Deep Learning
(MO434/MC934)

Prof. Alexandre Xavier Falcão

July 6, 2021

1 Overview

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, but not necessary, prior knowledge in keras, 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 happen through Google Meet (`https://meet.google.com/`) every Monday from 10 AM to 12 AM. The communication with the students will happen through Google Classroom (`https://classroom.google.com`).

3 *Syllabus*

- (Deep) neural networks for classification and regression.
- The perceptron and backpropagation algorithms.
- Optimizers, regularizers, activation and loss functions.
- Introduction to image analysis and its applications.
- (Fully) Convolutional neural networks (CNNs, FCNs), information visualization, and applications in image analysis.
- Introduction to text analysis and its applications.
- Recurrent neural networks, attention, transformers and applications in text analysis.

The lectures are complemented with hands-on activities in keras&tensorflow and pytorch using jupyter notebooks.

4 Evaluation criteria

One project is assigned to every student and the student's performance is evaluated based on

- the student's participation during the lectures,
- the quality of the presented report about the project,
- the quality of the implemented code (jupyter notebook) and its documentation.

Each project receives a score in [0,10] and the students are 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 should have a maximum of 20 pages with letter-size 11pt, including figures, tables, graphics, and references. It 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, its 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.

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. F. Chollet. Deep Learning with Python. Manning, 2018.
3. A. Géron. Hands-on Machine Learning with Scikit-Learn, Keras & Tensorflow. O'Reilly, 2nd Ed., 2019.
4. S. Vajjala, B. Majumder, A. Gupta, and H. Surana. Practical Natural Language Processing. O'Reilly, 2020.
5. D. Sarkar. Text Analytics with Python. Apress, 2019.
6. S. Ravichandiran. Getting started with Google BERT. Packt, 2021.
7. K. Koutroumbas and S. Theodoridis. Pattern Recognition. 4th Ed., Academic Press, 2009.