



# Deep Learning - DS677 - Syllabus

Fall 2025

## General Information

*Course Number:* DS677; *Section:* 001; *CRN* 92421

*Mode of Instruction:* In Person/Face to Face

*Time:* Monday/Wednesday 11:30 AM - 12:50 PM

*Location:* Cullimore Hall Lecture 2

Instructor	Email	Office Location	Office Hours
Akshay Rangamani	<a href="mailto:akshay.rangamani@njit.edu">akshay.rangamani@njit.edu</a>	GITC 2112	By appointment - please email to setup Tue/Thu 4PM - 5:30 PM via zoom link

TA - TBD

\*I will respond to all course related emails within 48 hours. Please include **DS677** in the subject line to ensure I see it.

---

## About the Course

### Prerequisites/Co-requisites

CS675 or DS675 or instructor permission. Python programming knowledge. A background in Calculus, Linear Algebra, Probability, and Statistics is assumed. Hopefully this is not your first Machine Learning course.

### Course Description

Over the past decade, deep learning has tackled problems previously assumed to be the exclusive domain of biological organisms. Computers can (with appropriate caveats) identify objects,

transcribe speech, translate from one language to another, and converse with humans. Deep learning has powered all of these achievements. This course is designed to introduce students to the fundamentals of deep learning. We will start with an overview of the basic workhorses of deep learning including backpropagation, first order stochastic optimization algorithms, and regularization methods in the context of supervised learning. We will then tour widely used architectures in modern deep learning including convolutional networks, recurrent networks, and transformers, studying both the design principles as well as implementation. We will briefly delve into the challenges of developing a mathematical theory of deep learning. Having laid this foundation, we will then switch tracks to study the use of deep networks in representation learning and generative models. The course will end with an overview of some applications of deep learning. Quizzes and programming assignments will provide students with an opportunity to implement what they have learned in class. Students will also have the opportunity to complete a deep learning project over the course of the semester.

### **Learning Outcomes**

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

1. Implement different architectures in deep learning frameworks
2. Identify the type of architecture to use for a given type of data
3. Explain the role of initialization, regularization, and choice of optimization algorithms in training deep networks.
4. Debug failures in training deep networks
5. Propose deep learning solutions to a wide range of problems

### **Course Materials**

While we will not stick to a specific textbook, the course material will be based on the following books.

#### *Dive into Deep Learning*

A. Zhang, Z. Lipton, M. Li, A. Smola

The textbook is open, free, and available [here](#).

#### *Deep Learning*

I. Goodfellow, Y. Bengio, A. Courville

The textbook is open, free and available [here](#).

### **Coursework and Assessment**

**Class Participation: (5% of grade)** Students are expected to attend and participate in all classes.

**Quizzes: (10% 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.

**Homeworks: (40% of grade)** Homeworks will be given bi-weekly (up to week #10) 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. There will be 5 homeworks each worth 8%.

**Midterm (15% of grade)** There will be one in-class midterm exam (80 minutes), tentatively scheduled for week 8. This is intended to test your understanding of the foundational architectures and techniques used in deep learning.

**Project: (30% of grade)** Students will work on a final project that will demonstrate their proficiency in deep learning. Students are expected to work in groups of at most 2 people on an idea of their choice that uses deep learning in an innovative way. There will be a date early in the semester where students will have to submit a project proposal outlining the idea behind and scope of their project. Students will also prepare a short presentation to explain their approach to the project midway through the semester. Projects will be evaluated in two ways - through an in-class presentation, and a final report (5 pages, NeurIPS style). More details about the logistics of the project, as well as resources and ideas for course projects will be provided.

## Course Policies

**Collaboration:** You are allowed to discuss the homeworks with anyone in the class (including the instructor). However, all code and written work must be your own and you are expected to understand every line of what you submit. Please acknowledge your collaborators by including their names at the beginning of your submission.

**AI Assistants:** Since this is a deep learning course, you are encouraged to try and use the plethora of AI assistants that are out there. Probing their capabilities and limitations is part of your learning experience. However, when it comes to their role in your coursework, please use your common sense. The policy for using AI assistants is identical to the policy for collaborating with humans. You can converse with an agent to discuss topics in the course, or to help you understand an idea better. You cannot ask an AI assistant to solve a problem in an assignment for you (just like you cannot ask your friend to do your homework for you). If you do use an AI assistant during the course of an assignment or writing your project report, please describe how you used the agent in a few sentences. Remember that any work that you turn in is assumed to be your own, including any inconsistencies or hallucinations from AI assistants.

**Late Homeworks and Quizzes:** You have a total of 10 late days to use throughout the semester for submitting your homeworks and quizzes. These can be used at your discretion, without notifying the instructor. Beyond that, you will be docked 20% for each day that you are late. These are meant to cover the normal circumstances of life (all my courses have the homeworks due on the same day, I need to attend a conference, I forgot that my homework was due today). The late days can only be used for homeworks and quizzes and may not be used for the other components of the coursework.

**General:** Please do not hesitate to approach me with any concerns you may have about the course policies or resources that you may require to succeed in the course. I am invested in ensuring that every student in the class gets what they need out of the course, and will try my best to work through any issues, or atleast point you to the appropriate personnel.

