Introduction to Deep Neural Networks

0. Logistics

Fall 2020

Outline

- Introduction
- Objectives and syllabus
- Course logistics
- Homeworks, quizzes, projects, grading, oh my!
- Prep, teamwork and mentoring
  - And cheating…
- Challenges

Outline

- Introduction
- Objectives and syllabus
- Course logistics
- Homeworks, quizzes, projects, grading, oh my!
- Prep, teamwork and mentoring
  - And cheating…
- Challenges

Neural Networks are taking over!

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

- In many problems they have established the state of the art
  - Often exceeding previous benchmarks by large margins

Breakthroughs with neural networks

TECHNEWSWORLD EMERGING TECH SEARCH

Computing Internet IT Mobile Tech Reviews Security Technology Tech Blog Reader Services

Microsoft AI Beats Humans at Speech Recognition

By Richard Adhikari
Oct 20, 2016 11:40 AM PT

Print Email

G+1 5 Tweet 25 Share 45 Share 11 Share 0 share 104

Image: Adobe Stock

Microsoft's Artificial Intelligence and Research Unit earlier this week reported that its speech recognition technology had surpassed the performance of human transcriptionists.

Most Popular Newsletters News Alerts

How do you feel about Black Friday and Cyber Monday?

- They're great -- I get a lot of bargains!
- The deals are too spread out -- I'd prefer just one day.
- They're a fun way to kick off the holiday season.
- I don't like the commercialization of Thanksgiving Day.
- They're crucial for the retail industry and the economy.
- The deals typically aren't that good.

Vote to See Results

E-Commerce Times

- Black Friday Shoppers Hungry for New Experiences, New Tech
- Pay TV's Newest Innovation: Giving Users Control
- Apple Celebrates Itself in $300 Coffee Table Tome
- AWS Enjoys Top Perch in IaaS, PaaS Markets
- US Comptroller Gears Up for Blockchain and

