DS 542 Deep Learning for Data Science

Fall 2025 Syllabus

Official Course Description

In this course students will gain an understanding of the fundamentals in deep learning and then apply those concepts in exercises and applications in python. We'll start with the origins of artificial neural networks, learn about loss functions, understand gradient descent, back propagation and various training optimization techniques. Students will be familiar with canonical network architecture such as multi-layer perceptions, convolutional neural networks, recursive neural networks, LSTMs and GRU, attention and transformers. Through explanations, examples and exercises students will build intuition on how deep learning algorithms work and how they are implemented in popular deep learning frameworks such as PyTorch. Students will be able to define, train and evaluate deep learning models as well as adapt deep learning frameworks to new functionality. Students will gain exposure to pre-trained large language models and other foundation models and the concepts of few-shot learning and reasoning. Finally, students will be able to apply many of the techniques they learned in a final class project.

Learning Outcomes

Upon successful completion of this course, students will be able to:

1. Grasp foundational theories and practices in the deep learning arena.
2. Design and implement various neural network architectures using Python and PyTorch.
3. Employ regularization, optimization, and advanced training techniques to enhance model performance.
4. Analyze real-world datasets, applying suitable deep learning techniques to derive actionable insights.
5. Understand the benefits and drawbacks of various neural architectures in specific contexts.
6. Understand the pros and cons of pre-trained large language and other foundation models and how best to employ them
7. Complete multiple data-centric projects, showcasing end-to-end deep learning implementation.

Prerequisites

- Python Programming – Should be proficient in python and associated data science packages, or studiously working towards proficiency. See for example Scientific Python Lectures for lessons on python language and relevant packages, or The Python Tutorial for a tutorial on the core language.

- Packages such as NumPy NumPy - Learn and SciPy (SciPy User Guide have tutorials and documentation as well.
- The more proficient you are, the more effective you will be at the assignments and projects. We will dedicate some discussion sessions to ensure your environment is set up correctly and review some of the basics as well as answer any questions.
- Math Proficiency – In order to understand the foundational concepts, it is important to have proficiency in a number of areas of mathematics. These include linear algebra, first year calculus and trigonometry as well as some concepts from Real Analysis. We will cover these concepts in the class and some recitation sessions, but refreshing or building your foundation will help.

Weekly Cadence

- Lectures: Monday/Wednesday 2:30-4:15pm (STH B19)
- Discussions: Wednesday 12:20-1:10 (CDS 164), 1:25-2:15 (STH 113)
- Instructor: Jeffrey Considine
    - Office: CDS 1645
    - Office Hours: Monday 11-12, Tuesday 3-4, Wednesday 1-2, Thursday 1-2
    - email: jconsidi@bu.edu
- Teaching Assistant: Carrie Feng
    - Office Hours: Tuesday 11-12, Friday 11-12
    - Office Hours Location: TBD
    - email: carrieff@bu.edu

Course Infrastructure

- Piazza for course announcements and questions.
- Gradescope for assignments and exams (entry code KDXJWJ). All regrade requests must be submitted through Gradescope within a week of grades being posted.
- You are expected to have a personal laptop that you can bring to discussion ready for coding and other basic tasks. If this is an issue, reach out to the instructor as soon as possible.

Textbooks and Reference Materials

- The primary textbook for this course will be Understanding Deep Learning, by Simon Prince. The book is available online.
- We'll also reference Introduction to Linear Algebra, Sixth Edition (2023), by Gilbert Strang.
- Slides will be posted online as well.
- Given the fast moving nature of this area, we'll also be citing many articles available online as well as other online reference materials in each lecture. As part of the class,

we will guide the students in constructing their own bibliography and give tips on how to efficiently and effectively read research papers.

Course Requirements

- Homework: For the first third of the class, there will be weekly Jupyter Notebook coding assignments and homework problems to help anchor the key concepts and python/pytorch coding patterns.
- Projects: For the last two thirds of the classes, there will be biweekly projects assigned where you will build models and systems similar to what we have covered in lecture.
- Discussions: You will practice the concepts covered in class and turn in a Jupyter notebook or answer questions at the end.

Course Assessment

- Homeworks: 25%
- Projects: 40%
- Discussions: 35%

Late submissions will usually be allowed for up to two days, but will be penalized 1% per hour based on last submission times recorded by Gradescope. This penalty will be applied separately from the raw scores that you see in Gradescope and Blackboard. See Gradescope for assignment-specific deadlines.

Most assignments will be focused on implementing techniques covered in class, but you will sometimes be asked questions with text answers. For example, you may be asked to explain, motivate or otherwise argue for an approach. In those cases, you are expected to give a concise and direct answer and not be unnecessarily verbose. Points may be deducted for poorly written responses.

Schedule

This schedule is subject to change, particularly towards the end of the course.

