INSTITUTE OF COMMUNICATION, INFORMATION AND PERCEPTION TECHNOLOGIES
Scuola Superiore Sant'Anna

Retis
Real-Time Systems Laboratory

# Neural Networks and Deep Learning:
# Part 1: Theoretical Foundations

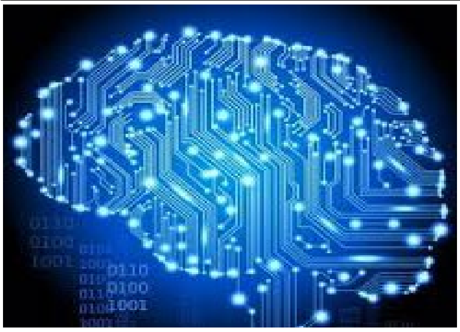


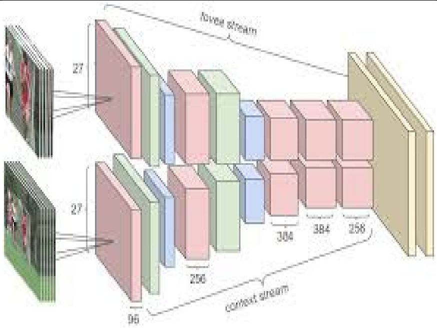


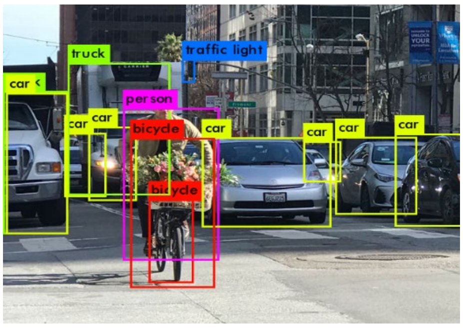


**The first lecture is scheduled on January 7, 2025 - 9:00**
Visit the course web page for registration and connecting to the channel.
http://retis.sssup.it/~giorgio/courses/neural/nn.html

**Giorgio Buttazzo**

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

## Course program

1. **Introduction to neural computing**. Motivations. From classical AI to machine learning: a paradigm shift. Typical machine learning problems. The neural approach. Evolution of neural networks. Main network structures and learning paradigms.
2. **Fully connected networks**. Hopfield networks. Associative memories: storage mechanism and storage capacity. Modern Hopfield networks. Application to optimization problems.
3. **Unsupervised learning**. The problem of high dimensionality. Dimensionality reduction and feature extraction: Principal component analysis and Self-organizing maps.
4. **Clustering algorithms**. K-means. Hierarchical clustering. DBSCAN. Clustering evaluation metrics. Examples and applications. Soft clustering: Fuzzy C-Means and Gaussian Mixture Model.
5. **Autoencoders**. Compressing, denoising, sparse, and contracting autoencoders.
6. **Reinforcement Learning**. Basic concepts and intuitions. Temporal credit assignment. The state-box learning paradigm. The ASE/ACE neural model. Q-learning and SARSA algorithms. TD-lambda extensions. Actor-Critic methods. Examples of control applications.
7. **Model-based Reinforcement Learning**. Definitions and taxonomy. Learning a model. The Dyna framework. Model-based RL in games: Minimax and Monte Carlo Tree Search. Off-line RL.
8. **Supervised learning - part 1**. The Perceptron: properties and limitations. Multi-layer networks. The Backpropagation algorithm. Stochastic gradient descent. Sample applications and demos.
9. **Supervised learning - part 2**. Important remarks. Applications to signal prediction, control, and system identification. Performance evaluation metrics. Managing overfitting and class imbalance.
10. **RBF networks**. Radial basis functions (RBF). Objectives and intuitions. Network architecture. Training algorithm. Multi-scale RBF networks.

INSTITUTE OF COMMUNICATION, INFORMATION AND PERCEPTION TECHNOLOGIES
Scuola Superiore Sant'Anna

Retis
Real-Time Systems Laboratory

# Neural Networks and Deep Learning: Part 2: Deep Networks

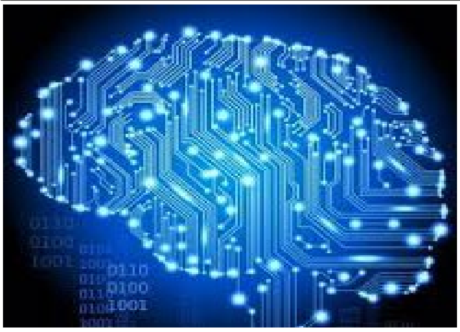
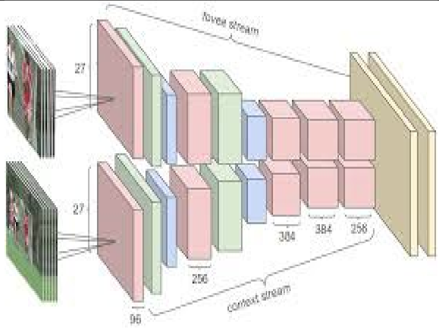

