PhD Course 2024-2025

Deep learning

Clément Gorin
Panthéon-Sorbonne University
clement.gorin@univ-paris1.fr

Course's objective

Many important questions in economics remain unanswered, partly because the data necessary to address them is encoded into high-dimensional data structures such as images or text. Applied scientists have become increasingly interested in using deep learning models to transform these data into simpler representations, which can be used as inputs for subsequent analysis. After introducing predictive modelling, this course provides a comprehensive understanding of neural network models, as well as some of the most capable structures to process image and language data. From a strong base in theory and mathematical formalisation, focus is kept on intuition and efficient implementation using Python, both to illustrate abstract statistical concepts using simulated data and to implement the models studied in class.

Outline

The course is organised in four three-hours sessions. Each session articulates a theoretical lecture and practical applications using Python. The sessions are organised as follows:

Lecture 1: Predictive modelling
Function approximation – Inference and prediction – Bias and variance – Regularisation – Resampling methods

Lecture 2: Neural networks
Model structure – Learned representations – Estimation – Backpropagation – Better optimisation

Lecture 3: Image modelling
Images as data – Image processing – Convolutional networks – Applications (e.g. image classification, segmentation)

Lecture 4: Language modelling
Text as data – Embeddings networks – Recurrent networks – Applications (e.g. sentiment analysis, semantic distances)

Course materials

Lectures, practical exercises, papers and other resources will be made available on the Dropbox of the course, which is updated before every lecture. Students should bring their personal computer with administrator rights and a working internet connection. For the practice sessions, we will use a Miniconda or Anaconda distribution of Python, which provides basic packages and simplify management. A Conda environment containing the necessary packages will be provided before the class.

Assessment

Students are assessed on the basis of a final project. The project entails selecting a research question involving a prediction task, preferably related to the PhD thesis, and addressing it using the tools and models covered in class. The final submission must be formatted as a poster summarising the analysis, with an optional appendix, and the replication code.

Handbooks

- Hastie, Trevor, Robert Tibshirani, and Jerome Friedman. *The elements of statistical learning*. Springer, 2009.
- Goodfellow, Ian, Yoshua Bengio, and Aaron Courville. *Deep learning*. MIT Press, 2016.
- Nielsen, Michael. *Neural Networks and Deep Learning*. Determination Press, 2019.

Papers

- Breiman, Leo. *Statistical modeling: The two cultures*. Statistical Science, 16(3): 199-231, 2001
- LeCun, Yann, Yoshua Bengio, and Geoffrey Hinton. *Deep learning*. Nature, 521: 436-444, 2015.
- Sendhil, Mullainathan and Spiess Jann. *Machine learning: An applied econometric approach*. Journal of Economic Perspective, 31(2): 87-106, 2017.
- Gentzkow, Matthew, Bryan Kelly, and Matt Taddy. *Text as data*. Journal of Economic Literature, 57(3): 535–574, 2019.
- Lones, Michael A. *How to avoid machine learning pitfalls*: A guide for academic researchers. *CoRR, 2021.*

Specific references will be given in class.