University of Pavia
Department of Electrical, Computer and Biomedical Engineering

Artificial Intelligence Reading Group (AIRG)

Deep Learning:
a theoretical introduction

Marco Piastra

University of Pavia
Department of Electrical, Computer and Biomedical Engineering

Episode 1: 26 May 2017, 14:15pm
Aula Seminari ExElettronica – Piano D

Episode 2: 9 June 2017, 14:15pm
Aula Seminari ExElettronica – Piano D

Episode 3: 16 June 2017, 14:15pm
Aula Seminari ExElettronica – Piano D

Abstract – Deep Learning techniques are at the heart of the recent, purported revolution in Artificial Intelligence that has opened many brand new perspectives for engineering applications. On a closer look, however, such techniques seem more evolutionary, in terms of substantial improvement over existing methods, than revolutionary. This cycle of reading groups is intended as an introductory mini-course to the subject, with the purpose of analyzing Deep Learning in a broader theoretical context and exploring the reasons of its unquestionable success in several specific application cases. The forthcoming book "Deep Learning", by I. Goodfellow, Y. Bengio and A. Courville will be assumed as reference for these episodes. At present, a draft of the book is freely accessible online at http://www.deeplearningbook.org/.

Organizer	**Ph.D. Coordinator**
Virginio Cantoni	Paolo Di Barba

The reading groups will be held in English.
The slides will be made available at:
http://vision.unipv.it/AI/AIRG.html

For more information: marco.piastra@unipv.it

University of Pavia
Department of Electrical, Computer and Biomedical Engineering

Artificial Intelligence Reading Group (AIRG)

Deep Learning:
a theoretical introduction

Program

Episode 1: The Context (date and time to be defined)
- Feed-forward networks (multi-layer perceptron)
- Universal approximation property: hidden layers
- Depth vs. breadth in feed-forward networks
- Gradient-based learning: backpropagation
- From backpropagation to automatic differentiation

Episode 2: Enabling Techniques (date and time to be defined)
- Supervised vs. unsupervised learning?
- Stochastic Gradient Descent (SGD)
- Graphical models: directed graphs
- Undirected graphs and energy-based models
- Restricted Boltzmann Machine (RBM)
- Gibbs sampling, ancestral sampling

Episode 3: Inside Deep Learning (date and time to be defined)
- Training Restricted Boltzmann Machines
- Deep Belief Networks
- Layer-wise pre-training: contrastive divergence
- Fine tuning
- Other variants: Multimodal DBNs, Deep Boltzmann Machines
- Appendix: Hints about practical applications
- Appendix: TensorFlow, Theano and similar frameworks