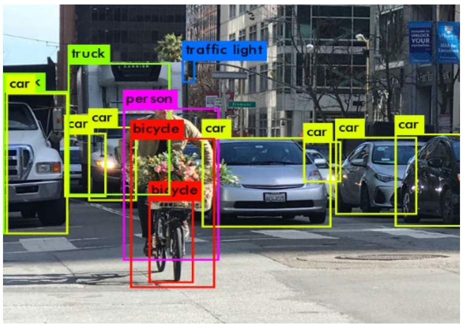


**The first lecture is scheduled on February 4, 2025 - 9:00**
Visit the course web page for registration and connecting to the channel.
http://retis.sssup.it/~giorgio/courses/neural/nn.html

**Giorgio Buttazzo**

**Objectives:** This module presents the foundations for understanding deep neural networks and deep learning algorithms. Topics include convolutional networks for classification, detection and segmentation, deep reinforcement learning, generative adversarial networks and transformers.

## Course program

1. **Towards deep networks**. Problems in training deep networks: overfitting and vanishing gradient. Solutions for deep learning: regularization, and dropout methods. Softmax output layer.
2. **Convolutional networks**. Basic principles and intuitions. Network architecture. Convolution operation. Convolution over a volume. Pooling layers. Normalization methods.
3. **Networks for classification**. Examples of CNNs and their key novelties: LeNet-5, AlexNet, VGG-Net, GoogLeNet, ResNet, Inception family, SqueezeNet, DenseNet, SENet, NASNet.
4. **Networks for object detection**. Two-stage detectors (the R-CNN family). Single-stage detectors (YOLO and SSD). Performance metrics. Pyramid networks (FPN, RetinaNet, the YOLO family).
5. **Networks for image segmentation**. Semantic, instance, and panoptic segmentation tasks. Fully Convolutional Network, U-Net, SegNet, ICNet, Mask R-CNN, YOLO-ACT.
6. **Deep Reinforcement Learning**. Deep Q-learning models: DQL, Double-DQL, Dueling-DQL, Rainbow DQL, Recurrent DQL. Policy gradient and actor-critic methods: Reinforce, DDPG, PPO.
7. **Recurrent neural networks**. Sequence to sequence models. Backpropagation through time. Natural language processing: GRU, LSTM, and Bidirectional networks. Language modeling. Word embedding. Machine translation. The attention mechanism.
8. **Generative adversarial networks**. Generative autoencoders, GANs, Style Transfer. Variational Autoencoders, Diffusion models.
9. **Transformers**. Encoder-Decoder architectures. Positional encoding. Self-attention and multi-head attention mechanisms. Masked attention.

INSTITUTE OF COMMUNICATION, INFORMATION AND PERCEPTION TECHNOLOGIES
Scuola Superiore Sant'Anna

# Neural Networks and Deep Learning: Part 3: Advanced Topics

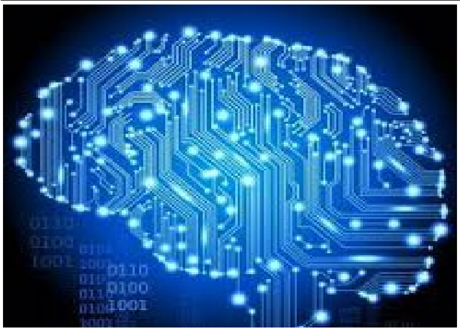
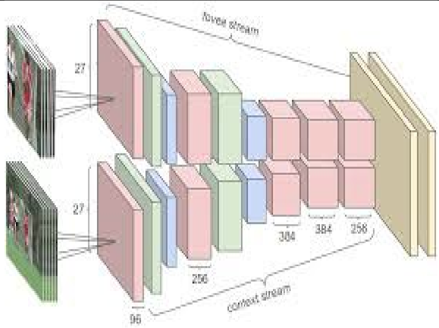

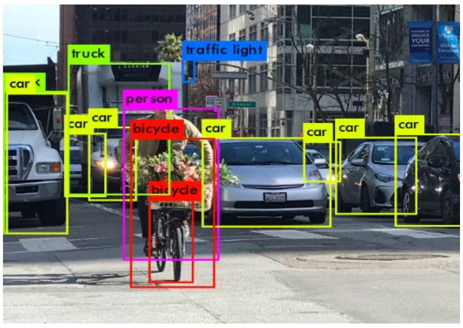


**The first lecture is scheduled on March 4, 2025 - 9:00**
Visit the course web page for registration and connecting to the channel.
http://retis.sssup.it/~giorgio/courses/neural/nn.html

**Giorgio Buttazzo, Federico Nesti, Giulio Rossolini**

**Objectives:** This module presents recent techniques proposed to improve previous models and overcome their limitations. Topics include model compression, semi-supervised learning, deep reinforcement learning, neural object tracking, adversarial attacks and defense methods.