Breakthroughs with neural networks

Image segmentation & recognition

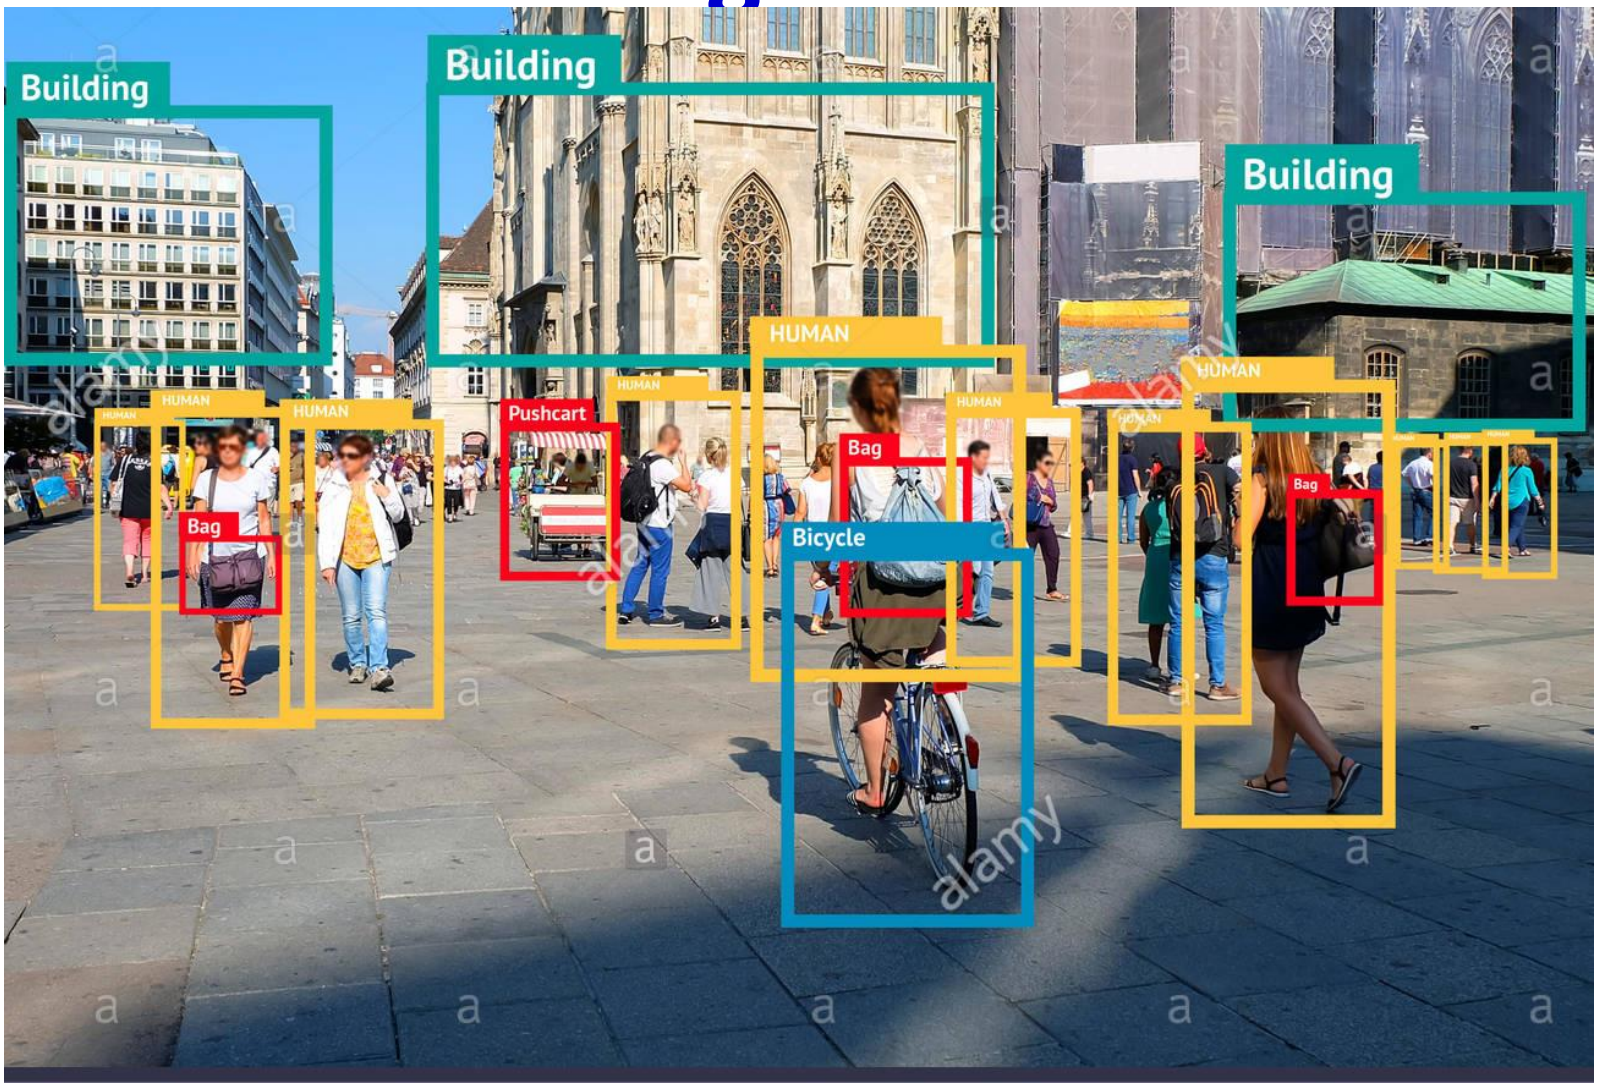

Image recognition

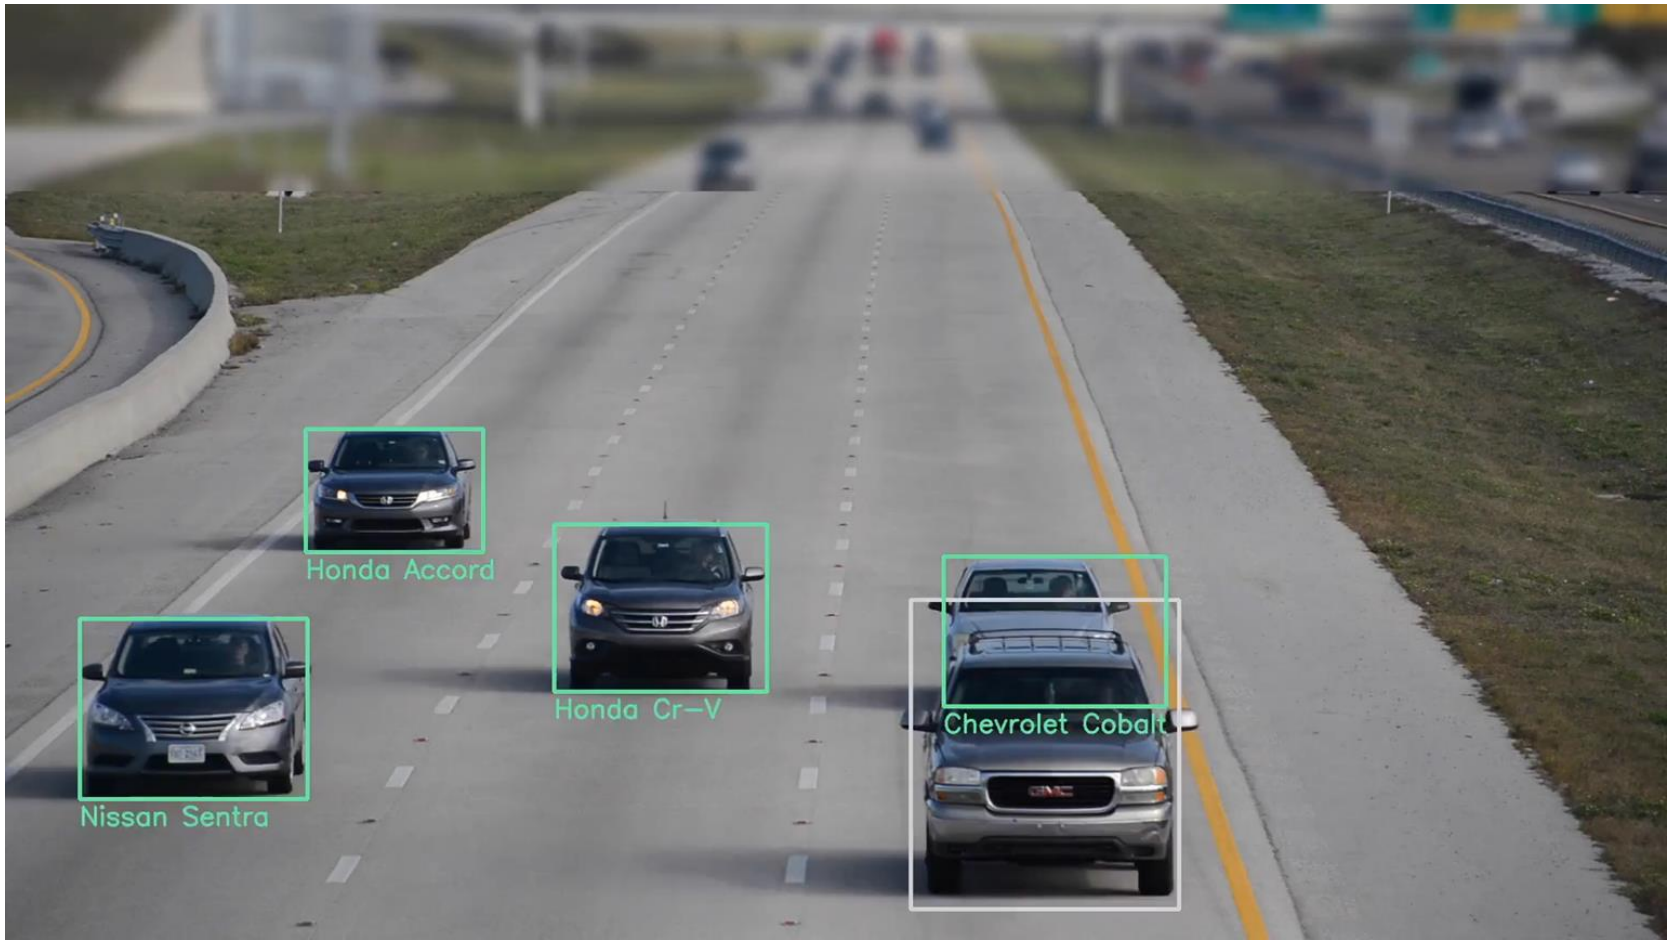

https://www.sighthound.com/technology/

Breakthroughs with neural networks

AlphaGo

