



## Instance-level Visual Recognition

Georgia Gkioxari

Facebook, US

### Abstract

In this lecture I will discuss the recent advances on instance-level visual recognition. I will cover in detail the most recent work on object detection, instance segmentation and human pose prediction from a single image. I will give a detailed tour of our recently released codebase, Detectron, that allows for a fast and efficient implementation of the models described in the lecture. Going beyond traditional object recognition tasks, I will introduce recent efforts in recognizing object interactions from images and I will describe our recent state of the art model for this task. Last, I will discuss how our efforts on image-based recognition can be extended to tackle video related tasks, and in particular video pose tracking from real videos.