Discriminative Clustering for Image Co-segmentation

Francis Bach
(ENS et Projet Willow INRIA), January 29, 2010

Purely bottom-up, unsupervised segmentation of a single image into two segments remains a challenging task for computer vision. The co-segmentation problem is the process of jointly segmenting several images with similar foreground objects but different backgrounds. In this work, we combine existing tools from bottom-up image segmentation such as normalized cuts, with kernel methods commonly used in object recognition. These two sets of techniques are used within a discriminative clustering framework: we aim to assign foreground/background labels jointly to all images, so that a supervised classifier trained with these labels leads to maximal separation of the two classes. In practice, we obtain a combinatorial problem which is relaxed to a continuous convex optimization problem, that can be solved efficiently for up to dozens of images (using appropriate manifold optimization techniques). We show that our framework works well on images with very similar foreground objects, which are usually considered in the literature, as well as on more challenging problems with objects with higher intra-class variations.