




Dilation and Upconvolution

Receptive Field – Striding



Segmentation


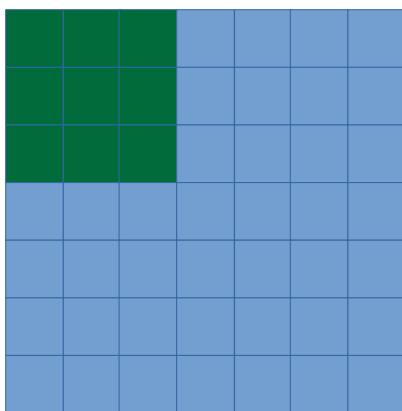
Instead of classifying the whole image, classify each pixel



We need to keep all of the spatial resolution

Dilation

Increase the receptive field without increasing kernel size or losing resolution




Holes

A trous


a	b	c
d	e	f
g	h	i

a	b	c
d	e	f
g	h	i


Receptive Field – Dilation



Upconvolution



Dilate the input

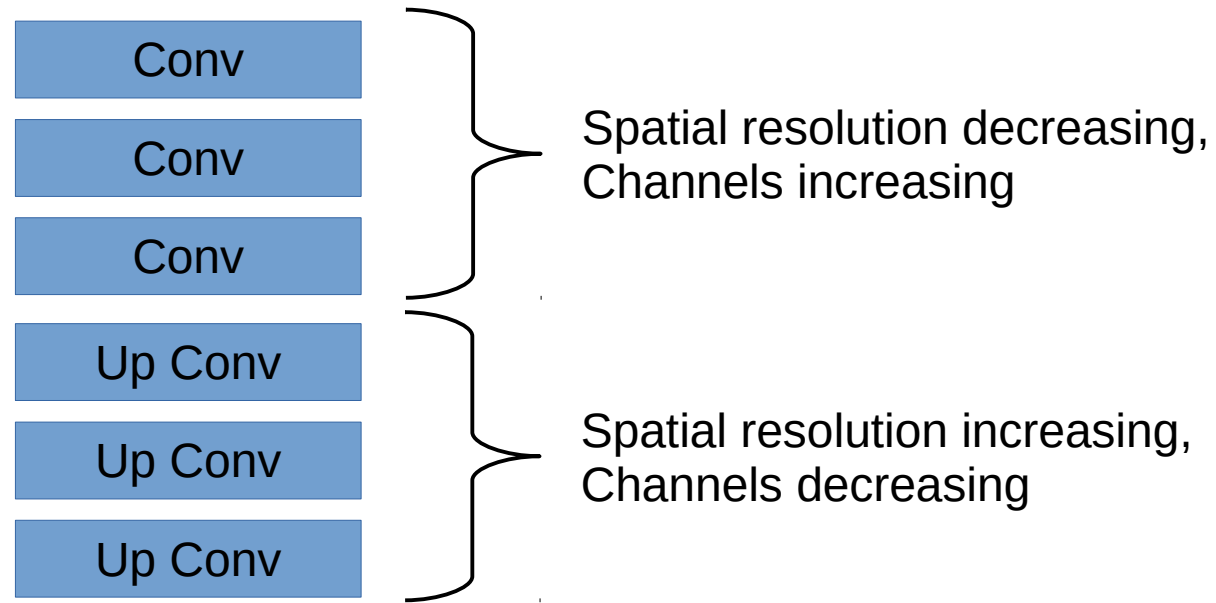


Transpose convolution

“Deconvolution”

