Deep Learning

Instructor: Yannet Interian
Email: yannet.interian@esade.edu

Course Description

This course offers a practical introduction to deep learning, combining theoretical foundations with hands-on PyTorch implementation. You'll explore essential neural network architectures like Feedforward Networks, CNNs, and Transformers, along with core concepts such as loss functions, optimization, and regularization. By the end, you'll be able to build and train deep learning models for real-world computer vision and natural language processing challenges.

Course Learning Objectives

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

1. **Design and train deep learning models using PyTorch** to solve tasks such as image classification, sequence modeling, and natural language understanding, applying architectural principles of CNNs, RNNs, and Transformers.

2. **Apply optimization and regularization techniques** (e.g., Adam, learning rate scheduling, Dropout, BatchNorm) to improve model generalization and mitigate overfitting or gradient instability.

3. **Manage end-to-end deep learning workflows**, including data preprocessing, model training, and evaluation on CPU/GPU, using PyTorch's data and compute utilities.

4. **Independently address real-world business or societal challenges using deep learning**, justifying architecture choices, evaluating model performance, and communicating findings clearly.

Course Materials

Required Textbook

The primary textbook for this course is:

- Prince, S. J. D. (2023). Understanding Deep Learning. MIT Press. Available at `https://udlbook.github.io/udlbook/`

Optional Recommended Books

The following books are recommended as supplementary resources for different learning styles and deeper dives into specific topics and PyTorch implementation:

- Bishop, C. M. Deep Learning: Foundations and Concepts. Cambridge University Press.
- Stevens, E., Antiga, L., & Viehmann, T. (2020). Deep Learning with PyTorch. Manning Publications.

Online Resources

We will also frequently utilize official PyTorch documentation and tutorials, the Hugging Face documentation, and various online articles and research papers as supplementary learning materials. Links to these resources will be provided on the course website.

Grading

In this course, weekly exercises are shorter, focused tasks designed to reinforce specific concepts and provide immediate practice. In contrast, programming assignments are larger, more comprehensive projects that require you to integrate multiple concepts and apply problem-solving strategies to build complete solutions. Your final grade will result from these items:

- *5%* Quizzes
- *25%* Assignments
- *30%* Midterm exam
- *40%* Final exam

Course Schedule

Week	**Topics and Key Concepts**	**Assignments Due**
Week 1 03-05/09	**Lecture 1**: Introduction to Deep Learning **Lab:** torch tensors, image manipulations	
Week 2 10-12/09	**Lecture 2**: - Shallow Neural Networks **Lab:** fitting linear models and shallow networks	
Week 3 17-19/09	**Lecture 3**: Deep Neural Networks **Lab:** experiments on learning rates, size of hidden layer, augmentations	
Week 4 26/9	**Lecture**: **Holiday** **Lab:** text classification	HW1 due 25/09
Week 5 01-03/10	**Lecture 4**: Embedding Models **Lab:** dataset and dataloaders, text classification experiments	
Week 6 08/10	**Lecture 5**: Losses and Metrics **Lab: No Class**	HW2 due 09/10
Week 7	Midterm Exam	
Week 8 22-24/10	**Lecture 6**: Fitting Models, Gradients and Initialization **Lab:**	
Week 9 29-31/10	**Lecture 7**: Training Basics: Convolutional Neural Networks **Lab:**	
Week 10 05/11	**Lecture 8**: Recurrent Neural Networks (RNNs) and Regularization **Lab: No Class**	HW3 due 06/11
Week 11 12-12/11	**Lecture 9**: Transformers Part I **Lab:**	
Week 12 19-21/11	**Lecture 10**: Transformers Part II **Lab:** .	
Week 13 28/11	**Lecture**: **No Class** **Lab:**	HW4 due 28/11