

Abstract Peter Gehler

Learning Bilateral Filters

In this talk I will present our recent work on learning filters for sparse and high dimensional signals. In particular we generalize the parametrization of the widely used bilateral filter and derive an algorithm so filter parameters can be learned from data. This generalization has interesting consequences for a wide range of computer vision problems and learning architectures. In particular it allows for CNN architectures that have no spatial but higher dimensional organization of intermediate representations. These can be used to naturally handle sparse input data. I will present results on several computer vision tasks, including image segmentation, and image filtering applications. For the task of semantic image segmentation I will demonstrate how this type of filtering can replace the recent so-called dense CRF inference approaches. It effectively removes the need for CNN-CRF combinations that are in common for segmentation.