End-to-End Deep Learning (CSC 277/477) - Syllabus
Fall 2024

Instructor & Lecturer: Prof. Christopher Kanan

4 Credit Hours

Catalog Description: Deep learning systems are now being widely productionized at large corporations and many AI-centric start-ups have been created. Productionizing AI systems requires more than just algorithmic considerations. We need to organize the data for training these systems, measure the bias present in these systems after training them, monitor them over time, and more. This course covers these topics, including, but are not limited to, deploying AI systems, MLOps, model versioning, dataset curation, data management, AI ethics/fairness, detecting and mitigation of bias, detecting out-of-distribution inputs, domain shift, data-centric AI, real-time machine learning, continual machine learning, monitoring AI systems after deployment, managing AI projects/teams, and training and inference on edge-devices.

Course Frequency: Offered in Fall

Prerequisites: At least one course that covers neural networks, e.g., CSC 242, CSC 298/578, CSC 266/466), (CSC 249/449), or instructor permission. Students are expected to be familiar with Python, one or more deep learning toolboxes (e.g., PyTorch, TensorFlow, Jax, etc.), and machine learning basics. Students should have at least a high-level understanding of neural networks.

Class Location & Time: Tuesday/Thursday, 3:25 PM - 4:40 PM, Gavett Hall 202

No textbook is needed for this course, but you may wish to refer to these resources:

- Deep Learning Tutorials: https://uvadlc-notebooks.readthedocs.io/en/latest/
- Full-Stack Deep Learning: https://fullstackdeeplearning.com/
- PyTorch Lightning: https://www.pytorchlightning.ai/index.html

Slack Link: https://join.slack.com/t/slack-om29363/shared_invite/zt-2p7qx3a7q-YVB~~twD_VTq~lX~Mks7zQ

Student Outcomes:

1. Demonstrate understanding of the latest advancements in deep learning for real-world applications.
2. Learn about industry-level deployment of deep learning models, especially large models.
3. Gain experience in conducting deep learning R&D.

Instructor Contact:

Name: Prof. Christopher Kanan
Office Hours: TBD Wegmans 3017
Email Address: ckanan@cs.rochester.edu

Illness Policy: All lecture slides will be posted online. If you are feeling unwell, you are encouraged to not attend class or at the very least wear a well fitted mask.

Teaching Assistants & Graders:

Role	Name	Office Hours	Email
Grad TA	Hakki Motorcu	Tuesdays 9AM - 11AM via Zoom.	hmotorcu@UR.Rochester.edu
Grad TA	Junyu Chen	Mondays 10AM - 12 PM via Zoom.	jchen175@ur.rochester.edu

Evaluation and Grading: The final course grade will be weighted as follows:

Homework:	40%
Exams:	30%
Project:	30%

We will follow standard grading guidelines to assign the percentage into a letter grade. The professor may choose to "curve" the class by giving all students the same number of additional points.

Homework: Your homework submissions must cite any references used (including articles, books, code, websites, and personal communications). All solutions must be written in your own words, and you must program the algorithms yourself. You are responsible for starting them early to ensure that you complete them by the deadline. If you start the day before, you will probably do poorly on the assignment. It takes time to implement, train, and evaluate neural networks. Depending on the hardware and implementation, some problems may take hours or more to train the network. All assignments will be done in Python with PyTorch. Your homework solutions must be typed and output to PDF format. Use LaTeX to write up your answers. Your solutions should include all diagrams, written explanations, code, and program outputs.

