Syllabus - DS677 Deep Learning

Spring 2025

Code: DS677
Time: Friday, 2:30 PM - 5:30 PM
Location: CKB 124
Mode: Face-to-Face

Instructor: Hai Phan
Office: GITC 2105
Email: phan@njit.edu
Office Hours: Friday **(**9:30am-noon). If these hours do not work with your schedule, appointments are also available by email.

Teaching Assistant/Grader: TBD

Note: Your messages will be answered by the end of the next day. Grades for all items will be getting posted during the week after their due date. For issues with your grades, contact the grader and cc the instructor.
Tutoring. NJIT provides a tutoring service. Please contact one of the available tutors. Please check the website for updates regularly as they may change the information.

General Information

Prerequisites/Co-requisites
CS675 or DS675 or instructor permission.

Course Description
This course covers current topics in data science. These include but are not limited to, parallel programming on GPU and CPU multi-cores, deep learning, representation learning, optimization algorithms, and algorithms for big datasets. Students will present recent papers in data science, work on programming assignments, and complete a machine learning/deep learning/data science project.

Extended Course Description
Deep Learning (DL) is a subfield of Machine Learning that has delivered disruptive technologies and created AI algorithms that outperform humans in various tasks. It paves the way for broader advances in science. DL consists of specialized techniques that exploit the abundant availability of data and computational power to build models composed of multiple processing layers and learn data representations at various levels of abstraction. Only a few years back, the development of DL models required significant experience. Still, introducing open-source DL libraries like TensorFlow and PyTorch has opened the area to scientists and professionals with more diverse backgrounds.

The course begins with a review of Artificial Neural Networks that guides you through PyTorch and enables you to build novel ANN architectures. Then, it presents the evolution of progressively deeper

architectures for Convolutional Neural Networks that addressed various training difficulties and led to very successful image classification models. In this spatial context, you will learn about Generative Adversarial Networks behind the fascinating "Deep Fake" images and videos. The course then takes you to the emerging applications of Recurrent Neural Networks in temporal data, including Natural Language Processing. You will also learn about Graph Neural Networks and their applications in analyzing real-world networks. The course may also touch upon selected topics like the ability of deep networks to generalize, techniques for "pruning" deep networks to make them more computationally efficient, and successful applications of DL methods in the Sciences.

Course Learning Outcomes

By the end of the course, students will be able to:

1. Program in widely used parallel frameworks for Deep Learning (DL)
2. Recognize problems amenable to DL methods
3. Describe and explain a wide variety of DL methods for various data types
4. Adapt existing DL resources to novel data and applications
5. Appreciate and evaluate new developments in the field of DL
6. Appreciate the broader impact of DL in the Sciences

Required Materials

There is no required course textbook. The course will draw material from several sources, including the instructor's own notes. Some open resources include:

- A. Zhang , Z. Lipton, M. Li, A. Smola. *Dive into Deep Learning*
- Prince, S.J. *Understanding Deep Learning*. MIT press. (Dec 5, 2023)

Grading Policy

The grading policy is designed to reflect the NJIT Grading Legend.

Final grades for all assignments will be based on the following percentages:

Quizzes (Short Quizzes - 10% - Midterm Exam 15%)	**25%**
Class participation	**10%**
Homeworks	**30%**
Projects Milestone-1 (10%), Milestone-2 (5%) Milestone 3 (20%)	**35%**

Letter to Number Grade Conversions

Raw numerical scores will be converted to letter grades using the following bounds.

A	B+	B	C+	C	F
≥93	≥85	≥70	≥60	≥50	<50

Course Work

Quizzes (10% of grade): Weekly Canvas quizzes reinforcing the material of each module, will help you keep up with the most important theoretical concepts. These quizzes are not proctored.
Class Participation (10% of grade): You are expected to attend classes and participate in classes by listening and understanding class contents and asking related questions. You are also expected to participate in weekly Canvas discussions prompted by the instructor, with meaningful questions and answers related to the week's topics or assignments.
Homeworks (30% of grade): Assignments will be given biweekly to give you an opportunity to apply course concepts for that week. Four homework assignments of equal grading weight.
Mid-term Exam (15% of grade): In-person exam, 90 minutes. Students are expected to bring a fully charged laptop, as the exam will be on Canvas with LockDown browser. Each student is allowed to bring at most 5 pages of notes. In the event the exam has to take place online, Respondus Monitor will be used for proctoring.
Project (35% of grade): The project will consist of three milestones, with weights [10%, 5%, 20%]. You will have opportunities to iterate and revise your work based on peer, TA, and instructor's feedback.

Course Topic Schedule (tentative)

The topics covered in this course include the following, presented in the approximate order in which they will be taught. This list of topics is to be considered as a *reference* that can be adjusted through the course of the semester to address changing needs.