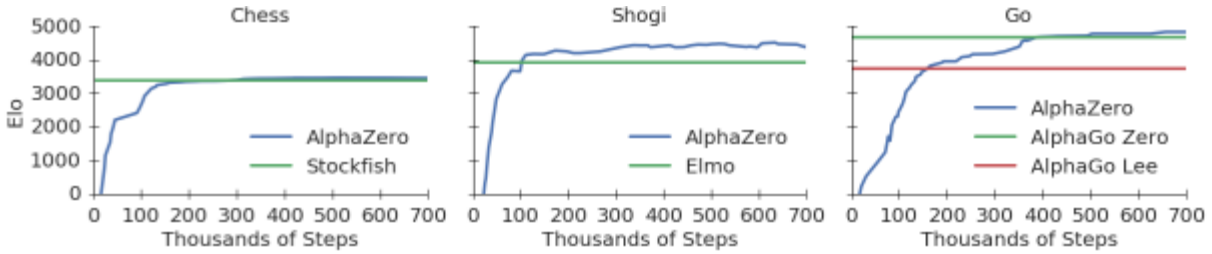

Figure 1: Training *AlphaZero* for 700,000 steps. Elo ratings were computed from evaluation games between different players when given one second per move. **a** Performance of *AlphaZero* in chess, compared to 2016 TCEC world-champion program *Stockfish*. **b** Performance of *AlphaZero* in shogi, compared to 2017 CSA world-champion program *Elmo*. **c** Performance of *AlphaZero* in Go, compared to *AlphaGo Lee* and *AlphaGo Zero* (20 block / 3 day) (*29*).

