Case study: Dilated convolutional networks
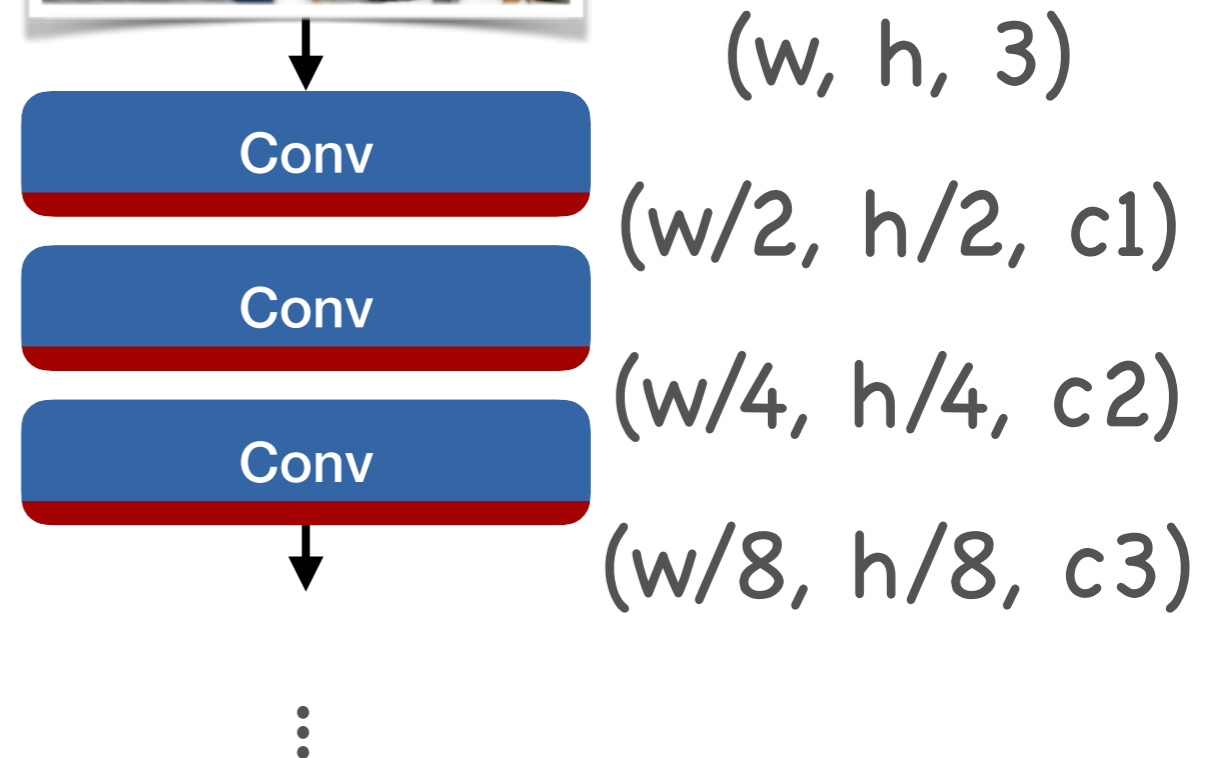
Dilated convolutional networks

- FCNs have a fairly low resolution

  - Solution: Replace strided convolutions with dilated convolutions

(w, h, 3)

Conv

(w/2, h/2, c1)

Conv

(w/4, h/4, c2)

Conv

(w/8, h/8, c3)

⋮

Semantic image segmentation with deep convolutional nets and fully connected CRFs, Chen etal. ICLR 2015
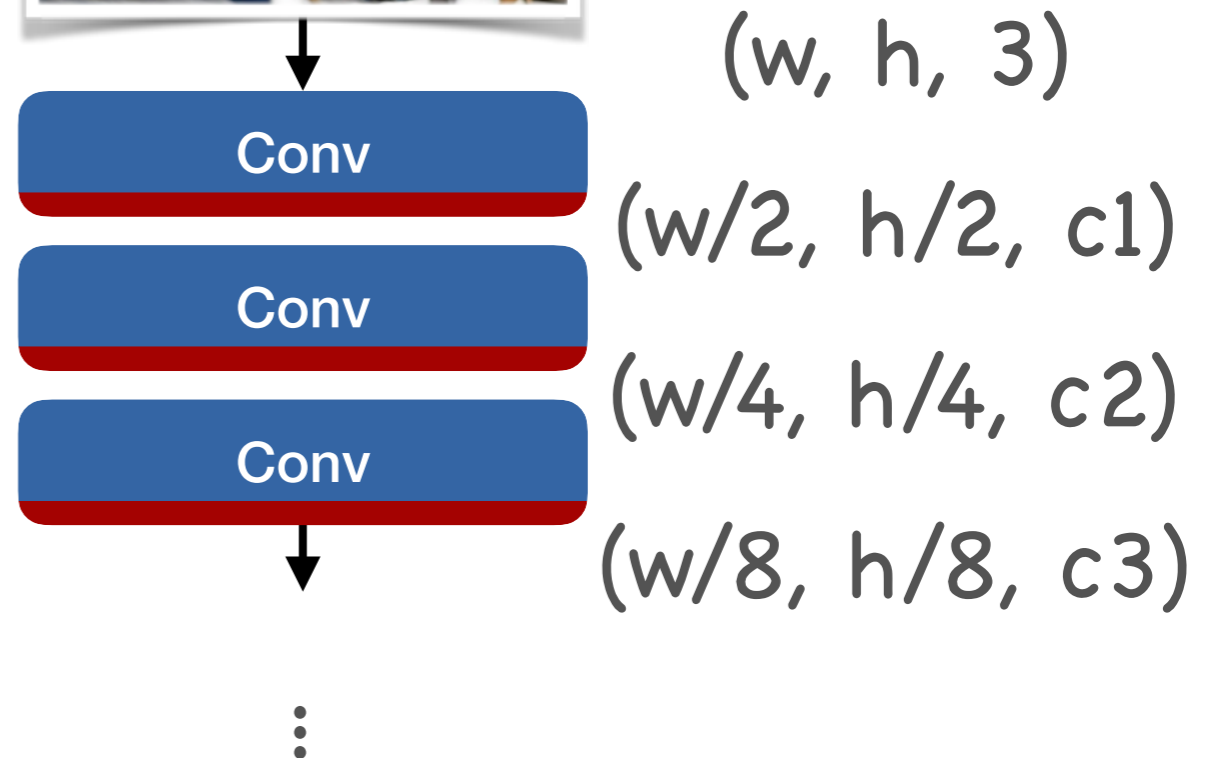
Extension: Context module

- Reason about class-wise interactions

  - Dilated layer on top of class output

- Alternatives: Graphical models

(w, h, 3)

(w/2, h/2, c1)

(w/4, h/4, c2)

(w/8, h/8, c3)

Multi-Scale Context Aggregation by Dilated Convolutions, Yu and Koltun, ICLR 2016

Extensions: Deformable convolutions

- Allow inputs to convolutions to be sampled irregularly

Deformable Convolutional Networks, Dai et al., ICCV 2017

Extensions: Deformable convolutions

Image source: Deformable Convolutional Networks, Dai et al., https://arxiv.org/pdf/1703.06211.pdf