Date	**Lecture Topic**	**Assignments**
Wednesday, September 3	Deep Learning Concepts and Ideas	Homework 1 released
Monday, September 8	Supervised Learning	Homework 2 released
Wednesday, September 10	Loss Functions	Homework 1 due
Monday, September 15	Gradients and Gradient Descent	Homework 3 released
Wednesday, September 17	Shallow Neural Networks	Homework 2 due

Monday, September 22	SCC tutorial	Homework 4 released
Wednesday, September 24	Deep Neural Networks	Homework 3 due
Monday, September 29	Fitting Models	Homework 5 released
Wednesday, October 1	Backpropagation and Initialization of Neural Networks	Homework 4 due
Monday, October 6	Measuring Performance and Generalization	Project 1 released
Wednesday, October 8	Regularization	Homework 5 due
Tuesday, October 14	Convolutional Neural Networks	
Wednesday, October 15	Residual Networks	
Monday, October 20	Recurrent Neural Networks	Project 1 due, Project 2 released
Wednesday, October 22	Transformers part 1	
Monday, October 27	Transformers part 2	
Wednesday, October 29	Training, Tuning and Evaluating LLMs	
Monday, November 3	Vision Transformers	Project 2 due, Project 3 released
Wednesday, November 5	Contrastive Learning	
Monday, November 10	Adversarial Inputs and Generative Adversarial Networks	
Wednesday, November 12	Unsupervised Learning and Variational Autoencoders	
Monday, November 17	Normalizing Flows	Project 3 due, Project 4 released
Wednesday, November 19	Diffusion Models	
Monday, November 24	Using Pre-trained Models	
Monday, December 1	Scaling Tradeoffs	
Wednesday, December 3	Reasoning and World Models	
Monday, December 8	Object Detection and Segmentation	Project 4 due
Wednesday, December 10	Benchmarks	

Student Code of Conduct

All students are expected to abide by University conduct policies as detailed in the following links:

- Boston University Student Codes of Conduct
- College of Arts & Sciences Codes of Conduct
- Boston University Student Responsibilities

Academic Honesty

You may discuss homework assignments with classmates, but you are solely responsible for what you turn in. Collaboration in the form of discussion is allowed, but all forms of cheating (copying parts of a classmate's assignment, plagiarism from books or old posted solutions) are NOT allowed. We – both teaching staff and students – are expected to abide by the guidelines and rules of the Academic Code of Conduct.
Graduate students must also be aware of and abide by the GRS Academic Conduct code.
You can probably, if you try hard enough, find solutions for homework problems online. Given the nature of the Internet, this is inevitable. Let me make a couple of comments about that:

1. If you are looking online for an answer because you don't know how to start thinking about a problem, talk to the TA or instructor, who may be able to give you pointers to get you started. Piazza is great for this – you can usually get an answer in an hour if not a few minutes.
2. If you are looking online for an answer because you want to see if your solution is correct, ask yourself if there is some way to verify the solution yourself. Usually, there is. You will understand what you have done much better if you do that. So … it would be better to simply submit what you have at the deadline (without going online to cheat) and plan to allocate more time for homeworks in the future.

Generative AI Assistance (GAIA) Policy

In general, we follow the policy outlined in the CDS GAIA Policy.
Extracting and paraphrasing from the student responsibilities of that policy. Where there is conflicting information between the CDS policy and below, the policy below should take precedence.
Students shall:

1. Give credit to AI tools whenever used, even if only to generate ideas rather than usable text, illustrations or code.
2. When using AI tools on written assignments, unless prohibited, add an appendix showing
    1. the entire exchange, highlighting the most relevant sections;
    2. a description of precisely which AI tools were used (e.g. ChatGPT private subscription version or DALL-E free version),

    3. an explanation of how the AI tools were used (e.g. to generate ideas, turns of phrase, elements of text, long stretches of text, lines of argument, pieces of evidence, maps of conceptual territory, illustrations of key concepts, etc.);
    4. an account of why AI tools were used (e.g. to save time, to surmount writer's block, to stimulate thinking, to handle mounting stress, to clarify prose, to translate text, to experiment for fun, etc.).
    5. Optional but recommended: Employ AI detection tools and originality checks prior to submission, ensuring that their submitted work is not mistakenly flagged.
3. When using AI tools on coding assignments, unless prohibited
    1. Add the prompt text and tool used as comments before the generated code. Clarify whether the code was used as is, or modified somewhat, moderately or significantly.
4. Not use AI tools during in-class examinations, or assignments, unless explicitly permitted and instructed.
5. Use AI tools wisely and intelligently, aiming to deepen understanding of subject matter and to support learning.

As these generative assistive tools become widely deployed and pervasive, we believe they will become integral to most people's workflow. However, for foundational concepts, as are taught in this course, it is in your best interest and worth it to struggle some in creating your answers and solutions. It is just as important to learn what doesn't work, and which paths are dead ends, as it is to learn what does work. When you are posed with new and unique problems, that intuition you develop will be vital in choosing directions. More pragmatically, some of the most coveted jobs at the most selective companies require technical interviews where they expect you to know these foundational concepts without assistance.

And finally, to reiterate, it is vitally important, and a core part of academic integrity, to cite when you are using Generative AI Assistive technologies. Arguably, not citing and risking plagiarism is worse than taking short cuts and using and then citing GAIA.

Accommodations for Students with Disabilities

If you have a disability and have an accommodations letter from the Disability & Access Services 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 BU Disability & Access Services, I encourage you to find more information at https://www.bu.edu/disability/.
This syllabus provides a general plan for the course; deviations may be necessary depending on the progress of the class.

Acknowledgements

This syllabus and the general course design is based on the first offering of this course by Tom Gardos.