## Tentative Course Schedule

Class	Date	Topic	Reading Material	Homeworks and Project Deadlines
0	9/3	Introduction, Logistics, and Historical notes, Linear Algebra, Probability, Convex Optimization	Introduction (Chapter 1) of <a href="#">Patterns, Predictions, Actions</a> Chapter 22 of D2L	
1	9/8	MLPs, Backpropagation	Section 2.5 of Dive into DL, Section 6.5 of Deep Learning Textbook	HW0 out
2	9/10	Activation functions, Losses, Initialization	Chapter 5 of Dive into DL, Sections 6.2, 6.3, 6.4, 8.4 of Deep Learning Textbook	
3	9/15	Optimization	Chapter 12 of Dive into DL, Chapter 8 of Deep Learning Textbook <a href="https://distill.pub/2017/momentum/">https://distill.pub/2017/momentum/</a> <a href="https://www.ruder.io/optimizing-gradient-descent/">https://www.ruder.io/optimizing-gradient-descent/</a> Convex Optimization for ML - <a href="https://ee227c.github.io/">https://ee227c.github.io/</a>	HW0 due; HW1 out
4	9/17	Regularization and Normalization	Sections 3.6, 3.7, 4.6, 5.5 of Dive into Deep Learning, Chapter 7 of Deep Learning textbook	
5	9/22	Convolutional Layers	Chapter 7 of Dive into Deep Learning, Chapter 9 of Deep Learning <a href="#">A guide to convolution arithmetic for deep learning</a>	
6	9/24	Modern Convolutional Neural Networks	<a href="#">Class 6 Worksheet</a> Chapter 8 of Dive into Deep Learning	HW1 due; HW2 out
7	9/29	Recurrent Neural Networks	Sections 9.4-7, 10.1-4 of Dive into Deep Learning, Chapter 10 of Deep Learning <a href="#">Supervised Sequence Labeling with RNNs</a>	Submit Project Proposal
8	10/1	Language Modeling and RNNs	Sections 9.2, 9.3, 10.8 of Dive into Deep Learning <a href="#">Hugging Face Transformers</a>	

9	10/6	Seq2Seq, Translation and Attention and Memory Mechanisms	Sections 10.5-7, 11.1-5 of Dive into Deep Learning	HW2 due; HW3 out
10	10/8	Transformers	Sections 11.6 - 9 of Dive into Deep Learning  <a href="#">The annotated transformer</a> , <a href="#">The illustrated transformer</a>	
11	10/13	Modern Transformer Architectures	<a href="#">Class 11 Worksheet</a>	
12	10/15	Pre-Midterm Review		HW3 due
13	10/20	<b>In-Class Midterm</b>		
14	10/22	Post-Midterm Review		
15	10/27	Adversarial Attacks + Robustness		
16	10/29	Project Scope Presentations		
17	11/3	Representation Learning		
18	11/5	Representation Learning		HW4 out
19	11/10	Theory of Deep Learning		
20	11/12	Theory of Deep Learning		
21	11/17	Generative Modeling		HW4 due
22	11/19	Generative Modeling		
23	11/24	Graph Neural Networks		
24	12/1	Diffusion Models		
25	12/3	State Space Models		
26	12/8	Final Project Presentations		
27	12/10	Final Project Presentations		
	12/17			Final Project Reports Due

## Academic Support System

NJIT Academic Calendar <https://www5.njit.edu/registrar/calendars/> . This contains drop/add and other important dates.

### IST Service Desk

The IST Service Desk is the central hub for computing information and first point of contact for getting help and reporting issues related to computing technology at NJIT. Students can put in a ticket with the service desk: <https://servicedesk.njit.edu/CherwellPortal/IST> or call (973) 596-2900 Monday - Friday from 8:00am – 9:00pm

### Academic Advising Success Center

Is for undergraduate students only. Ying Wu College of Computing has its own advising team.

### 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 academic code of integrity policy that is found at: <http://www5.njit.edupoliciesitespoliciesfiles/academic-integrity-code.pdf>.

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](mailto:dos@njit.edu)

### Academic Support and Student Affairs

“From questions about becoming a student at NJIT – to student engagement – to searching for information on career development, the Division of Academic Support and Student Affairs Staff is here to help.”

### Office of Accessibility Resources and Services

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.

[Additional Tutoring Centers](#)

[The Writing Center](#)

[Bookstore](#)

[Canvas Help Page](#)

[Center for Counseling and Psychological Services](#)

“The NJIT Center for Counseling and Psychological Services (C-CAPS) is committed to assisting students in the achievement of their academic goals as well as benefiting from their personal experience on campus. College life can be personally challenging and stressful at times. We believe that the educational process is an important component of the development of the individual as a whole person. Our goal is to optimize the college experience and improve the quality of the lives of our students by promoting their mental health and facilitating students' personal, academic and professional growth.”

[The Learning Center](#)

“Our mission is to assist students both in the classroom and beyond by providing tutorial services, academic coaching, academic and personal enrichment workshops and staff and peer support so students can meet the demands of their coursework and are prepared for life after graduation.”

[Robert W. Van Houten Library](#)

“The Van Houten Library offers electronic and print resources essential to the mission of New Jersey's science and technology university, including a core collection of academic books, databases, and journals, as well as research and consultation services.”

[Student Disability Services](#)

“The Disability Support Services office works in partnership with administrators, faculty and staff to provide reasonable accommodations and support services for students with disabilities that have provided our office with documentation to receive services.”

[Student Financial Aid Services](#)

“Student Financial Aid Services (SFAS) at NJIT is committed to providing you with every opportunity to obtain funding to support your undergraduate educational costs at NJIT.”