ICVSS 2007 – International Computer Vision Summer School 2007

Detection, Recognition and Segmentation in Context

Sicily ~ 9-14 July 2007

Supervised and Unsupervised Methods for Learning Invariant Feature Hierarchies

Yann LeCun

Yann LeCun, Courant Institute, New York University

Abstract

The first part of the lecture will give an introduction to supervised training methods for discriminative Energy-Based Models, including discriminative HMMs, graph transformer networks, conditional random fields and max-margin Markov nets. I will show how such models can be trained to segment and recognize compositional objects.

The second part will address the issue of training "deep architectures". Many models of object recognition in humans and animals are hierarchical, comprising multiple levels of increasingly abstract, global, and ivariant features. By contrast, many artificial object recognition systems are "shallow", comprising a single layer of features, followed by a generic classifier.

One example of deep architecture for object recognition is the convolutional network model. Supervised learning algorithms for convolutional nets will be reviewed. However, purely supervised methods have the disadvantage of requiring many labeled training samples.

In recent years, a number of unsupervised learning algorithms for deep architectures have emerged, which greatly reduce the number of necessary labeled samples. Hinton's restricted Boltzmann machines and contrastive divergence algorithm will be described. Finally, unsupervised algorithms that can learn sparse and invariant feature hierarchies will be introduced.

Real-time demos of a face detector and an object recognition system based on convolutional nets will be shown.