



Visual Recognition over Multiple Categories: A Scalable Framework

Amnon Shashua
The Hebrew University of Jerusalem, Israel

Abstract

Single category visual recognition, i.e., detecting an object category from an image has witnessed great success in both academic and industrial contexts. For example, systems today exhibit impressive performance both in accuracy and realtime efficiency when detecting faces, cars or pedestrians from images. However, human vision can effortlessly handle thousands of object classes. The issue of how to cope with the classifier's accuracy and resource complexity when dealing with a growing number of object categories largely remains an open problem. In the class, I will review the basic elements of 2-class learning - specifically the hinge-loss and logistic loss - and then work on the extension to multiple classes. We will then work out the methods for applying the theory to real-world problems involving multiple classes including character recognition and multiclass visual recognition.

Syllabus: Visual Recognition, Multiclass Learning