## Course program

1. **Model compression**. Weight quantization. Model pruning. Model distillation for transferring the knowledge of large DNNs to small DNNs. Optimized architectures for mobile devices.
2. **Semi-supervised learning**. K-nearest neighbors. Self-training algorithms. Co-training. One-shot learning. Siamese networks. Zero-shot learning.
3. **Contrastive learning**. Principles of contrastive learning. Contrastive learning models: SimCLR, MoCo, Swav. Supervised Contrastive Learning.
4. **Neural networks for object tracking**. Generic object tracking. Tracking by detection: DeepSORT, FairMOT, BYTEtrack.
5. **Trustworthy AI**. Safety, security, and predictability issues in deep neural networks. Adversarial attacks and general defense mechanisms.
6. **Adversarial attacks and defenses**. White-box attacks: one-step vs. iterative approaches. Black – box attacks. Real-world adversarial attacks. Defense methods.
7. **Explainable AI**. Definition and taxonomy. Common approaches for image processing and tabular data. Gradient based methods.
8. **Anomaly detection and domain generalization**. Out-of-distribution detection. Domain Generalization and Domain Adaptation.
9. **Attention mechanisms in computer vision**: CNN-based strategies and Visual Transformers. Recent Transformer architectures in Computer Vision

INSTITUTE OF COMMUNICATION, INFORMATION AND PERCEPTION TECHNOLOGIES
Scuola Superiore Sant'Anna

Retis
Real-Time Systems Laboratory

# Neural Networks and Deep Learning: Part 4: Implementation Issues

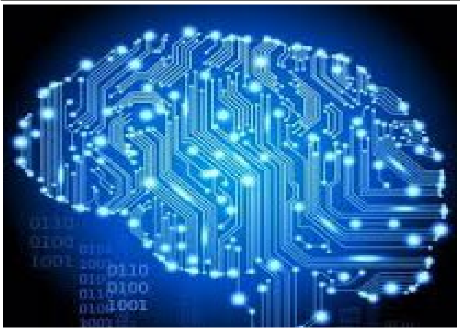


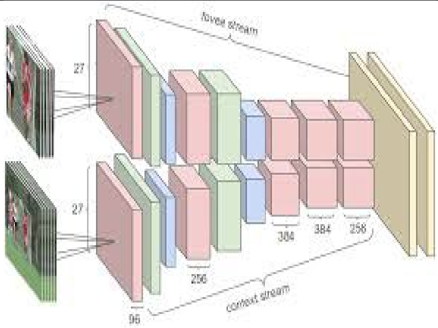


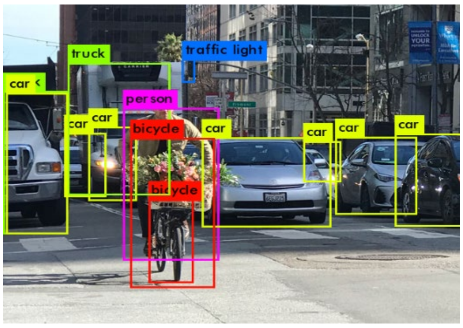


**The first lecture is scheduled on April 8, 2024 - 9:00**
Visit the course web page for registration and connecting to the channel.
http://retis.sssup.it/~giorgio/courses/neural/nn.html

**Giorgio Buttazzo, Alessandro Biondi, Daniel Casini, Federico Aromolo**

**Objectives:** The aim of this part is to discuss practical and implementation issues useful to deploy neural networks on a variety of embedded platforms using different languages and development environments.

## Course program

1. **Implementing Neural Networks from scratch in C**. General implementation principles. Implementing Reinforcement Learning (ASE/ACE and Q-learning). Implementing feed-forward networks.
2. **Development frameworks**. Overview of the existing frameworks. Common data sets. Common frameworks: Keras, Tensorflow, Caffe, and Pytorch. Examples of neural network implementations.
3. **Functional components on autonomous driving**. Basic blocks for perception, prediction, planning, control, and actuation.
4. **The Apollo framework for autonomous driving**. Overview of the framework. Neural networks in Apollo. Neural models for perception and prediction.
5. **Simulators for autonomous driving**. Overview of existing simulators. The CARLA simulator. Practical examples.
6. **DNN optimization for embedded platforms.** Weight quantization, pruning, distillation, to reduce execution times and contain memory footprints in resource constrained platforms.
7. **Accelerating deep networks on GPGPUs**. GPU programming in CUDA. Overview of the Nvidia TensorRT framework. Executing a DNN modelled in Caffe in TensorRT.
8. **GPU-based real-time neural vision**. How to accelerate a neural network on TensorRT to detect objects from a video camera.
9. **Accelerating deep networks on FPGA**. Technologies and approaches. Overview of the Xilinx Deep Processing Unit to execute convolutional networks.