Course on
Neural Networks and Deep Learning


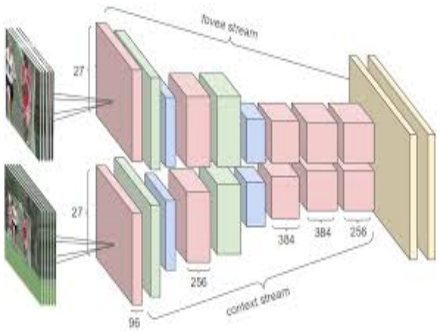

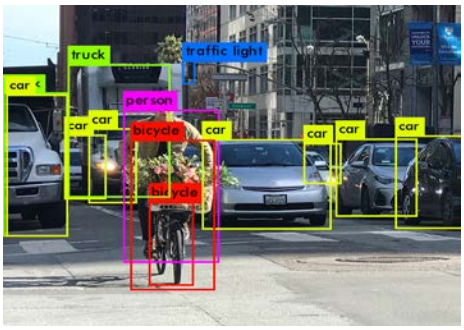

January – February 2019 – Gray Room
Retis Lab – TeCIP Institute
Via Moruzzi, 1 - Pisa

Giorgio Buttazzo, Alessandro Biondi, Daniel Casini

Objectives: The aim of the course is to provide the key concepts and methodologies on the main models and paradigms of existing neural networks, explaining how to use them for pattern recognition, image classification, signal prediction, system identification, and adaptive control.

Program (the first lecture is on **January 16, 2019** - **9:00** at the TeCIP Institute - Gray Room)

1. **Introduction to neural computing**. Motivations. Main network models and learning paradigms. Fully connected networks and associative memories. Hopfield networks. Application to optimization problems. Implementation issues. Examples and demos.
2. **Competitive learning**. Self-organizing maps. Kohonen networks: network model, learning algorithm and main network properties. Implementation issues. Examples and applications.
3. **Reinforcement Learning**. The state-box learning paradigm. Temporal credit assignment. The ASE/ACE neural model. Q-learning and SARSA algorithms. Implementation issues. Examples and applications.
4. **Supervised learning**. The Perceptron: model, properties and limitations. Multi-layer networks. The Back Propagation algorithm. Convergence and generalization. Applications of multi-layer networks to signal prediction, control and system identification. Recurrent networks. Implementation issues. Examples and applications.
5. **Towards deep networks**. Advantages of increasing the number of neural layers. Problems in training deep networks: overfitting and vanishing gradient. Solutions for deep learning: better loss functions, better activation functions, regularization, and dropout methods.
6. **Deep network models**: Boltzmann Machines, Restricted Boltzmann Machines, Autoencoders, Convolutional Networks. Implementation issues. Examples and applications.
7. **Specific deep neural networks**: LeNet-5, Alex-Net, VGG-Net, GoogLeNet, ResNet, SqueezeNet, SENet.
8. **Neural Networks for object detection**. Sliding windows, OverFeet, R-CNN, Yolo.

Practical Lectures

9. **Frameworks for training and inferring Deep Neural Networks**. How to model neural networks in Tensorflow and Caffe. Example of LeNet implementation on Caffe.
10. **Real time inference on Deep Neural Networks on GPGPUs**. This lecture will show how to use the TensorRT API by Nvidia to execute a DNN modelled in Caffe.