  - Learning in neural networks
    - Training, concepts, practical issues
  - Architectures and applications
- **Practical**
  - Familiarity with training
  - Implement various neural network architectures
  - Implement state-of-art solutions for some problems
- **Overall: to set you up for further work in your research area**

Course Objectives: **Topics**

- **Basic network formalisms:**
  - Multilayer Perceptrons
  - Convolutional Neural Networks
  - Recurrent Neural Networks
  - Boltzmann Machines
- **Some advanced formalisms**
  - Generative models: VAEs
  - Adversarial models: GANs
  - Graph Neural Networks
  - Transformers
- **Topics we will touch upon:**
  - Computer vision: recognizing images
  - Text processing: modelling and generating language
  - Machine translation: Sequence to sequence modelling
  - Modelling distributions and generating data
  - Speech recognition

Reading Materials

- List of books on the course webpage
- Additional reading material will also appear on the course pages

Outline

- Introduction
- Who should take this course?
- Objective and syllabus
- Course logistics
- Homeworks, quizzes, projects, grading, oh my!
- Preps, teamwork and mentoring
  - And cheating…
- Challenges

Instructor and TAs

- Instructor:
  - Bhiksha Raj
  - Rita Singh
- TAs:
  - List of TAs, with email ids on course page
  - We have TAs from the Pittsburgh and Kigali Campuses
  - Remote and in-person OH offerings and times are subject to change
- Office hours: see course webpage https://deeplearning.cs.cmu.edu

Lectures

- In-class lectures
  - Live streaming for remote sections
  - Lectures will also be recorded and uploaded to the course website
- Importance of viewing lectures
  - Even if you think you know the topic
  - Quizzes: your marks depend on viewing lectures
  - We monitor attendance (more on this later)

some lectures may be conducted entirely over zoom.

Additional Logistics

- Discussions: Piazza (please be updated with the recent threads)
- Compute Infrastructure:
  - Amazon tokens (AWS) will be provided
  - Google Cloud Platform (GCP)
  - Kaggle

Lecture Attendance

- You get marks for attendance (i.e. class participation)
  - Our performance metrics over the semesters show a distinct correlation between attendance and course scores
  - We also note a distinct *inverse* correlation between attendance and the amount of help you require on piazza and during office hours
  - To encourage attendance, we assign 1 mark for attendance
    - 1% of your total grade for 11685 and 11785
    - 1.3333% of your total grade for 11485
    - This can be the difference between a B and an A
- We will track lecture attendance
  - More on next slide

Lecture Attendance: Rules

- You must either attend the lectures in person, or watch the streamed live video
  - With some exceptions: see below
  - **We get attendance stats from your participation during in-class polls**
  - Polls will be conducted via Piazza for in-person lectures. It is recommended to use the Piazza app available both on Android and iOS.
  - Lectures conducted over zoom will use zoom polls
  - **Use of electronic devices during in-person lectures is not permitted except during polls**
- Students in SV, and in other time zones where your local time is before 8am or after 5pm for the class, may alternately watch recorded lectures on mediatech instead
  - Mediatech records who watched and for how long. Entire lectures must be viewed.
  - If viewed on mediatech, the lectures of each week must be viewed before 8AM of the Monday following the following week
    - Otherwise, it doesn't count
- The precise mechanism for computing the class participation score from attendance will be decided at the end of the semester. **Our advice to you: just attend**

Lecture Schedule

- On website
  - The schedule for the latter half of the semester may vary a bit
    - Guest lecturer schedules are fuzzy..
- Guest lectures:
  - TBD

Recitations

- 14 recitations (does not include HW Bootcamps)
- Held on Fridays
- Hands-on, will cover implementation specifics while you follow along
  - Important to get the maximum out of the course
- Topic list on the course schedule
- *Strongly recommend attending all recitations*
  - *Even if you think you know everything*

Outline

- Introduction
- Who should take the course?
- Objective and syllabus
- Course logistics
- Homeworks, quizzes, projects, grading, oh my!
- Preps, teamwork and mentoring
  - And cheating…
- Challenges

Evaluation

- Scores are evaluated based on 3 types of tests:
  - Weekly Quizzes
  - Home works
  - Team Project (not for 11485)

Weekly Quizzes

- Questions are related to topics covered that week
  - In lecture slides
  - Released Fridays @ 11:59PM EST
  - Closed Sundays @ 11:59PM EST
  - (this may occasionally shift, don't panic!)
- 14 total quizzes
  - We accept the best 12 (you may skip 2)
    - This is expected to account for any circumstance-based inability to work on quizzes
- Each has 10 multiple-choice questions
  - We accept the best of 3 attempts per quiz

Weekly Quizzes

- Slides often contain a lot more information than is presented in class
- *Quizzes will contain questions from topics that are on the slides, but not presented in class*
- *Will also include topics covered in class, but not on online slides!*
- There will be questions based on latest research papers in the quiz – the links to the papers will be provided

