Neural Networks for Data Science Applications

Master's Degree in Data Science

Lecture 0: About the course

Lecturer: S. Scardapane

SAPIENZA
UNIVERSITÀ DI ROMA

- **Master's Degree in Data Science**, code 10589627, 2nd year, optional group D, SSD ING-IND/31.
- **Timetable**: Wednesday, 5-7 PM, Room A5-A6 (Via Ariosto), Friday, 8-11 AM, Room A5-A6 (Via Ariosto).
- **Office hours**: by appointment, remotely or in-person (Via Eudossiana 18, DIET Department, 1st floor, room 102).

Official course website:
`https://www.sscardapane.it/teaching/nnds-2024/`.

Register to the **Google Classroom** from the website for all updates (mandatory).

Exam dates (tentative)

1. **Sessions 1-2**: January 20 and February 20.
2. **Sessions 3-4**: June 20 and July 18.
3. **Session 5**: September 19.
4. **Session E1**: March 20 (**reserved**, see regulations).
5. **Session E2**: October 20 (**reserved**, see regulations).

1. One **mid-term** homework (5 points) (*can be recovered during the final project*).
2. One **end-of-term** homework (10-15 points).
3. One **oral examination** on the program (15 points).

The EoT homework can be sent before *any* exam date. The marks for the two homeworks can be kept during the academic year, irrespective of the oral. *Lode* will be given only to exceptional (top 5%) homeworks and orals.

- Fundamental tools underlying neural networks: optimization, gradient descent, automatic differentiation.
- Basic blocks to build modern neural networks (convolution, attention, normalization, ...).
- Proficiency in a real-world deep learning library (**TensorFlow**, **JAX**).
- Capability of navigating the current literature and ecosystem in autonomy, and understanding some critical limitations (e.g., bias, brittleness).

1. **Preliminaries** (tensors, linear algebra, optimization).
2. **Supervised learning** as numerical optimization.
3. Linear models and **fully-connected** models.
4. **Convolutional models** (for sequential, spatial, and temporal data).
5. Blocks to train **deeper models** (dropout, batch normalization, ...).
6. **Attention** models for sets.
7. **Graph** models (e.g., graph convolutional networks).
8. Optional topics depending on time and material.

1. Several practical lectures with **TensorFlow** and **JAX** (hands-on coding from scratch).
2. When possible, a showcase of other libraries (e.g., HuggingFace Datasets).
3. Optional topics depending on time and material.

Slides are self-contained, but the material is expanded in a textbook:

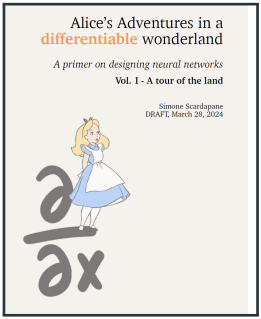

The book is available for free as a PDF or via Amazon for a printed copy:

`https://sscardapane.it/alice-book/`

The book was recently published – would be happy for feedback or ideas for completed exercises or additional sections!

Additional useful textbooks:

- **Dive into Deep Learning**, online, much more practical.
- **Understanding Deep Learning**, high-quality illustrations and descriptions.
- **Patterns, Predictions, and Actions**, slightly broader on the machine learning side.
- **Deep Learning - Foundations and Concepts**, for a beginner Bayesian treatment.

From mobile, you can also check out **The Little Book of Deep Learning**.