Carnegie Mellon University
Master of
Software Engineering

17-[644: Applied Deep Learning]
[M/W 11:50am-1:05pm, 3SC 265]
[A4, Spring 2021, 6 Units]

Instructor	Email	Office Location & Hours
Prof. Vincent Hellendoorn	vhellendoorn@cmu.edu	Online, Mon 10-11am

Course Description. Deep neural networks have made in-roads in virtually every industry, propelled by exponential increases in compute power, data, and fundamental progress in modeling. Knowledge of these models is fast becoming a key asset for software engineers, as modern software systems increasingly involve neural components (think: self-driving cars, augmented reality, robotics). This course equips the next generation of software engineers with knowledge of neural models, the software engineering challenges involved in using these, and hands-on experience with their applications. It teaches both a rich vocabulary of general, essential concepts (including architectures), and recent work on applications of these models. Coursework includes regular in-class quizzes and a group project with regular milestones aimed at constructing a neural solution for an existing application that will be used to teach the various stages (and their pitfalls) of building and deploying deep learners.

Prior Knowledge. Basic knowledge of programming (esp. Python) and software engineering concepts. Familiarity with basic machine learning -- this course will not cover fundamental ML topics.

Learning Objectives. After completing this course, you will be able to:

- Assess the deep learning needs and options in a variety of real-world problem settings.
- Construct datasets and training pipelines that effectively meet performance targets.
- Accurately and realistically evaluate model performance and validity.

Learning Resources. Publicly available resources on deep learning, including online textbooks, research papers, and blog posts. Implementations and datasets for hands-on experience.

Assessments & Grading. Most classes after week 1 will begin (or end) with a short quiz reviewing the recent material. There will be 10 such quizzes, accounting for 25% of the total course grade. Students learn more by applying and explaining ideas to others. In this course, students will apply their knowledge by incrementally implementing and improving in a class

project. This project will be completed in groups of 3 to 5 and spans the length of the course: students will sign up for a topic (based on recent research applications) in week 1 and present+report on their final result in week 7. These projects will be implemented on Github where we will make use of features such as pull requests and releases to make each member's contribution transparent.

Assessment	Final Grade %
Quizzes (10 total)	25%
Project milestone 1	20%
Project milestone 2	20%
Project final report & presentation	30%
Class participation	5%

Grade	Percentage Interval
A	90-100%
B	80-89%
C	70-79%
D	60-69%
R (F)	59% or below

Project details: Partial progress will be assessed at two milestones (in weeks 3 & 5). Each of these sub-deliverables counts for 20% of the final grade (see table above) and will consist of: a) a new release on Github (with associated commits & reviews), b) a brief (3-5 minute) in-class presentation on the project status, and c) a short (1-page) written report of the work done and corresponding results. To ensure that grades reflect individual contributions as well as the overall result, 75% of the grade for each milestone (so 15% of final) is based on the overall quality of the deliverable and 25% is based on individual participation, as gauged by contributions on Github. At a minimum, students are expected to both submit and review one or more commits/Pull Requests.

- **Milestone 1:** load, preprocess, and document the statistics of the (project-specific) problem data. Train a simple model to assess & report the baseline quality. Submit a ~2 page report outlining the basic problem, data and existing solutions. Prepare a 1-3 minute presentation outlining the problem space and existing solutions, plus preliminary plans.
- **Milestone 2:** take first steps to improving performance by making task-appropriate changes, such as augmenting the input features, using more modern models, better regularization, etc. Report appropriate metris on final result. Expand previous report to ~4 pages, adding experimental details and preliminary results. Prepare a 1-3 minute presentation on the state of the project (results, bottlenecs).
- **Final submission:** complete experiments involving both baselines, the proposed enhancements, and ablations to analyze the impact of the proposed changes on a property of interest (performance, explainability, model size, etc.). Submit the final 6-8 page, ICLR-formatted report with the customary sections (introduction, overview of approach, details, results, related work), and present for ca. 15-20 minutes on this work. Presentations may be by one or more team members.

Class participation: Enrich in-class discussions with your insight, relevant experience, critical questions, and analysis of the material. The quality of contribution is more important than the quantity.

Course and Grading Policies

- **Late-work policy**: All work is expected to be handed in at the indicated due date and time (see the course schedule page). For fairness to the whole class, no late submissions will be accepted for the group work.
- **Quizzes:** You will be allowed to drop your two lowest quiz grades. This includes quizzes in classes you were unable to attend.
- This semester involves regular use of technology during class — both for in-person and remote students. Research has shown that divided attention is detrimental to learning, so I encourage you to close any windows not directly related to what we are doing while you are in class. Please turn off your phone notifications and limit other likely sources of technology disruption, so that you can fully engage with the material, each other, and me. This will create a better learning environment for everyone.