Planned Assignment Topics (subject to change):

- Homework 0: A quick assignment to assess prerequisite knowledge. The other assignments will be longer.
- Homework 1: Testing for deep learning models, MLOps, experiment tracking.
- Homework 2: LLMs
- Homework 3: Memory Efficient Training / Distributed Training Methods
- Homework 4: Model calibration, OOD detection / selective classification, continual learning.
- Homework 5: Data-centric AI

Project: You are required to complete a project related to the topics covered by this course. For CSC 477, projects should be done individually, but for CSC 277 you may have up to 3 team members. Exceptions to team group limits may be given with strong justification. You may use the programming language and toolboxes of your choice. Your write-up should rigorously evaluate your project. Run your early ideas by Prof. Kanan and other staff via email or in person. The schedule for the project is as follows:

1. **Project Proposal**: The project proposal should clearly state what you plan to do. It should be at least 3 pages long (not including references). It should contain a list of three to six milestones

and deadlines. You should list what software you will be using or will build upon. Describe the datasets you will use and how will you know if the project is successful. Describe the related work. The proposal should be a well organized document in continuous English, and it should not be merely an outline. You should be able to reuse much of the text for the final report. It should be submitted as a PDF (under 10MB).

2. **Revised Project Proposal (optional)**: The revised proposal is an opportunity to improve your grade if you fail to do the project proposal effectively. You may submit a revised proposal that takes into account the comments received by the instructor and TA. The new grade will replace the original score, but the maximum score for the revised proposal is 80%.
3. **Project Report**: The project report will describe the project, i.e., what you did and the result. *It should be at least 6 pages long (not including references and figures)*. The report should be formatted in CVPR format (exceptions will be granted if you are submitting to a conference or journal). It should be submitted as a PDF (under 10MB).

Policy on Late Work: No credit will be given for the project report if it is turned in late. For the project proposal, a late submission will be treated as submitting a revised project proposal capping the score at 80%. For homework assignments, full points will be awarded only if the assignment is turned in at most one day late. Late homework assignments will be accepted up to 7 days late with a 20% penalty imposed, meaning the highest possible score will be 80% for any assignment that is 2-7 days late. No credit will be given for assignments turned more than 7 days late. An exception to this policy is that no assignments will be accepted after the project report deadline.

Programming Environment: For homework assignments, this course uses Python and PyTorch (Lightning). For the class project, you may use the programming language and framework of your choice, but most of our expertise is in PyTorch.

Academic Honesty and Integrity: All assignments and activities associated with this course must be performed in accordance with the University of Rochester's Academic Honesty Policy. You are expected to read, understand, and follow the policy. Additionally:

- In general, homework is to be completed independently. However, you are encouraged to study together and to discuss information and concepts covered in lecture and the sections with other students. You can give "consulting" help to or receive "consulting" help from such students. However, this permissible cooperation should never involve one student having possession of a copy of all or part of the work done by someone else. Should copying occur, both the student who copied work from another student and the student who gave material to be copied will both automatically receive a zero for the assignment. Penalty for violation of this Code can also be extended to include failure of the course and University disciplinary action.
- Posting homework and project solutions to public repositories on sites like GitHub is a violation of the College's Academic Honesty Policy, Section V.B.2 "Giving Unauthorized Aid."
- During examinations, you must do your own work. Talking or discussion is not permitted during the examinations, nor may you compare papers, copy from others, or collaborate in any way. Any collaborative behavior during the examinations will result in failure of the exam, and may lead to failure of the course and University disciplinary action.

Prior Course Materials: Unauthorized use of course materials from previous semesters (e.g., material you have received from others), is strictly prohibited.

New Course Materials: Course materials (slides, lectures, assignments, etc.) may not be re-distributed or posted elsewhere online. Redistribution of copyright protected material outside this course may be prohibited by law.

Notes on Plagiarism: Plagiarism is a serious offense and is in violation of university policy.

