Deep Learning - DS677 - Syllabus

Spring 2024

Course Modality:
Online (Virtual Campus). For more information on using Canvas and other supported learning tools, visit the IST Service Desk Knowledge Base.

Instructor Information

Instructor	Email	Office Hours
Ioannis Koutis	ikoutis+677@njit.edu	Tuesday 5 pm -7 pm: online slots (please reserve a few hours ahead) Class meet up: (evening time)

*I will respond to all emails/Inbox messages within 48 hours.
Quizzes, homework, and discussions will be graded weekly.

Teaching Assistant: Hamideh Sabaei (hs833@njit.edu)

General Information

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

Course Description
This course covers current topics in data science. The topics 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 do 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 a set of specialized techniques that exploit the abundant availability of data and computational power to build models that are composed of multiple processing layers and learn representations of data at multiple levels of abstraction. Only a few years back, the development of DL models required significant experience, but the introduction of open-source DL libraries like TensorFlow and PyTorch has opened the area to scientists and professionals with more diverse backgrounds. The course opens 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 also learn about Generative Adversarial Networks that are 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 the analysis of real-world networks (e.g., social, or biological 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

Dive into Deep Learning
A. Zhang, Z. Lipton, M. Li, A. Smola
The textbook is open, free, and available here.

Grading Policy

The grading policy is designed to reflect the NJIT Grading Legend

Final Grade Calculation

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

Quizzes (Short Quizzes - 10% - Final Summary Quiz 15%)	**25%**
Discussion Forums and Misc Participation	**15%**
Exercises/ Learning Activities	**25%**

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

In some cases a letter grade can be upgraded to the next letter to reflect natural clusters of performance.

Course Work

Assignment and Projects

Quizzes: (25% of grade) There will be weekly short multiple-choice quizzes, worth 10% of the total grade. These are meant to help you keep up with the most important theoretical concepts. These quizzes are not proctored and the two weakest scores will be dropped. There will be one 90-minute summary quiz, worth 15%. This is required and proctored. It is meant to simulate an interview environment and assess your overall understanding of the material.

Discussion Forums and Participation Activities: (15% of grade) When all students participate in a discussion, it creates an active learning environment that will help you better understand the materials and be more successful in the class. You are expected to participate in two types of forums: (i) Weekly discussion forums in Canvas, with Q&A about the week's material (9%). Your contributions are due by Sunday, 11:59 pm. (ii) A permanent discussion forum on successful applications of the material we cover (5%). Other participation activities are worth 1%.

Exercises/Learning Activities: (25% of grade) Assignments will be given bi-weekly (up to week #12) to give you an opportunity to apply course concepts for that week. These activities are designed to help you practice and prepare for the project. The weakest grade will be dropped automatically.

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 and instructor feedback.

Feedback

Assignment solutions will be distributed for each assignment, along with general class-level feedback from the grader. Occasionally, and when needed, you will also 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

Generally speaking, assignments, project milestones, and participation exercises will be accepted late without penalty, but only up until the time grading has begun, or solutions or other feedback have been released to members of the class. At that time, the class will be informed that submissions have closed, that no further submissions will be accepted and that any missing work will be given a mark of zero. Students should be aware that grades, solutions, and/or feedback may be released at any time after the deadline, at the sole discretion of the lecturer and without prior warning. The only way to ensure that work will be accepted without penalty is to submit it before the deadline. No individual extensions will be granted. However, special consideration may be given in rare cases when a student is unable to complete an assignment for serious, unavoidable reasons — these must be communicated and documented promptly.

Exam Information and Policies

This course has at most two proctored quizzes. The quizzes will be on the LockDown browser with web proctoring. The final exam has a set date and time by the University's registrar. The majority of your grade is based on authentic assessment, meaning that you will be assessed and graded on your ability to deliver real-world outputs as well as your participation and feedback to other students.

Collaboration and External Resources for Assignments

Some of the assignment problems will be quite challenging. You are advised to first try to solve all the problems on your own. For problems 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 are allowed to exchange general ideas and approaches only: unless you are given explicit permission to do so in the assignment statement, the full solutions themselves 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). In the case 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 1 or 2 modules. Each week, the students should attend the week's lecture. Whenever that is not possible, the students are advised to watch the posted videos. 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.

Course Topic Schedule

Week	Topic	Textbook Reading	Due Work
1	• Introduction to CS677 • Review: Optimization and Pytorch	chapter 1 chapter 3	
2	• Softmax Regression • Regularization, Dropouts, Initialization	chapters 4,5 chapters 3,5	hwk #1
3	• Introduction to CNNs	chapter 7	project initiation
4	• Modern CNNs	chapter 8	hwk #2
5	• Introduction to RNNs and Language Modeling	chapter 9	
6	• Modern RNNs and Language Translation	chapter 10	hwk #3
7	• Attention - Transformers	chapter 11	project milestone #1
8	• Graph Neural Networks	online tutorial	hwk #4
9	• Review Week	chapter 12	summary quiz project milestone #2
10	• Applications in Vision	chapter 14	hwk #5
11	• Word Embeddings	chapter 15	
12	• NLP applications	chapter 16	hwk #6

13	• Recommender Systems	chapter 17	project milestone #3
14	• Various Topics in Deep Learning	chapter 20	broader impacts of DL
15	• Final Week		peer feedback on projects

- Homeworks become available two weeks before their due date
- Homeworks and milestones are due on Sunday, at 23:55
- There are also multiple-attempts weekly quizzes due on Sunday at 23:55

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.