Breakthroughs with neural networks

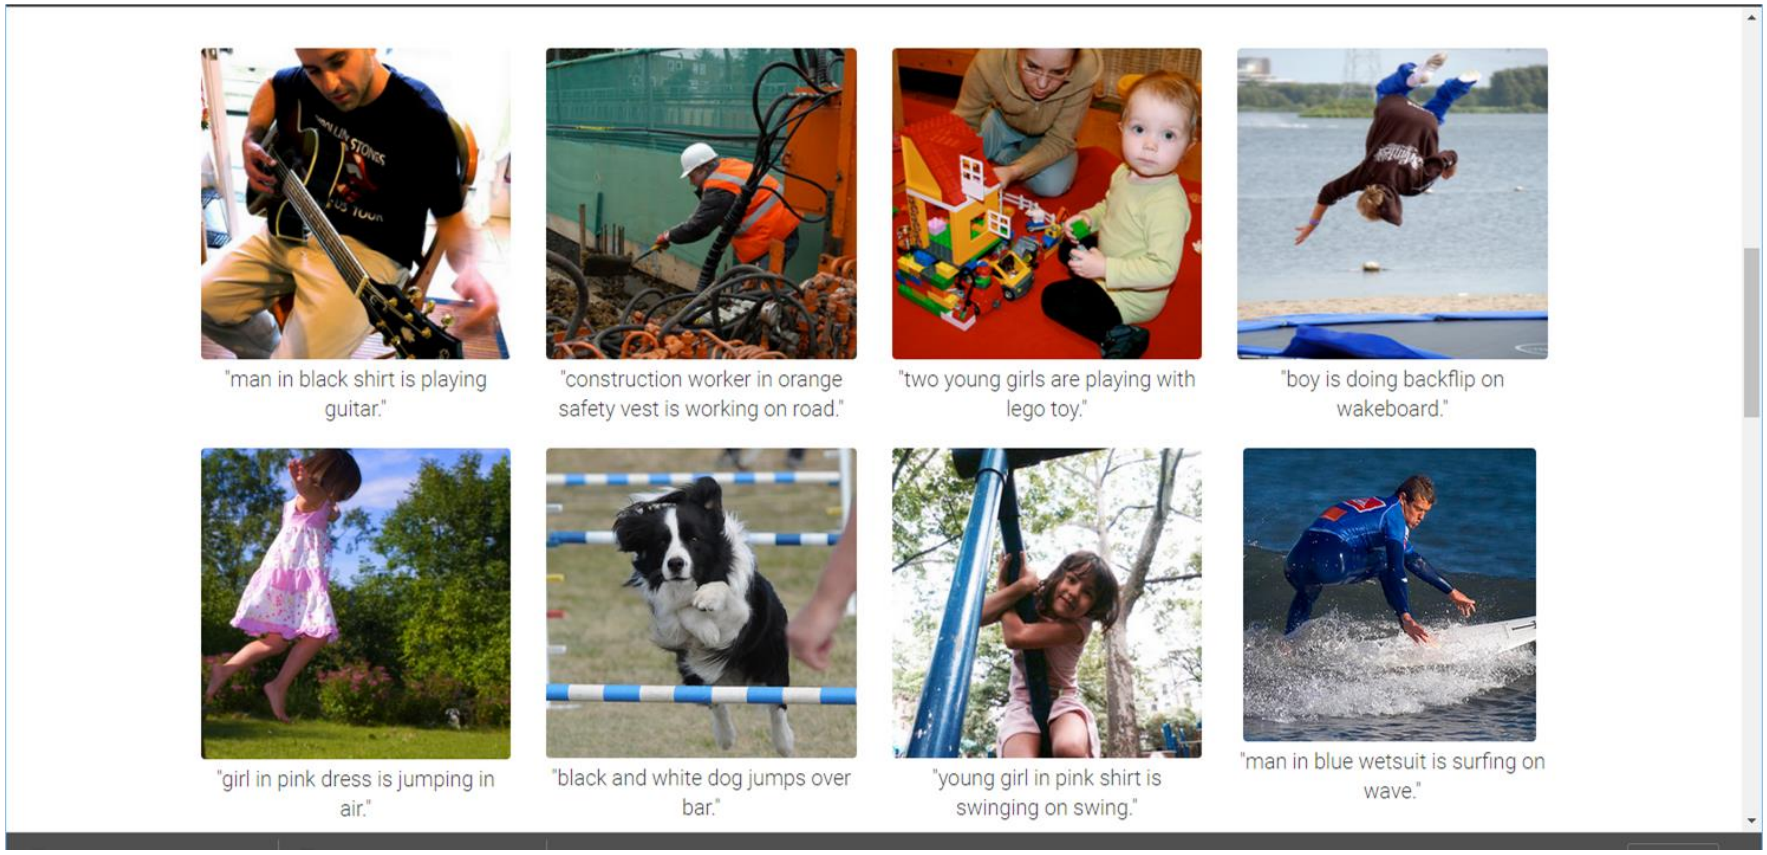

- Captions generated entirely by a neural network

Breakthroughs with neural networks

ThisPersonDoesNotExist.com uses AI to generate endless fake faces

Hit refresh to lock eyes with another imaginary stranger

By James Vincent | Feb 15, 2019, 7:38am EST

SHARE

A few sample faces — all completely fake — created by ThisPersonDoesNotExist.com

– https://www.theverge.com/tldr/2019/2/15/18226005/ai-generated-fake-people-portraits-thispersondoesnotexist-stylegan

Successes with neural networks

- And a variety of other problems:
  - From art to astronomy to healthcare...
  - and even predicting stock markets!

