Part-based models for finding people and objects

Deva Ramanan
University of California at Irvine, USA

Abstract

Object recognition is a core challenge in vision, with people being perhaps the most interesting objects to recognize. Much progress has been made in recent years, as evidenced by large-scale benchmark competitions such as the PASCAL Visual Object Challenge. This tutorial will review some of the more successful techniques based on part-based models, attempt to uncover the crucial components that make them work.

Specifically, it will revisit rather classic vision issues of model-based shape, statistical appearance models, and occlusion-reasoning. While some would argue that detection systems have plateaued, I will argue that we are just getting started; we now have tools ask more interesting questions regarding three-dimensional shape, detailed human pose, and spatial object interactions in truly unconstrained, "in-the-wild" imagery. The tutorial will conclude with such opportunities for future research directions.

Keywords: Part-based models, Object recognition, Action Recognition