Week	Date	Topic	Textbook Readings	Due Work
Week 1	9/6	Introduction and course overview	chapter 2 chapter 3	
Week 2	9/13	Basic Building blocks in Deep Learning (Softmax Regression, Regularization, Dropout, Initialization, optimization, etc.)	chapters 4,5 chapters 3,5	hwk #1
Week 3	9/20	Convolutional Neural Networks (CNNs)	chapter 7	project initiation
Week 4	9/27	Modern CNNs	chapter 8	hwk #2
Week 5	10/4	RNNs and Language Modeling	chapter 9	
Week 6	10/11	Modern RNNs and Language Translation	chapter 10	hwk #3
Week 7	10/18	Transformers	chapter 11	project milestone #1
Week 8	10/25	Midterm exam		Summary quiz
Week 9	11/1	Vision Transformers	chapter 12	hwk #4 project milestone #2
Week 10	11/8	Selective State Space Models		

Week 11	11/15	Graph Neural Networks	online tutorial; chapter 14	hwk #5
Week 12	11/22	Reinforcement Learning	chapter 17	project milestone #3
Week 13	11/27	Deep Generative Modeling	chapter 20	hwk #6
Week 14	12/6	New Frontiers in Deep Learning		
Week 15	12/13	Project Presentation		Project presentation

* All homework quizzes (multiple attempts), and milestones are due on Sunday at 23:59.

Feedback

Assignment solutions will be distributed for each assignment, along with general class-level feedback from the grader. Occasionally, and when needed, you will receive individualized comments directly on your assignment notebook. You can also always directly inquire about a specific grade item. In that case, please email both the instructor and the grader.

Policy for Late Work

When a student is unable to complete an assignment or for other serious reasons, these must be communicated and documented promptly. In any other case, each hour of delay after the due date will incur a 2% score reduction. No extensions will be granted. However, each student's lowest programming score and the two lowest quiz assignment scores will be dropped.

Exam Information and Policies

This course has two proctored quizzes. These will take place in the classroom, and presence is required. The quizzes will be on LockDown browser, so please make sure you bring your computers charged. Most of your grade is based on authentic assessment, meaning that you will be assessed and graded on your ability to deliver real-world outputs and your participation and feedback to other students.

Collaboration and External Resources for Assignments

Some of the assignment problems will be challenging. You are advised to try to solve all the problems first. For issues that persist, you are welcome to talk to the instructor during office hours or raise questions in the Weekly Discussion Forum. In consulting with others, you can only exchange general ideas and approaches unless you are given explicit permission to do so in the assignment statement; the complete solutions must be worked out by you alone.

Generative AI Tools and Other External Resources

Sometimes, you may come across code, text, or other helpful information online, or you may be able to generate it using AI tools such as ChatGPT or other Large Language Models (LLMs). In most cases, you will be allowed to integrate this information into your solution. However, if you do, you must always give the appropriate credit and citations (e.g., links) for the material you use (especially when you use the code and text you found online). If you use an LLM, you must say that you did so and present the entire transcript of your ‘conversation’ with it, which should show what you asked and how you guided it or were guided by it to the delivered solution. Your ‘conversation’ with it must be entirely yours and sufficiently different from that of other students. Failure to give appropriate credit when using the work of others (whether human or AI) is considered plagiarism and may lead to disciplinary action under NJIT's Academic Integrity policy.

Academic Integrity

"Academic Integrity is the cornerstone of higher education and is central to the ideals of this course and the university. Cheating is strictly prohibited and devalues the degree that you are working on. As a member of the NJIT community, it is your responsibility to protect your educational investment by knowing and following the NJIT academic code of integrity policy.

Please note that it is my professional obligation and responsibility to report any academic misconduct to the Dean of Students Office. Any student found in violation of the code by cheating, plagiarizing or using any online software inappropriately will result in disciplinary action. This may include a failing grade of F, and/or suspension or dismissal from the university. If you have any questions about the code of Academic Integrity, please contact the Dean of Students Office at dos@njit.edu"

Weekly Expectations

The course is organized into modules. Each week consists of one module. Each week, the students should attend the week's lecture. The students are also expected to read the corresponding sections of the textbook, and participate in a class discussion forum as prompted by the instructor. The students must also be aware of any assignments due at the end of each week.

Additional Information and Resources

Netiquette

Throughout this course, you are expected to be courteous and respectful to classmates by being polite, and active participants. You should respond to discussion forum assignments in a timely manner so that your classmates have adequate time to respond to your posts. Please respect opinions, even those that differ from your own, and avoid using profanity or offensive language.

Accessibility

This course is offered through an accessible learning management system. For more information, please refer to Canvas's Accessibility Statement.

Requesting Accommodations

The Office of Accessibility Resources and Services works in partnership with administrators, faculty, and staff to provide reasonable accommodations and support services for students with disabilities who have provided their office with medical documentation to receive services.
If you are in need of accommodations due to a disability, please contact the Office of Accessibility Resources and Services to discuss your specific needs.

Resources for NJIT Online Students

NJIT is committed to student excellence. To ensure your success in this course and your program, the university offers a range of academic support centers and services. To learn more, please review these Resources for NJIT Online Students, which include information related to technical support.