



UNIVERSITÀ DI PISA
DIPARTIMENTO DI INGEGNERIA DELL'INFORMAZIONE
Dottorato di Ricerca in Ingegneria dell'Informazione

Doctoral Course

“Deep Learning for signal processing, vision and control”

Prof. Davide Bacciu

Dipartimento di Informatica, Università di Pisa

bacciu@di.unipi.it

Short Abstract: This course aims to provide an introduction to the design and use of deep learning models and reinforcement learning approaches for sensor data processing, machine vision and robotics. The first part of the course introduces the basic concepts and fundamentals of machine learning and neural networks. The second part presents advanced deep models and their use in monitoring, understanding, control and planning tasks, with focus on robotics and distributed sensing application scenarios. Presentation of the theoretical models and associated algorithms will be complemented by references to popular software frameworks and code. Given the course focus, much of the concepts and models presented will deal with sequential data (e.g. sensor or control timeseries) and visual data (images and video), with insights on relevant problems, including lifelong learning, reinforcement learning, federated learning and learning under resource constraints.

Course Contents in brief:

- Fundamentals of machine learning: generalization, model-selection, hyperparameters, regularization techniques, error function, maximum likelihood learning, basic concepts of probabilistic learning
- (Deep) Neural networks introduction: artificial neuron, backpropagation, optimization techniques, multi-layer perceptron, deep autoencoders, pretraining
- Convolutional neural networks: fundamental building blocks, advanced techniques, notable convolutional architectures, applications to image classification and semantic segmentation
- Recurrent neural networks: sequential data processing, early recurrent models, gated recurrent architectures, advanced memory models
- Reservoir computing: efficient recurrent neural models, memory-constrained recurrent model, sensor data processing, echo state networks, deep reservoir computing
- Generative deep learning: variational approximation, sampling in machine learning, variational autoencoders, generative adversarial networks

- Reinforcement learning: fundamentals, Markov decision processes, model-free and model-based algorithms, deep reinforcement learning, imitation learning
- Advanced topics and applications: lifelong and continual learning, federated learning in cloud/distributed environments, relational learning, deep learning for robotics, embedded learning systems.

Total # of hours of lecture: 20

References:

Two standard reference textbooks will be used (listed below). These will be complemented with selected article readings (2 to 4) for each lecture.

[1] Ian Goodfellow, Yoshua Bengio and Aaron Courville, Deep Learning, MIT Press

[2] Richard S. Sutton and Andrew G. Barto, Reinforcement Learning: An Introduction, Second Edition, MIT Press, 2018

CV of the Teacher

Davide Bacciu is senior Assistant Professor (Associate Professor, starting 01/07/2020) at the Department of Computer Science, University of Pisa, working on the general topics of machine learning and artificial intelligence. His research spans neural networks, deep learning, generative models, Bayesian learning, kernel methods, learning in structured and relational domains (sequences, trees and graphs), ambient intelligence, biomedical and life-science data processing, robotics, vision and computational creativity.

He has been supervising 14 Ph.D. students in Computer Science and Data Science, 8 research assistants on machine learning related projects and over 65 M.Sc. and B.Sc. theses in Computer Science, Mathematics and Robotics Engineering. He teaches courses on deep learning, neural networks and reinforcement learning at M.Sc. and Ph.D. level.

He is the coordinator of the H2020 project TEACHING (2020-22) on the development of humanistic intelligence for autonomous distributed applications, and he is coordinating an industrial research project on deep and Bayesian learning for perceived stress prediction from mobile data. He has coordinated the Italian national project LISTIT (MIUR-SIR program, 2016-2019) on the design of machine learning models for structured data. He has co-authored 98 research papers published in international journals, conference proceedings and book chapters. Currently, he is Associate Editor of the IEEE Transactions on Neural Networks and Learning Systems, Senior PC and PC member of top ML&AI conferences (NeurIPS, AAI, IJCAI, ICML, ECAI, IJCNN). He has delivered invited presentations in highly reputed research centers, including University of Heidelberg, Royal Statistical Society, King's College London and Italian Institute of Technology. He is the recipient of the 2009 E. R. Caianiello Award for the best Italian Ph.D. thesis in neural networks.

At a local level, he is the co-coordinator of the recently established Pervasive Artificial Intelligence Laboratory, a joint lab of Dipartimento di Informatica, Università di Pisa, and ISTI-CNR. At an Italian national level, he is Secretary (since 2017) and Steering Board member (since 2015) of the Italian Association for Artificial Intelligence. At an international level, he is a Senior Member of the IEEE, a member of the IEEE CIS Neural Networks Technical Committee and the Chair of the IEEE Task Force on Learning for Structured Data. Recently, he has been appointed Coordinator of the working group on Bioinformatics of the European CLAIRE-COVID19 task force.

Room and Schedule

Room: *From remote by using Microsoft Teams*

Schedule:

- 02/02/2021: 9:00-13:00
- 04/02/2021: 9:00-13:00
- 09/02/2021: 9:00-13:00
- 11/02/2021: 9:00-13:00
- 15/02/2021: 9:00-13:00