BOSTON UNIVERSITY

# Spring 2025 Deep Learning: Syllabus and Schedule

| Schedule | Syllabus | Lecture Slides | Piazza |
|---|---|---|---|

## Course Description:

This course is an introduction to deep learning, a branch of machine learning concerned with the development and application of modern neural networks. Deep learning algorithms extract layered high-level representations of data in a way that maximizes performance on a given task. For example, when asked to recognize faces, a deep neural network may learn to represent image pixels first with edges, followed by larger shapes, then parts of the face like eyes and ears, and, finally, individual face identities. Deep learning is behind many recent advances in AI, including Siri's and Alexa's speech recognition, Facebook's tag suggestions and self-driving cars. We will cover a range of topics from basic neural networks, convolutional and recurrent network structures, deep unsupervised and reinforcement learning, and applications to problem domains like speech recognition and computer vision.
**Prerequisites**: a strong mathematical background in calculus, linear algebra, and probability & statistics, as well as prior coursework in machine learning and programming experience in Python.

**Lecture**:

ENG EC523 A1
T/Th 1:30-3:15pm, PHO 205

**Instructor:**

Brian Kulis
office hours:  Tues 10-11am, outside PHO 441

**Teaching Assistant:**

Zijian Chen

**Graders/Additional Staff:**

Amruth Niranjan
Mete Gumusayak
Christian So
Brennan Mahoney

**Grader/TA Office Hours**:

Mondays 4:30-5:30pm, Location: PHO 442
Fridays 12:30-1:30pm, Location PHO 442

**How to Contact us:** Please use Piazza for all communication; if your question is only directed to the instructors, please make a post to "Individual Student(s) / Instructor(s)" and select "Instructors".

**Piazza: https://piazza.com/bu/spring2025/engec523**
We will be using piazza for online discussions, questions, and to post assignments.

**Gradescope:** we will be using Gradescope for submitting and grading assignments.

## Schedule*

| | Topic | Details | Homework |
|---|---|---|---|
| Tue Jan 21 | 1. Course overview | What is deep learning? DL successes; syllabus & course logistics. | |
| Thu Jan 23 | 2. Math/ML Review I | Probability, distributions, maximum likelihood, empirical risk minimization. | |
| Tue Jan 28 | 3. Math/ML Review II | Generalization, train/validation/test splits, stability, stochastic gradient descent. | HW1 out (machine learning prereq) |
| Thu Jan 30 | 4. Neural network basics I | Classification and regression tasks, perceptron, universal approximation. | |
| Tue Feb 4 | 5. Neural network basics II | MLP, activation functions, surrogate loss functions, softmax, and compression. | |
| Thu Feb 6 | 6. Neural network basics III | Automatic differentiation and backpropagation, matrix derivatives. | |
| Tue Feb 11 | 7. Training I | Mini-batching, regularization, adversarial examples, dropout, batch norm, layer norm. | HW1 Due<br>HW2 out |
| Thu Feb 13 | 8. Quiz 1 | Material Through Neural Network Basics III | SCC Info |
| Thu Feb 20 | 9. Training II | Momentum and acceleration, physical interpretation of accelerated gradient descent, stochastic gradients and variance. | |
| Tue Feb 25 | 10. Training III | Adaptive gradient methods: adagrad, adam, Lars/Lamb, and large batch sizes. | HW2 due<br>HW3 out |
| Thu Feb 27 | 11. CNNs | Convolutional neural networks, including AlexNet, VGG, and Inception; **Reading**: Goodfellow Ch9.1-9.3. | |
| Tue Mar 4 | 12. CNNs II and Advanced Architecture Design I | Modern Conv Nets, ResNet | Project Proposal Due<br>Template<br>How to make a group submission |
| Thu Mar 6 | 13. Advanced Architecture Design II | | |
| Tue Mar 11 | SPRING BREAK | | |
| Thu Mar 13 | SPRING BREAK | | |
| Tue Mar 18 | 14. Deep Unsupervised Learning I | Autoencoders | HW 3 due |
| Thu Mar 20 | 15. Deep Unsupervised Learning II | Variational Autoencoders | |
| Tue Mar 25 | 16. Deep Unsupervised Learning III | Diffusion Models | HW 4 out |

| | | | |
|---|---|---|---|
| Thu Mar 27 | 17. Deep Unsupervised Learning IV | Generative Adversarial Networks | |
| Tue Apr 1 | 18. RNNs | Recurrent neural networks; sequence modeling; backpropagation through time; vanishing/exploding gradient problem; gradient clipping, long-short term memory (LSTM). | |
| Thu Apr 3 | 19. Transformers I | Embeddings, word vectors, self-attention, transformers. | |
| Tue Apr 8 | 20. Transformers II | GPT, BERT, pretraining, masked language modeling task, few-shot learning. | HW4 due<br>HW5 out |
| Thu Apr 10 | 21. Graph Neural Networks | | Project Status Report Due Thu 11:59pm<br>Template |
| Tue Apr 15 | 22. Audio I | Keyword Spotting, Audio Synthesis | |
| Thu Apr 17 | 23. Audio II | Automatic Speech Recognition | |
| Tue Apr 22 | 24. Self-supervised Learning | self-supervised learning (slides from this tutorial) | HW5 due |
| Thu Apr 24 | 25. Quiz 2 | Cumulative | |
| Tue Apr 29 | Project Presentations I | Presentation Schedule<br><br>Send slides (or slide link) to my email | Slides due Apr 29 12:00pm NOON |
| Thu May 1 | Project Presentations II | Presentation Schedule<br><br>Send slides (or slide link) to my email | Slides due Thu May 1 12:00pm NOON<br>Slide Template |
| | | | |
| Tuei May 6 | Final Project reports/code due (No lecture) | Upload your reports on GradeScope, **use the template**<br>Sample Projects: reports, slides | Due Fri 11:59pm MIDNIGHT<br>Report Template |

