


Multi-scale Context Aggregation by Dilated Convolutions

F. Yu, and V. Koltun (ICLR 2016)



- ▶ Convolutional network module that is specifically designed for dense prediction (semantic segmentation)
- ▶ Dilated convolutions to systematically aggregate multi-scale contextual information without losing resolution
- ▶ "The dilated convolution operator can apply the same filter at different ranges using different dilation factors."
- ▶ Front end module: VGG16 with deconvolutions (FCN) by removing the last two pooling and striding layers
- ▶ Front end is already too good: outperforms both FCN-8s and the DeepLab, and even DeepLab+CRF
- ▶ Identity initialization for the context module
- ▶ Trained on Microsoft COCO and VOC-2012 and tested on VOC-2012