Efficient Novel-Class Recognition and Search

Lorenzo Torresani
Dartmouth College, USA

Abstract

Over the last decade, research in computer vision and machine learning has produced dramatic improvements in object categorization, the problem of determining "what is present in a picture". These advances stem primarily from two innovations: first, new low-level image features specifically designed to capture different visual cues, such as shape, texture, and scene layout; second, methods to learn powerful nonlinear classifiers that can effectively combine these multiple, complementary image descriptors.

After reviewing some of the key models and techniques in this area, we will study the challenges in scaling these approaches to the problem of class recognition in large image collections. We will focus specifically on scenarios where the classes to be recognized are not known in advance. The motivating application is "object-class search by example" where a user provides at query time a small set of training images defining an arbitrary novel category and the system must retrieve images belonging to this class from a large database. This setting poses challenging requirements on the system design: the object classifier must be learned efficiently at query time from few examples; recognition must have low computational cost with respect to the database size; finally, compact image descriptors must be used to allow storage of large collections in memory.

We will study methods that address these requirements by learning compact image descriptors optimized to yield good categorization accuracy with efficient classifiers. We will also review how data structures and methods from text-retrieval can be adapted to enable real-time search of an object-class in a large image collection.

Syllabus: Novel-Category Recognition, Classification using Multiple Kernel Learning, Learning Compact Image Descriptors