Machine learning methods for large-scale image search and recognition

Rob Fergus
New York University, USA

Abstract

Existing object recognition methods are not practical to use on huge image collections found on the Internet and elsewhere. Recently, a number of computer vision approaches have been proposed that scale to millions of images. At their core, many of them rely on machine learning techniques to learn compact representations that can be used in efficient indexing schemes. My talk will review this work, highlighting the common learning techniques used and discuss open problems in this area.

Syllabus: *Large scale computer vision, Internet vision, Distance metric learning, Hashing techniques*