- If you are unsure of what constitutes plagiarism in written documents, a good description can be found here: https://rochester.edu/college/gradstudies/assets/pdf/Plagiarism_Misconduct.pdf
- Plagiarism does not just occur in written documents; it also occurs in code. Many of the algorithms we will code and problems we will solve have been solved by others who have posted code (in various programming languages) online. It is unacceptable (and it is considered plagiarism) to copy code developed by others and submit it as your own. (This includes code that is written by your fellow students!) Even making minor changes, such as changing variable names, function names, formatting, etc., is not enough to allow you to claim your submission as your own because the underlying structure of the code remains unchanged. You may also be in violation if you excessively rely on AI "Co-Pilot" systems to assist you with writing your code.
- If you do consult any online sources of code, you must properly attribute the corresponding sections in your code to their original source, as you would add quotations, footnotes, or references in a written document. The consequences of plagiarism, whether in code or in written documents, are at the discretion of the instructor, and can be as severe as automatic failure of the course.

Academic Accommodations: We are committed to providing reasonable accommodations to students with disabilities. Please see the professor about your required accommodations as early as possible in the term. The University of Rochester respects and welcomes students of all backgrounds and abilities. In the event you encounter any barrier(s) to full participation in this course due to the impact of a disability, please contact the Office of Disability Resources. The access coordinators in the Office of Disability Resources can meet with you to discuss the barriers you are experiencing and explain the eligibility process for establishing academic accommodations. You can reach the Office of Disability Resources at: disability@rochester.edu; (585) 276-5075; Taylor Hall; www.rochester.edu/college/disability.

Course Schedule: The following schedule lists dates for class topics. *The content in this schedule is tentative and subject to change.* It is your responsibility to attend class and to remain informed of any changes that may be announced.

Week	Date	Assignments	Class / Discussion Topics	Presenter
1	8/27	Homework 0 Assigned	Course Introduction	Kanan
	8/29		Review: Linear Algebra, Feedforward Networks, & Backpropagation	Kanan
2	9/3		Review: Neural Network Architecture Zoo & Convolutional Neural Networks	Kanan
	9/5	Homework 0 Due (9/9)	MLOps & Experiment Tracking	Kanan
3	9/10	Homework 1 Assigned	Foundation Models & Transformers	Kanan
	9/12		LLMs Part 1	Kanan
4	9/17		LLMs Part 2	Kanan
	9/19	Homework 2 Assigned	LLMs Part 3	Kanan
5	9/24		LLMs Part 4	Kanan
	9/26		Maximizing GPU Utilization and Mixed Precision Training	Kanan
6	10/1	Homework 3 Assigned	Activation Checkpointing & Intro to Distributed Model Training	Kanan
	10/3	Homework 1 Due (10/4)	Distributed Model Training: Data parallel, Model parallel, Sharding	Kanan
7	10/8		Exam 1 Review	Kanan
	10/10		**Exam 1**	N/A
8	10/15	Project Proposal Due (10/20)	No Class - Fall Break	N/A
	10/17	Homework 2 Due (10/21)	Llama 3 Deep Dive	Kanan
9	10/22	Homework 4 Assigned	AI Ethics & Safety	Kanan
	10/24		Hypothesis Testing and Confidence Intervals for ML Models Bias Mitigation Methods	Kanan
10	10/29		Uncertainty Estimation, Model Calibration, & OOD Detection	Kanan
	10/31	Homework 3 Due (11/4)	Conformal Prediction	Kanan
11	11/5	Homework 5 Assigned	Data-Centric AI Part 1	Kanan
	11/7		Data-Centric AI Part 2	Kanan
12	11/12		Model Monitoring & Continual Learning	Kanan
	11/14	Homework 4 Due (11/18)	No Class - CVPR Deadline / Prof. Kanan Traveling	N/A
13	11/19		Writing AI Papers	Kanan
	11/21		AI Startups, ML Teams & Project Management	Kanan
14	11/26		Final Lecture / Course Wrap-up	Kanan
	11/28		No Class - Thanksgiving	N/A
15	12/3		Exam 2 Review	Hakki & Junyu
	12/5		**Exam 2 - Last In-Class Event**	N/A
	12/9	Homework 5 Due		N/A
	12/13	Project Report Due	No extensions	N/A