Course Schedule. The following schedule provides a general overview of topics and assignments. Please refer to Canvas for specific lecture reading assignments and due dates.

Class	Date	Topic and Lectures
1	3/14	Overview: general DL pipeline, brief history of DL (incl. artificial neurons, perceptrons), applications & challenges in practice. **Consider reading:** this chapter on standard (feedforward) networks and how they transform their data.
2	3/16	DL basics: terminology, algebraic basics (vectors/matrices/tensors, stochastic gradient descent) **After class:** sign up for course project topic **Consider reading:** Chapter on linear algebra (+ notation) (if you're in need of a refresher), Machine learning basics (especially for those with little ML background).
3	3/21	Inputs 1a: data mining basics, input analysis/statistics and modeling ramifications, hardware accelerators, embeddings. ***Start of assignment 1***
4	3/23	Architectures 1a: common layers (MLP, CNN, RNN), auto-encoders, encoder/decoders
5	3/28	Evaluation 1a: losses (L2, [B/C]CE) and relation to metrics ([top-k] accuracy, precision/recall)

		Consider reading: Deep Learning book chapter 8.
6	3/30	*Inputs 1b: layer initializations, input scaling, normalization; multi-modal learning, sub-tokenization, intro to regularization* ***Prepare:*** *brief project progress-updates 1* ***After class:*** *assignment 1 due* *Start of assignment 2*
7	4/4	Architectures 1b: attention (+ high-level Transformers), pointers
8	4/6	Evaluation 1b: validation sets + bias/variance & regularization, impact of dimensions, hyper-parameter search/tuning.
9	4/11	Advanced input topics: intrinsic dimensionality, PCA, T-SNE
10	4/13	Advanced architecture topics: graphical models, Long-range Transformers, BERT/GPT-X, Bayesian learning. ***After class:*** assignment 2 due *Start of assignment 3 (final)*
11	4/18	Advanced evaluation topics: advanced metrics (KLD, top-K loss), beam search, explainability/transparency. Reading: Deep learning book, chapter 19 **Prepare:** brief project progress-updates 2
12	4/21	Advanced ML Topics: Bayesian learning & variational inference; meta-learning, reinforcement learning Reading: Deep learning book, chapter 20
13	4/25	Advanced ML Topics ct'd/Presentations
14	4/27	Presentations & final discussions **Final report due**

Accommodations for Students Disabilities. If you have a disability and have an accommodations letter from the Disability Resources office, I encourage you to discuss your accommodations and needs with me as early in the semester as possible. I will work with you to ensure that accommodations are provided as appropriate. If you suspect that you may have a disability and would benefit from accommodations but are not yet registered with the Office of Disability Resources, I encourage you to contact them at access@andrew.cmu.edu.

Academic Integrity. Honesty and transparency are important to good scholarship. Plagiarism and cheating, however, are serious academic offenses with serious consequences. If you are

discovered engaging in either behavior in this course, you will earn a failing grade on the assignment in question, and further disciplinary action may be taken.

For a clear description of what counts as plagiarism, cheating, and/or the use of unauthorized sources, please see the University's Policy on Academic Integrity.

If you have any questions regarding plagiarism or cheating, please ask me as soon as possible to avoid any misunderstandings. For more information about Carnegie Mellon's standards with respect to academic integrity, you can also check out the Office of Community Standards & Integrity website.

Student Wellness. As a student, you may experience a range of challenges that can interfere with learning, such as strained relationships, increased anxiety, substance use, feeling down, difficulty concentrating and/or lack of motivation. These mental health concerns or stressful events may diminish your academic performance and/or reduce your ability to participate in daily activities. CMU services are available, and treatment does work. You can learn more about confidential mental health services available on campus at the Counseling and Psychological Services website. Support is always available (24/7) from Counseling and Psychological Services: 412-268-2922.

This semester is unlike any other. We are all under a lot of stress and uncertainty at this time. Attending Zoom classes all day can take its toll on our mental health. Make sure to move regularly, eat well, and reach out to your support system or me if you need to. We can all benefit from support in times of stress, and this semester is no exception.

Respect for Diversity. It is my intent that students from all diverse backgrounds and perspectives be well served by this course, that students' learning needs be addressed both in and out of class, and that the diversity that students bring to this class be viewed as a resource, strength, and benefit. It is my intent to present materials and activities that are respectful of diversity: gender, sexuality, disability, age, socioeconomic status, ethnicity, race, and culture. Your suggestions are encouraged and appreciated. Please let me know if any of our class meetings conflict with your religious observations so that I can make alternate arrangements for you.