Deep Learning
(MO434A/MC934B)

Prof. Alexandre Xavier Falcão

January 1, 2023

1 What is this all about?

This course provides the knowledge to construct and use deep neural networks for image and text analysis. The course starts from the basic concepts to understand, train and test neural networks for classification and regression. It introduces image analysis and then evolves to (Fully) Convolutional Neural Networks for image classification, object detection, and (semantic/instance) segmentation. In the sequence, it provides an introduction to text analysis and then covers Recurrent Neural Networks, Attention, Transformers and applications in text analysis.

Prior knowledge in optimization, linear algebra, statistics, machine learning, image/text processing and analysis is important, but the basic concepts are provided whenever they are required.

It is important the student can code in Python and desirable prior knowledge in PyTorch, and other packages usually used in python scripts for image and text processing, graphics display, and machine learning.

2 Location and schedule

The lectures will take place at the classroom CC 52 of the building IC3 (i.e., room 352) every Tuesdays from 4:00 PM to 6:00 PM, being the last 20 minutes reserved to assist the students with the practical tasks (exercises in the notebooks).

3 *Syllabus*

- Introduction to deep learning.
- Fundamentals of Deep Neural Networks (DNNs).
- The art of training DNNs.
- Fundamentals for image analysis using DNNs.
- Convolutional Neural Networks (CNNs).
- Network visualization, image classification and object detection.
- Fully Convolutional Neural networks (FCNs) and image segmentation.
- Fundamentals for text analysis.
- Recurrent Neural Networks (RNNs), attention and transformers.
- Applications in text analysis and image classification with transformers.

The lectures are complemented with hands-on activities using jupyter notebooks.

4 Evaluation criteria and instructions for the final report

Students with less than 75% of attendance will be graded E. Students with 75% or more than 75% of attendance will be graded based on

- their participation during the lectures,
- practical exercises in several notebooks,
- a practical project that will be presented during the course.

A report describing the results of the exercises and project shall be delivered at the end of the course. The clarity and correctness of the codes will be considered as well as the quality and completeness of the report.

Scores in [0,10] will be assigned to each student and graded as follows.

- grade A for score in [8.5,10];
- grade B for score in [7.0,8.5);
- grade C for score in [5.0,7.0);
- grade D for score in [0.0,5.0);

The report may have from 20-40 pages with letter-size 11pt, including figures, tables, graphics, and references. They should present the following organization.

- Cover page: provide the name of the discipline, name of the student, academic identification number (RA), and delivery date, followed by a summary of the implemented techniques and their main results.
- Subsequent pages: present the literature that has been studied to implement the project and exercises, difficulties, given solutions, implemented algorithms, and their results with illustration and discussion.
- Final page: present a conclusion about what has been accomplished with this course and provide suggestions to improve it.

The students shall send a tar.bz2 file with the report and codes related to each exercise. Datasets can be shared by informing their link in google drive. **Please, do not include datasets in the tar.bz2 file**.

5 Bibliography

The books used to prepare this course are listed below.

1. V.K. Ayyadevara and Y. Reddy. Modern Computer Vision with Pytorch. Packt, 2020.
2. A. Zhang, Z.C. Lipton, M. LI, and A.J. Smola. Dive into Deep Learning: Interactive deep learning book with code, math, and discussions, `https://d2l.ai/`, 2020.
3. F. Chollet. Deep Learning with Python. Manning, 2018.
4. A. Géron. Hands-on Machine Learning with Scikit-Learn, Keras & Tensorflow. O'Reilly, 2nd Ed., 2019.
5. S. Vajjala, B. Majumder, A. Gupta, and H. Surana. Practical Natural Language Processing. O'Reilly, 2020.
6. D. Sarkar. Text Analytics with Python. Apress, 2019.
7. S. Ravichandiran. Getting started with Google BERT. Packt, 2021.
8. K. Koutroumbas and S. Theodoridis. Pattern Recognition. 4th Ed., Academic Press, 2009.