Statistical aspects of Deep Neural Networks

Phd Programme in Economics, Statistics and Data Science, University of Milano-Bicocca

The Ph.D. in Economics and Statistics and the Department of Economics, Management and Statistics of the University of Milano-Bicocca are pleased to announce the course *Statistical aspects of Deep Neural Networks* held by Professor Omiros Papaspiliopoulos, Bocconi University, during the period October 5-27, 2021.

Instructor

Omiros Papaspiliopoulos is Full Professor at the Bocconi University, Milano. Previous to that, he was ICREA Research Professor based at Universitat Pompeu Fabra. He have held other positions at Warwick, Oxford and Lancaster, and have also worked in Berlin, Osaka, Paris, Madrid and Lima. In 2010 he was awarded the Royal Statistical Society's Guy Medal. In 2013 he founded the first Master in Data Science in Europe at Barcelona Graduate School of Economics, and in 2016 the Data Science Center, which operated at the intersection of scientific research and consulting, both of which he directed until 2021. Currently he is co-editor of Biometrika and associate editor SIAM/ASA Journal of Uncertainty Quantification. Google scholar profile: `https://scholar.google.com/citations?user=Rspjrv8AAAAJ&hl=en`.

Enrollment

The course *Statistical aspects of Deep Neural Networks* is part of the *Statistical Learning* course of the PhD in Economics and Statistics, but it is open to all PhD students, researchers, academics, professionals and practitioners with strong quantitative background and motivation (see Prerequisites). Lectures will be delivered **in-class**.

Application period: September 8, 2021 - September 17, 2021. In order to apply you must fill in the following form:

`https://forms.gle/BRetfcwHikHunxbG7`

Admissions are conditional on place availability. **No fees are requested**. The ECOSTAT Administration Office will contact non-admitted candidates only.

Course objectives

The course covers the foundations of deep neural networks with a special emphasis and priority on more statistical aspects of this research agenda. Nevertheless, also practical aspects are addressed and illustrations in Pytorch are given. Every lecture discusses a couple of (typically very recent) research articles. Background on modern statistical machine learning is useful to follow the course. The course is organised along the following themes:

- A. From regression to random features to double descent to feedforward neural networks

  We discuss high-dimensional regression, the random features models (which are half-way between linear regression and neural networks), the benign overfitting and double descent phenomena and introduce the fundamental neural network structure. We show theoretical results on the performance of DNNs as nonparametric regression models.

- B. From PCA to autoencoders

  We cast basic unsupervised learning algorithms as neural networks and introduce the concept of disentangled representations

- C. Optimisation for neural networks

  We discuss architectural choices and their effect on learning, we present results on the optimisation landscape related to neural network learning, we present back-propagation for computing gradients and discuss gradient-based optimisation

- D. Overview of software and implementations in Pytorch

- E. From deep neural networks to stochastic processes

  We overview the link between neural networks and kernel methods, in particular Gaussian processes in the wide limit, and residual networks and neural ODEs and SDEs in the deep limit

- F. Regularisation, stability and adversarial training

  We discuss weight decay, early stopping, dropout and their interconnection, we highlight the instability problem both for forward and inverse regression and discuss adversarial training

- G. Connection between simulation-based inference and neural networks

  We discuss a number of interesting connections between neural networks and simulation-based inference, among others also GANs

- H. Special architectures

  We discuss convolutional and recurrent neural networks and more recent ideas such as transformers

Main references:

- 1. Deep learning book, https://www.deeplearningbook.org/, Chapter 1, 4, 6, 7, 8, 13, 14, 15, 16, 20
- 2. Bishop's Pattern Recognition and ML, Chapter 5
- 3. Ripley's Pattern Recognition and Neural Networks, Chapter 5, 6 and 10
- 4. Neal's Bayesian Learning for Neural Networks, Chapter 1 and 2
- 5. Hastie et al.'s Elements of Statistical Learning, Chapter 5 (basis expansions), Sections 14.4, 14.5, 14.7
- 6. Various research articles, among others also the following:
- Baldi and Hornik (1998, Neural Networks) Learning from Examples Without Local Minima
- Schmidt-Hieber (2020, Annals of Stats) Nonparametric regression using DNN with Relu activation function
- Eckle and Schmidt-Hieber (2019, Neural Networks) A comparison of deep networks with Relu activation function and linear spline-type methods
- Mei and Montanari (2019) The generalisation error of random features regression
- Rahimi and Recht (2007, NIPS) Random features for large-scale kernel machines
- Madry et al. (2019) Towards Deep Learning models resistant to adversarial attacks
- Wager, Wang, Loang (2013, NIPS) Dropout training as adaptive regularization

Schedule

- Tuesday 5 October 11.00-13.00 and 14.00-15.30
- Wednesday 6 October 11.00-13.00 and 14.00-15.30
- Tuesday 12 October 11.00-13.00 and 14.00-15.30
- Wednesday 13 October 11.00-13.00 and 14.00-15.30
- Tuesday 19 October 11.00-13.00 and 14.00-15.30
- Wednesday 20 October 11.00-13.00 and 14.00-15.30
- Tuesday 26 October 11.00-13.00 and 14.00-15.30
- Wednesday 27 October 11.00-13.00 and 14.00-15.30

Lectures will be delivered in-class at the University of Milano-Bicocca (aula “de Lillo”, 2nd floor, building U7), but there may be changes due to coronavirus.

Prerequisites
The overall flavour of the course overall will be theoretical and methodological. The material should be understandable for anyone with a thorough knowledge of undergraduate Statistics, probability, real analysis and linear algebra.

Contacts
For more information:
Prof. Matteo Manera, Coordinator of ECOSTAT, e-mail: matteo.manera@unimib.it
Prof. Aldo Solari, ECOSTAT faculty member, e-mail: aldo.solari@unimib.it
Dr. Tommaso Rigon, DEMS faculty member, e-mail: tommaso.rigon@unimib.it

For administrative issues:
Dr. Silvia Locatelli, ECOSTAT Administration Office, e-mail: phdecostat@unimib.it