Homework

- There will be four in-term homeworks and a Recitation-0
  - Recitation 0: recommended preparatory material for the course
  - Homeworks 1-4: neural-net exercises

- Homeworks 1-4 all have two parts:
  - Part 1: Autograded problems with deterministic solutions
    - You must upload them to autolab
  - Part 2: Open problems posted on Kaggle
  - Bonus and Autograd Bonus HWs:
    - There will be bonus homeworks associated with some (or all) of the homeworks; there will also be a separate autograd bonus homework.
    - These marks will not contribute to final grading curves and give you the chance to make up for marks misse elsewhere

Homework: Part 1

- HW Part-1s evaluate your ability to code neural nets on your own from scratch
  - If you implement all mandatory and bonus questions of part 1 of all homeworks, you will, hopefully, have all components necessary to construct a little neural network toolkit of your own
    - "mytorch" ☺️
- The homeworks are autograded
  - Be careful about following instructions carefully
    - The autograder is setup on a computer with specific versions of various packages
    - Your code must conform to their restrictions, if not the autograder will often fail and give you errors or 0 marks, even if your code is functional on your own computer
  - Note: you get some bonus points if you submit the homework early

Homework: Part 2

- HW Part-2s test your ability to solve complex problems on real-world data sets
  - Unlike part-1s, initial early submissions are required and worth 10/100 points (to show you have gotten started)
  - The remaining 90/100 points are determined by your final score relative to the cutoffs

Homework: Part 2

- The standard cutoffs are (After 10 pts ESD)
  - High cutoff : 90 pts
  - Medium cutoff : 70 pts
  - Low cutoff : 50 pts
  - Very Low cutoff : 30 pts
  - No submission : 0 pts
- Your score is linearly interpolated between cutoffs. Ex, if an 85% accuracy model is for the high cutoff and 75% model is for the medium, a model with 80% accuracy will get 82.5/90 pts
- There is a 5th bonus cutoff for extra credits (10 pts) determined by the highest scoring student
  - Ex. If the highest accuracy is 92%, the scoring student will get 110/100 and anyone between 85% to 92% will get 100+/100 (assuming they have 10 bonus points from early submission)

Homework Deadlines

- There are separate deadlines for HW Part-1s and HW Part-2s

- HW Part-2s have multiple deadlines
  - *Early Submission D*eadline: worth 10% of Part-2 and affirms you have started the assignment
  - *On-time Submission Deadline:* determines your eligibility for full cutoff marks
  - *Late/Slack Deadline:* everyone gets up to 10 total slack days **(does not apply to initial submissions, all bonus points, nor HW Part-1s)** -- you can distribute them as you want across your HWs, and once you use up your slack days, all subsequent late submissions will accrue a 10% penalty (on top of any other penalties)

- Note: Kaggle: Kaggle leaderboards stop showing updates on full-submission deadline
  - We will continue to accept submissions until you run out of slack days (the drop-dead deadline)

- ***Please refer to the course webpage for a complete set of policies***

Course Project

- If you're taking 11-485 this slide does not apply to you
- If you're taking 11-785, you will be required to do an original course project
- If you're taking 11-685, you will be required to do an assigned course project

- Projects are intended to exercise your ability to comprehend and implement ideas beyond those covered by the HWs

- Projects can range from
  - Implementation and evaluation of cutting-edge ideas from recent papers
    - *Verification* of results from "hot" published work
  - Problem-solving that might lead to publication, if completed well
  - New models/learning algorithms/techniques, with proper evaluation
  - … etc.

Course Project

- Teams: think about forming project teams as soon as possible (we may help, too)
  - 11-785 teams are of 3-4 students maximum
  - 11-685 teams are of 2 students maximum
- Project requirements:
  - Report:
    - Submit a project proposal by the proposal deadline listed on the website
    - Submit an initial report 3⁄4 way through the semester
    - Submit an updated report draft 3 days **before the presentation due date**
    - Submit a final report during finals week, at the end of the semester
  - Video:
    - Make a **5 min** video presentation of the project at the end of the semester
    - The video can be presented by one, some, or all team members
    - Video should explain the problem, proposed solution, and evaluation
    - Defend your project by answering questions about your video, on Piazza
- *Each team will be assigned a mentor from among the TAs, who will monitor your progress and assist you if possible.*
- **The project is often the most enjoyable portion of the course**

Grading

Weekly Quizzes		**24**
14 Quizzes, bottom two dropped		24
Assignments		**50**
HW1 – Basic MLPs	(AL + Kaggle)	12.5
HW2 – CNNs	(AL + Kaggle)	12.5
HW3 – RNNs	(AL + Kaggle)	12.5
HW4 – Sequence to Sequence Modelling	(AL + Kaggle)	12.5
Team Project (Not for 11-485)		**25**
Proposal		-
Mid-term Report		5
Preliminary Full Report		-
Project Presentation		10
Peer Reviewing		binary multiplier
Final report		10