*schedule is tentative and is subject to change.

# Syllabus

### Course Prerequisites

This is an upper-level undergraduate/graduate course. All students should have the following skills:

- Calculus, Linear Algebra
- Probability & Statistics
- Ability to code in Python
- Background in machine learning (e.g. EC 414, EC 503, CS 542)

### Textbook

There is no required textbook for the course. Some recommended textbooks:

- Ian Goodfellow, Yoshua Bengio, Aaron Courville. Deep Learning. MIT Press, 2016.

- Aston Zhang, Zack C. Lipton, Mu Li, and Alexander Smola. Dive into Deep Learning, 2020.
- Christopher M. Bishop. Deep Learning: Foundations and Concepts. Springer, 2023.

Other recommended supplemental textbooks on general machine learning:

- Duda, R.O., Hart, P.E., and Stork, D.G. **Pattern Classification**. Wiley-Interscience. 2nd Edition. 2001.
- Theodoridis, S. and Koutroumbas, K. **Pattern Recognition. Edition 4**. Academic Press, 2008.
- Russell, S. and Norvig, N. **Artificial Intelligence: A Modern Approach**. Prentice Hall Series in Artificial Intelligence. 2003.
- Bishop, C. M. **Neural Networks for Pattern Recognition**. Oxford University Press. 1995.
- Hastie, T., Tibshirani, R. and Friedman, J. **The Elements of Statistical Learning**. Springer. 2001.
- Koller, D. and Friedman, N. **Probabilistic Graphical Models**. MIT Press. 2009.

Recommended online courses

- http://cs231n.stanford.edu/ CS231n: Convolutional Neural Networks for Visual Recognition
- http://web.stanford.edu/class/cs224n/ CS224n: Natural Language Processing with Deep Learning
- http://rll.berkeley.edu/deeprlcourse/ CS 294: Deep Reinforcement Learning
- http://distill.pub/ Very nice explanations of some DL concepts

## Deliverables/Graded Work

There will be five homework assignments, each consisting of written and/or coding problems, and a final project. The worst homework grade will be dropped, but one must complete HW 1. The project will be done in teams of 3-4 students and will have several deliverables including a proposal, progress update(s), final report and a final in-class/virtual presentation. We will also have two written in-class exams (dates TBD). The course grade consists of the following:

- Homeworks (hw1 and best 3 of 2-5) 35%
- Project (including all components) 35%
- Written Quizzes 30% (10% Quiz I + 20% Quiz II)

## Software/Hardware

Programming assignments and projects will be developed in the Python programming language. We will also use the pytorch deep learning library for some homeworks and for the project. Students are expected to use the Shared Computing Cluster (SCC) and/or their own machines to complete work that does not require a GPU. For the projects, we will provide GPU resources.

## Late Policy

Late work will incur the following penalties

- Project deliverables: 20% off per day up to 2 days
- Homework 20% off per day, up to 3 days
- We will automatically drop the lowest scoring homework (except hw1)

## Academic Honesty Policy

The instructors take academic honesty very seriously. Cheating, plagiarism and other misconduct may be subject to grading penalties up to failing the course. Students enrolled in the course are responsible for familiarizing themselves with the detailed BU policy, available here. In particular, plagiarism is defined as follows and applies to all written materials and software, including material found online. Collaboration on homework is allowed, but should be acknowledged and you should always come up with your own solution rather than copying (which is defined as plagiarism):

> **Plagiarism**: Representing the work of another as one's own. Plagiarism includes but is not limited to the following: copying the answers of another student on an examination, copying or restating the work or ideas

of another person or persons in any oral or written work (printed or electronic) without citing the appropriate source, and collaborating with someone else in an academic endeavor without acknowledging his or her contribution. Plagiarism can consist of acts of commission-appropriating the words or ideas of another-or omission failing to acknowledge/document/credit the source or creator of words or ideas (see below for a detailed definition of plagiarism). It also includes colluding with someone else in an academic endeavor without acknowledging his or her contribution, using audio or video footage that comes from another source (including work done by another student) without permission and acknowledgement of that source.

## Religious Observance

Students are permitted to be absent from class, including classes involving examinations, labs, excursions, and other special events, for purposes of religious observance. In-class, take-home and lab assignments, and other work shall be made up in consultation with the student's instructors. More details on BU's religious observance policy are available here.