Carnegie Mellon University
Master of
Software Engineering

17-[644: Applied Deep Learning]
[M/W 2:20-3:40pm, Online]
[A4, Spring 2021, 6 Units]

Instructor	Email	Office Location & Hours
Prof. Vincent Hellendoorn	vhellend@andrew.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 current systems are quickly starting to include many neural components, and the practice of software engineering itself is starting to benefit from neural program assistance (incl. automated bug finding, translation between programming languages).
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, aimed primarily at applications for and in software engineering itself. Coursework includes regular assignments and a group project 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.

Use of Zoom in the Class. In our class, we will be using Zoom. The link will be made available on the course Canvas page. Please make sure that your Internet connection and equipment are set up to use Zoom and able to share audio and video during class meetings. Let me know if there is a gap in your technology set-up as soon as possible, and we can see about finding solutions.

Sharing video: In this course, being able to see one another helps to facilitate a better learning environment and promote more engaging discussions. Therefore, our default will be to expect students to have their cameras on during lectures and discussions. However, I also completely understand there may be reasons students would not want to have their camera son. If you have any concerns about sharing your video, please email me as soon as possible and we can discuss possible adjustments. Note: You may use a background image in your video if you wish; just check in advance that this works with your device(s) and internet bandwidth.

Assessments & Grading. 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 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. This project will be completed in groups of 2 or 3. These projects will be implemented on Github where we will make use of features such as pull requests and releases.

Assessment	**Final Grade %**
Project milestone 1	20%
Project milestone 2	20%
Project milestone 3	20%
Project final report & presentation	30%
Class participation	10%

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

Partial progress will be assessed at three *milestones* (in weeks 3, 4 & 6). 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. The latter is graded based on the scope and quality of individual contributions and full credit will only be awarded *selectively*, to contributions showing exceptional commitment and exploration. This will be the main opportunity for students to differentiate themselves in terms of the final grade. At a minimum, students are expected to both submit and review one or more Pull Requests (+ associated Issues).

The final presentation and report grade will be equal for the whole group. These will be graded on academic quality, with the frame of reference for a perfect grade being a draft of a conference submission -- indeed, high-quality reports may be shepherded for submission to an ML venue. Throughout this course, students will read a range of deep learning papers, and the final result should reflect this acquired expertise.

- **Milestone 1**: load, preprocess, and document the statistics of the (project-specific) problem data. Train a simple model to assess & report the quality of *naive baselines*.
- **Milestone 2:** improve performance by adopting modern models, regularization, hyper-parameter tuning, and other improvements. Report appropriate metris on final result and *ablations*.
- **Milestone 3:** integrate a task-appropriate advanced feature into the project; demonstrate that this novel combination significantly impacts a property of interest (performance, explainability, model size, etc.), quantitatively and through an analysis of examples.
- **Final presentation & report:** *submit* a 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**, to enrich the discussion 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. For fairness to the whole class, no late submissions will be accepted for the group work. In the first week of classes, you should receive a course schedule for each course; please use them to plan ahead.

  Each student is allowed one late submission for the individual homework assignments. You should immediately notify the course TA(s) before the submission deadline that you will submit late. Late work must be submitted as soon as circumstances allow, ordinarily within 24 hours of the due date. If you have any questions you should raise them immediately rather than waiting for conflicts to arise.

- **Participation policy**. Class participation will be graded by in-class engagement, including asking relevant questions based on a critical review of required readings, lectures, and comments made by your peers. The lack of attendance, and the use of mobile devices, including phones and laptops, will count against your participation grade.

  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.

Attendance. [Tips on writing an attendance policy. You may feel that this is covered in the participation policy in the paragraph directly above. If so, you can delete this section.]

Recording of Class Sessions. All synchronous classes will be recorded via Zoom so that students in this course (and only students in the course) can watch or re-watch past class sessions. Please note that breakout rooms will not be recorded. I will make recordings available on Canvas as soon as possible after each class session (usually within 3 hours of the class meeting). Recordings will live in our Canvas website (https://canvas.cmu.edu/courses/22123)**.** Please note that you are not allowed to share these recordings. This is to protect your FERPA rights and those of your fellow students.

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	3/22	Overview: general DL pipeline, terminology, applications & challenges in practice.
2	3/24	DL basics: brief history of DL (incl. artificial neurons, perceptrons), algebraic basics (vectors/matrices/tensors, stochastic gradient descent) **After class:** sign up for course project topic
3	3/29	Inputs 1a: data mining basics, impact of input statistics (skew), hardware accelerators. *Start of assignment 1*
4	3/31	Architectures 1a: common layers (MLP, CNN, RNN), impact of dimensions/hyper-parameters.
5	4/5	Evaluation 1a: losses (L2, [B/C]CE) and relation to metrics ([top-k] accuracy, precision/recall) **After class:** assignment 1 due
6	4/7	Inputs 1b: layer initializations, embeddings, input scaling, normalization **Prepare:** brief project progress-updates 1 *Start of assignment 2*
7	4/12	Architectures 1b: auto-encoders, encoder-decoders, attention (+ high-level Transformers), pointers
8	4/14	Evaluation 1b: validation sets + bias/variance recap, regularization (dropout, layer/batch-norm), hyper-parameter search/tuning. **After class:** assignment 2 due

9	4/19	Advanced input topics: intrinsic dimensionality, multiple modalities, BPE **Prepare:** brief project progress-updates (3mins each) *Start of assignment 3*
10	4/21	Advanced architecture topics: graphical models, Long-range Transformers, BERT/GPT-X, Bayesian learning.
11	4/26	Advanced evaluation topics: advanced metrics (KLD, top-K loss), beam search, real-world efficacy, explainability/transparency. **Prepare:** brief project progress-updates 2 **After class:** assignment 3 due
12	4/28	Advanced ML Topics: meta-learning, reinforcement learning
13	5/3	Advanced ML Topics ct'd/**Presentations**
14	5/5	**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.