

Course 5	Neural Networks and Deep Learning
Program	<ol style="list-style-type: none"> 1. Introduction <ol style="list-style-type: none"> 1.1. Biological neurons 1.2. Artificial neurons 1.3. Rosenblatt perceptron 1.4. Multilayer perceptron 1.5. Other types of neural networks: <ol style="list-style-type: none"> 1.5.1 Radial basis function networks 1.5.2 Recurrent networks 2. Training algorithms <ol style="list-style-type: none"> 2.1. Machine learning 2.2 ML, MAP and Bayesian learning 2.3. The backpropagation algorithm 3. Learning and Optimization <ol style="list-style-type: none"> 3.1. Training as Error Minimization 3.2. Gradient Descent 3.3. Higher Order Descent <ol style="list-style-type: none"> 3.3.1. Newton's and Gauss-Newton Methods 3.3.2. Levenberg-Marquardt Method 3.3.3. Conjugate Gradient and Quasi-Newton Methods 3.3.4. Momentum and Nesterov acceleration 3.3.5 Adadelta and Adam 3.4. Optimization and generalization 3.5. Online and Minibatch Learning 4. MLPs in practice <ol style="list-style-type: none"> 4.1. Activation functions 4.2. Data preprocessing 4.3. Network initialization 4.4. Tuning network structure 4.5. Learning rate, momentum, weight decay 4.6. Fast training: Extreme Learning Machines 4.7 Practical data modelling with neural networks 5. Deep Networks <ol style="list-style-type: none"> 5.1. Introduction to deep networks <ol style="list-style-type: none"> 5.1.1. The problem of learning several hidden layers. 5.1.2. Deep networks initialization 5.1.3. Neurons in deep networks 5.1.4. Dropout 5.1.5. Current approaches for building deep networks 5.2. Neural networks learning <ol style="list-style-type: none"> 5.2.1. Optimization methods for deep networks 5.2.2. Exploiting parallel and distributed computation 5.3. Convolutional neural networks <ol style="list-style-type: none"> 5.3.1. Convolution layers 5.3.2. Pooling layers 5.3.3. Image processing with Deep Networks 5.4. Recurrent neural networks <ol style="list-style-type: none"> 5.4.1. Learning in recurrent networks 5.4.2. Long-Short Term Memory layers 5.4.3. Embedding layers 5.5 Practical Sessions: <ol style="list-style-type: none"> Python with keras+tensorflow and Jupyter notebooks.

Bibliography

- Y. LeCun, L. Bottou, G. Orr and K. Muller: Efficient BackProp, in Orr, G. and Muller K.
- (Eds), *Neural Networks: Tricks of the trade*, Springer, 1998.
- R. Duda et al. *Pattern classification*.
- C. Bishop, *Neural networks for pattern recognition*.
- T. Hastie et al. , *The elements of statistical learning*
- Y. Bengio, I. Goodfellow, *Deep Learning*