Neural Networks and the Job Market

This guy didn't know about neural networks (a.k.a deep learning)

This guy learned about neural networks (a.k.a deep learning)

Outline

- Introduction
- Objectives and syllabus
- Course logistics
- Homeworks, quizzes, projects, grading, oh my!
- Prep, teamwork and mentoring
  - And cheating…
- Challenges

Course Objectives

- Understanding neural networks
- Comprehending the models that do the previously mentioned tasks
  - And maybe build them
- Familiarity with some of the terminology
  - What are these:
    - http://www.datasciencecentral.com/profiles/blogs/concise-visual-summary-of-deep-learning-architectures
- Fearlessly design, build and train networks for various tasks

- *You will not become an expert in one course*

Course objectives: Broad level

- Concepts
  - Some historical perspective
  - Types of neural networks and underlying ideas
  - Learning in neural networks
    - Training, concepts, practical issues
  - Architectures and applications
  - Will try to maintain balance between squiggles and concepts (concept >> squiggle)
- Practical
  - Familiarity with training
  - Implement various neural network architectures
  - Implement state-of-art solutions for some problems
- Overall: Set you up for further research/work in your research area

Course learning objectives: Topics

- Basic network formalisms:
  - MLPs
  - Convolutional networks
  - Recurrent networks
  - Boltzmann machines
- Some advanced formalisms
  - Generative models: VAEs
  - Adversarial models: GANs
- 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
  - Reinforcement learning and games
  - Speech recognition

Reading

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

Outline

- Introduction
- Objectives and syllabus
- Course logistics
- Homeworks, quizzes, projects, grading, oh my!
- Prep, teamwork and mentoring
  - And cheating…
- Challenges

Instructors and TAs

- Instructor: Bhiksha Raj
  - bhiksha@cs.cmu.edu
  - x8-9826
- TAs:
  - List of TAs, with email ids on course page
  - We have TAs for the
    - Pitt Campus
    - Kigali,
    - SV campus,
  - Please approach your local TA first
- Office hours: On webpage
- http://deeplearning.cs.cmu.edu/

Logistics: Lectures..

- Class is entirely online
- Lectures are held over zoom
- Recordings will be posted

- Important that you view the lectures
  - Even if you think you know the topic
  - Your marks depend on viewing lectures
  - We will monitor attendance (more on this later)

Additional Logistics

- Discussions:
  - On Piazza
- Compute infrastructure:
  - Everyone gets Amazon tokens
  - Initially a token for $50
  - Can get additional tokens of $50 up to a total of $150

Lecture Attendance

- We will track lecture attendance
  - You must either be in class or watch the videos uploaded to mediatech
    - If you're watching the videos on mediatech, we can track who watched what lecture and for how long,
    - If you're in class, we assume you attended the entire class
- We will report these numbers to you
  - With the primary objective of informing you about how much attention you're paying to the class
  - We highly highly recommend greater than 75% attendance
    - We assume you're in the class to *learn*

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
    - Scott Fahlman, Shinji Watanabe, Gerald Friedland, Graham Neubig
      - Or some subset thereof

Recitations

- We will have 13 recitations
  - Possibly a 14th if TAs and students are still enthusiastic after 15 grueling weeks
- Will cover implementation details and basic exercises
  - Very important if you wish 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*

Recitations Schedule

- Every Friday of the semester
- See course page for exact details!

