

Efficient Inference Algorithms for Scene Understanding Problems

Karteek Alahari
LEAR tem, INRIA Rhône-Alpes

One of the goals of computer vision is to interpret a scene semantically from an image or a video. This problem has manifested itself in various forms, including, but not limited to, object recognition, 3D scene recovery, image segmentation, human pose estimation. Addressing these problems by casting them as discrete labelling tasks has received significant attention over the past ten years or so. Although tremendous progress has been made with this framework, many challenges still remain, especially in the context of large-scale data. This talk will focus on our work addressing challenges in the context of large-scale modelling, parameter learning and inference. In particular, (i) image priors with higher-order and global constraints, which have the ability to encode significantly more sophisticated structural dependencies among image pixels, compared to traditional pairwise interactions; (ii) temporal constraints to capture rich motion and appearance cues in video.