*** Note: There is 1 mark for attendance**

Outline

- Introduction
- Who should take the course?
- Objective and syllabus
- Course logistics
- Homeworks, quizzes, projects, grading, oh
- Preparation, teamwork, mentoring
  - And cheating…
- Challenges

Preparation for the course

- This course is implementation heavy
  - There is a lot of coding
  - There is a lot of experimenting
  - You will work with some large datasets
- Language of choice: Python
- Toolkit of choice: PyTorch
  - You are welcome to use other languages/toolkits, but the TAs will not be able to help with coding/homework
- We hope you have gone through
  - Recitation-0 (optional)
    - Not graded

Course Communication

- Class Poll
- Question and Answering
- Announcements and Reminders
- Student Collaboration
- Instructor Notes

https://piazza.com/

Study groups

- ***Everyone must join a study group***
- ***Each study group will be assigned a TA mentor for help throughout the course***
- We recommend study groups of 3-4 people
- If you do not have a study group, we will form one for you
  - Please register on the forms posted on Piazza

Study groups

- What study groups may do:
  - Discuss homework problems and solutions
  - Discuss papers
  - Discuss class work
  - Discuss quizzes
- We encourage you to meet regularly to discuss IDL work
- Study groups may also go on to form project teams

Teamwork vs Cheating

- Learning happens best together
  - You may learn more from each other than from us
- While we encourage teamwork, *there are strict rules…*

Study groups vs. Cheating

- Every student must solve their quizzes by themselves
  - You may discuss the questions with your study groups/friends, but when you solve the quiz**, isolate yourself and do it alone**
- Every student must solve every homework by themselves
  - You may discuss the homeworks with your friends, and even help them debug their code, but when you finally solve it, every line of your code (except libraries that have been okayed by course staff) must be written by you
  - Your solution must be yours
- Plagiarizing code from the web or your friends constitutes cheating
  - Submitting any solutions not obtained by you constitutes cheating

Cheating

- You are here to learn DL yourself, not to demonstrate how well your friend, or that guy on the web has learned DL

- You are at CMU – which means you are among the brightest and best students in the *world*
  - You probably were among the top students in your peer group all your life, before you came here
  - It will be an insult to yourself and everything you ever stood for in your life to lower yourself from your own standards and start cheating
    - **So don't!!!**

- If you are unsure whether something you're doing constitutes cheating or not, come check with us

Study Group Mentors

- Every study group will be assigned a TA mentor
  - We will track your progress and reach out to you if you appear to be in trouble

- If you are in trouble, reach out to your TA mentor or the instructor
  - If you feel you're falling behind, reach out
  - If you feel you are struggling, reach out
  - If you feel pressured/unable to cope, reach out
  - We will try our best to help you
  - *Please watch Recitation 0I if you are stuck or feeling overwhelmed*

- We aim to make this a successful course for all of you
  - In our ID(ea)L world, everyone performs well enough to get an A
  - *Without lowering our standards* – i.e. we would like to bring *you* all up to where we believe you deserve an A
  - Everything about this course is geared to that objective

Outline

- Introduction
- Who should take the course ?
- Objective and syllabus
- Course logistics
- Homeworks, quizzes, projects, grading, oh my!
- Preps, teamwork and mentoring
  - And cheating…
- Challenges

This course is not easy

- A lot of work!

This course is not easy

- A lot of work!
- A lot of work!

This course is not easy

- A lot of work!
- A lot of work!!
- A lot of work!!!

This course is not easy

- A lot of work!
- A lot of work!!
- A lot of work!!!
- **A LOT OF WORK!!!!**

Not for chicken!

This course is not easy

- A lot of work!
- A lot of work!!
- A lot of work!!!
- **A LOT OF WORK!!!!**

- But a lot of fun :)

This course is not easy

- A lot of work!
- A lot of work!!
- A lot of work!!!
- **A LOT OF WORK!!!!**
- *Mastery-based* evaluation
  - Quizzes to test your understanding of topics covered in the lectures
  - HWs to teach you to implement complex networks
    - And optimize them to high degree
  - Projects to expose you to real world DL problems
  - **FYI: Anyone who gets an "A" in the course is technically ready for a deep learning job**

Recitation 0

- If you have any doubts about your preparations for this course, please, please, please, please go through the videos for Recitation-0 and complete the notebook exercises
  - Gain comfort with the coding required in the homeworks, early on

- HW1 Part 1 has components intended to help you *later* in the course
  - It's worth it
  - It’s the easiest HW!

Questions?

- Please post on piazza
- Feel free to post on piazza for any clarification or doubts
- We strive to respond to student queries within 5 minutes

Thank you for
deciding to take this journey
of Deep Learning
with us

- Prof. Bhiksha and TAs