Outline

- Introduction
- Objectives and syllabus
- Course logistics
- Homeworks, quizzes, projects, grading, oh my!
- Prep, teamwork and mentoring
  - And cheating…
- Challenges

Evaluation

- Performance is evaluated based on 3 types of tests
- Weekly Quizzes
- Homeworks
- Team Project

Weekly Quizzes

- 10 multiple-choice questions
- Related to topics covered that week
  - On both slides and in lecture
- Released Friday, closed Sunday night
  - This may occasionally shift, don't panic!
- There will be 14 total quizzes
  - We will consider the best 12
  - This is expected to account for any circumstance-based inability to work on quizzes
    - You could skip up to 2

Lectures and 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!*

Homeworks

- There will be one early homework (released before the start of the semester) and four in-term homeworks
  - Homework 0: Preparatory material for the course
  - Homeworks 1-4: Actual neural-net exercises

- Homeworks 1-4 all have two parts:
  - Part 1: Autograded problems with deterministic solutions
    - You must upload them to autolab
    - Will include mandatory parts and "bonus" parts
    - "bonus" questions will not contribute to final grading curves and give you the chance to make up for marks missed elsewhere
  - Part 2: Open problems posted on Kaggle

Homeworks 1-4 – Part 1

- Part 1 of the homeworks evaluate your ability to code in 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

Homeworks 1-4, Part 2

- Part 2 of every homework tests your ability to solve complex problems on real-world data sets

- These are open problems posted on Kaggle
  - You compete with your classmates on a leaderboard
  - **We post performance cutoffs for A, B and C**
    - **If you achieved the posted performance for, say "B", you will at least get a B**
    - **A+ == 105 points (bonus)**
    - **A = 100**
    - **B = 80**
    - **C = 60**
    - **D = 40**
    - **No submission: 0**
  - **Actual scores are linearly interpolated between grade cutoffs**
    - Interpolation curves will depend on distribution of scores

Homework Deadlines

- Multiple deadlines
- Separate deadline for Autograded deterministic component
- Kaggle component has multiple deadlines
  - *Initial submission* deadline : Making this is worth 5% of the marks of the HW
  - *Full submission deadline:* Your final submission must occur before this deadline to be eligible for full marks
  - *Drop-dead deadline:* Must submit by here to be eligible for *any* marks
    - Day on which solution is released

- Homeworks: Late policy
  - Everyone gets up to 7 total slack days (does not apply to initial submission)
  - You can distribute them as you want across your HWs
    - You become ineligible for "A+" bonus if you're using your grace days for Kaggle
  - Once you use up your slack days, all subsequent late submissions will accrue a 10% penalty (on top of any other penalties)
  - There will be no more submissions after the drop-dead deadline
  - Kaggle: Kaggle leaderboards stop showing updates on full-submission deadline
    - But will continue to privately accept submissions until drop-dead deadline

- ***Please see course webpage for complete set of policies***

Course project

- If you're taking 11-685/11-785/18-786, you will be required to do a course project
- Projects are done by teams of students
  - Ideal team size is 4
  - You are encouraged to form your teams early
- Projects are intended to exercise your ability to comprehend and implement ideas beyond those covered by the HWs

- Project can range from
  - Implementing and evaluating cutting-edge ideas from recent papers
    - *Verifying* results from "hot" published work
  - "Researchy" problems that might lead to publication if completed well
  - Proposing new models/learning algorithms/techniques, with proper evaluation
  - Etc.

