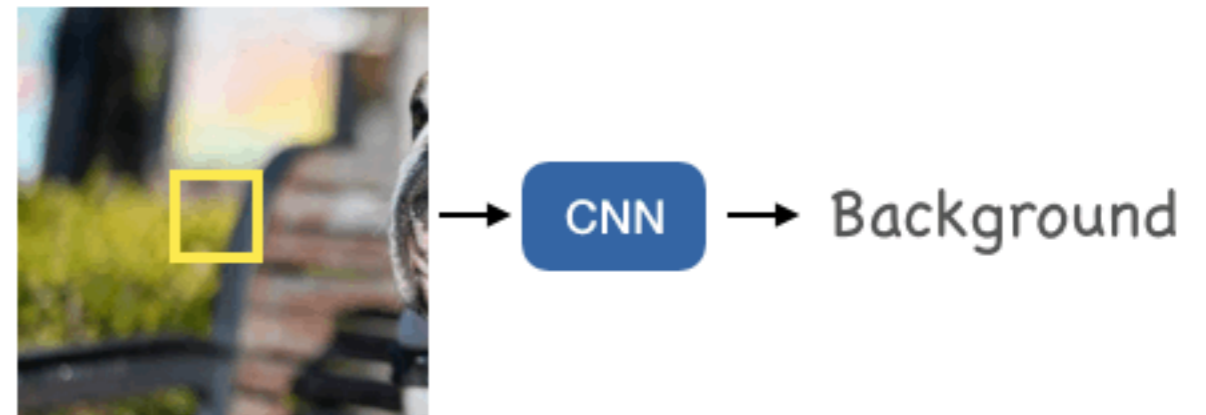


Case study: FCN

© 2019 Philipp Krähenbühl and Chao-Yuan Wu

Fully convolutional networks

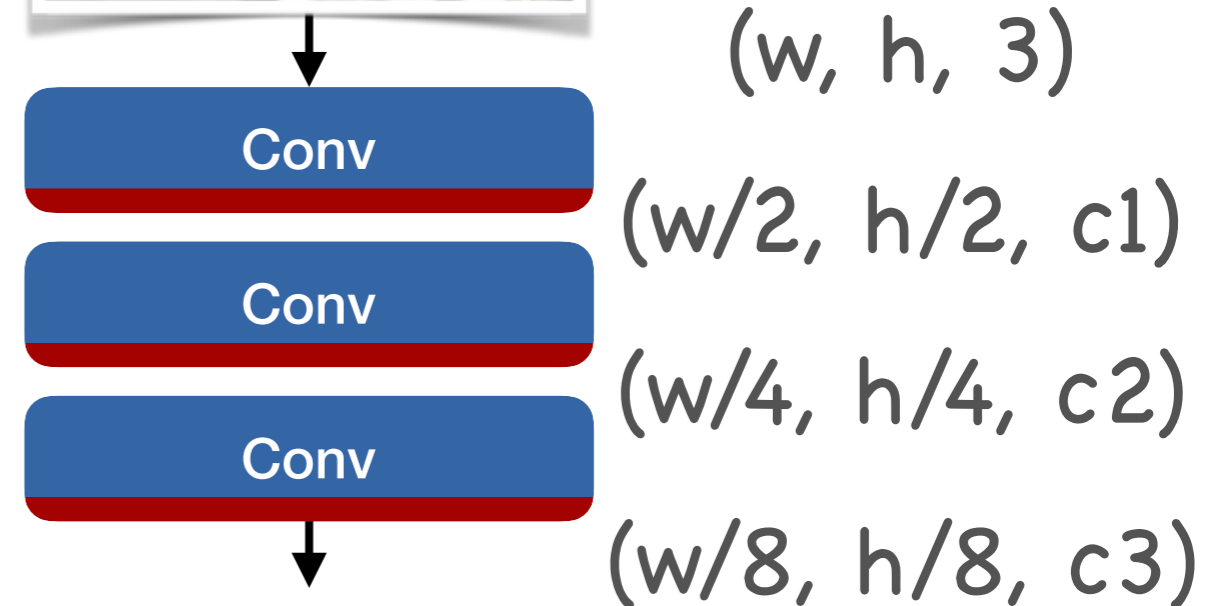
- Semantic segmentation by classifying each pixel
 - Fully convolutionally
 - Cross entropy loss at every pixel



Fully Convolutional Networks for Semantic Segmentation, Shelhamer et al., CVPR 2015

Output resolution

- Striding in network reduces output resolution
- Output segmentation 16–32x smaller
- Solution: Upsampling
 - Naive: Linear
 - Up-convolution



Practical considerations

- Output should be at center of receptive field
- Do the math, or NaN test to verify



How to build an FCN

- Take a classification network
 - Convert linear layers to convolutions
 - Add up-sampling / up-conv
 - Add optional skip connections
- Train fully convolutionally
 - Image-net pre-training helps



Conv

⋮

Conv

Average Pool

Linear (4096)

Linear (4096)

Linear (1000)