Course project

- Project teams must be formed by the first week of **October**
  - If you don't form your own teams, we will team you up
- Each team must:
  - Submit a project proposal by the second week of October
  - Submit a mid-way report ¾ way through the semester **(First week of November)**
  - Submit a preliminary full report three days **before the presentation due date**
  - Make a **5 min** video presentation of the project at the end of the semester
    - Can be presented by one, some, or all team members
    - Will be evaluated by the instructor, TAs, and your classmates
    - Ensure you explain the problem, proposed solution, and the evaluation clearly
    - Allocate enough time to make the presentation, it is not as easy as you think
    - Poor presentation can significantly affect your project score :)
  - Submit a final full report at the end of the semester
  - Templates for proposals and reports will be posted
- *Each team will be assigned a mentor from among the TAs, who will monitor your progress and assist you if possible.*
- **More details on project evaluations will be posted towards the end of the sem**
- **The project is often the most fun portion of the course**

Grading

Weekly Quizzes		**24%**
14 Quizzes, bottom two dropped		24%
Assignments		**51%**
HW0 – Preparatory homework	(AL)	1%
HW1 – Basic MLPs	(AL + Kaggle)	12.5%
HW2 – CNNs	(AL + Kaggle)	12.5%
HW3 – RNNs	(AL + Kaggle)	12.5%
HW4 – Sequence to Sequence Modelling	(Kaggle)	12.5%
Team Project (Not for 11-485)		**25%**
Proposal		TBD
Mid-term Report		TBD
Preliminary Full Report		TBD
Project Presentation		TBD
Peer Reviewing		TBD
Final report		TBD

Homeworks

- We have planned for a fifth homework to be released as a “make up” for incomplete projects, in the event of a catastrophe
  - Remember Spring 2020

- If released it will
  - Be worth only 60% of full project score
  - Be released after the project deadline
  - You will get scored for *either your project or HW5!*

Outline

- Introduction
- Objectives and syllabus
- Course logistics
- Homeworks, quizzes, projects, grading, oh my!
- Prep, teamwork and mentoring
  - And cheating…
- Challenges

Preparation for the course

- Course is implementation heavy
  - A lot of coding and experimenting
  - 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
    - Some support for TensorFlow

- We hope you have gone through
  - Recitation zero
  - HW zero
    - Carries marks

Teamwork

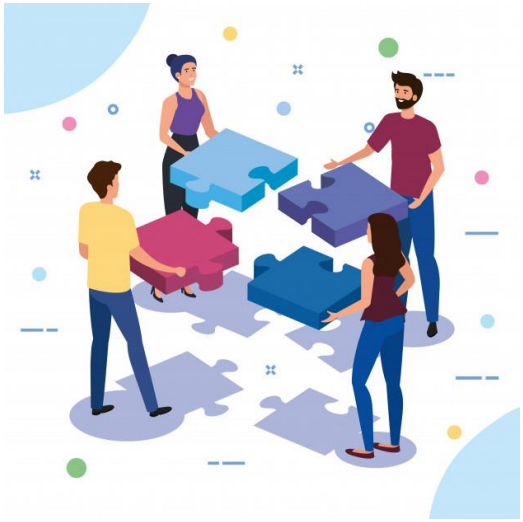

- Learning happens best together
  - You will learn more from each other than you will from us
- *We encourage teamwork*
  - But there are strict rules…

Study groups

- Please form study groups
- If you do not have a study group of your own, we will form one for you
  - Please register on Sean's forms
- *Everyone must be part of a study group*

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
- So what are the caveats? What may you *not* do

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, 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
  - And submitting 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, check with us

Mentoring

- 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 in trouble, reach out to your TA mentor and/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
- We aim to make this a successful course for all of you
  - In our ideal 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
- Objectives and syllabus
- Course logistics
- Homeworks, quizzes, projects, grading, oh my!
- Prep, 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 somewhat calibrated (over the years) to ensure it is doable

Over 60% of students got some flavor of A each of the past three semesters and they deserved it

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
- Target: Anyone who gets an "A" in the course is technically ready for a deep learning job

HW0 / Recitation 0

- Please, please, please, please, please go through the videos for recitation 0, and complete HW0.
  - These are essential for you to gain comfort with the coding require in the following homeworks

- HW1 part 1 also has many components intended to help you *later* in the course
  - So if it seems a bit dense, please bear with it, its worth it
- HW1 is the easiest HW!

Questions?

- Please post on piazza