Fractionally strided convolution

Upconvolution

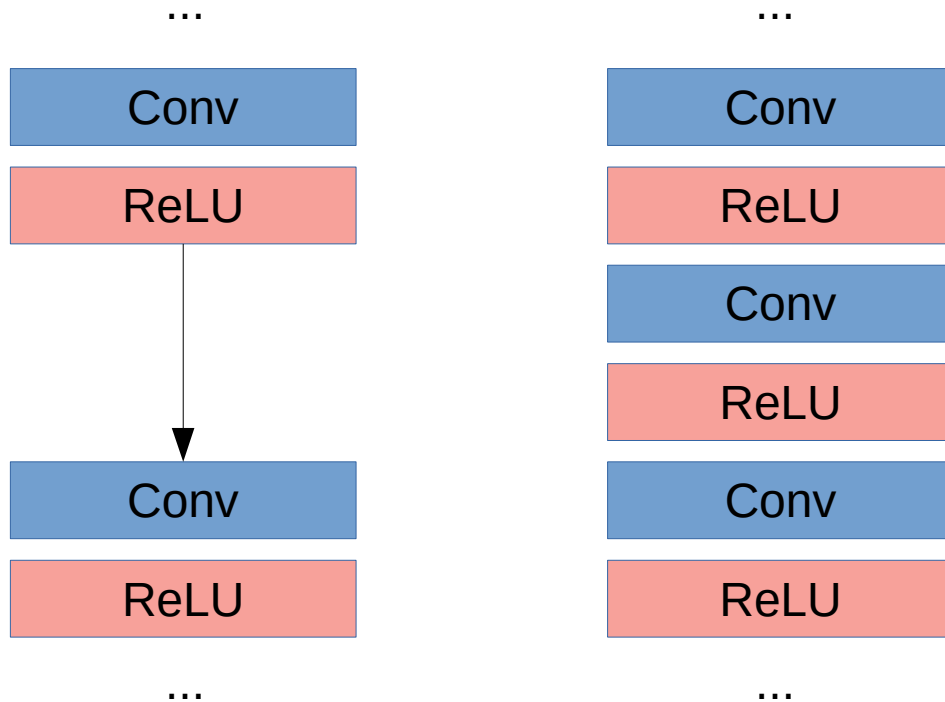




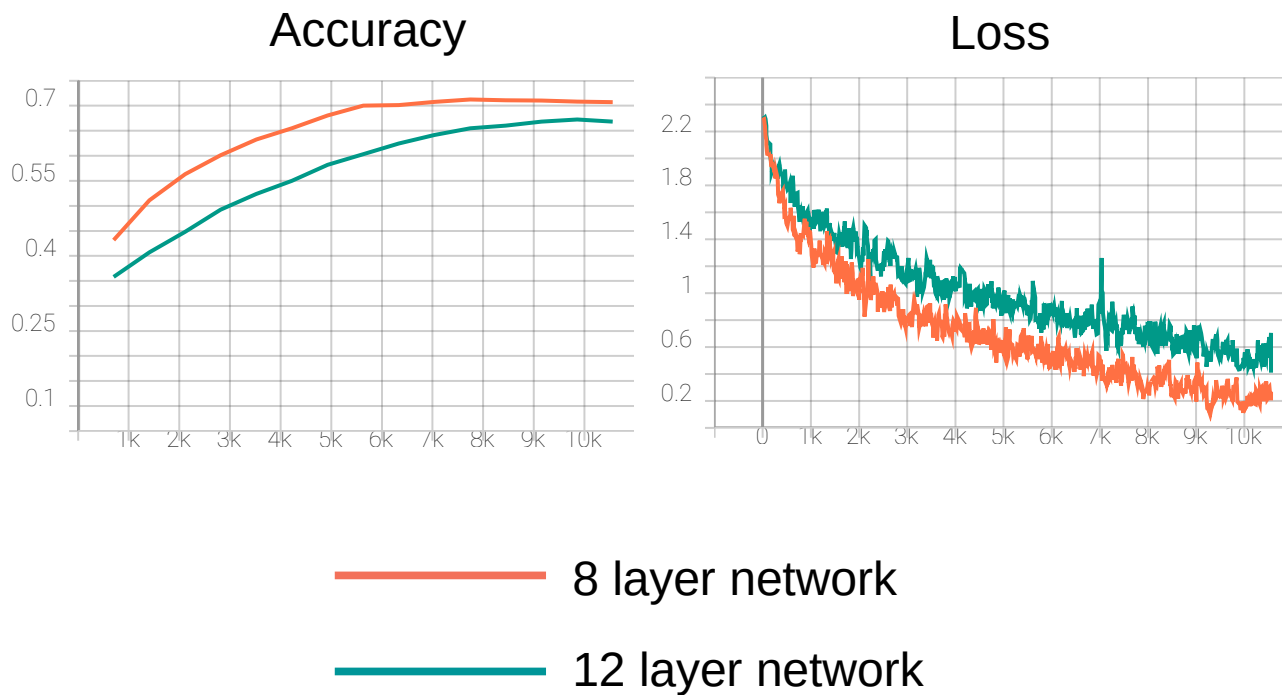
Residual Connections

Deep Networks are Hard to Train

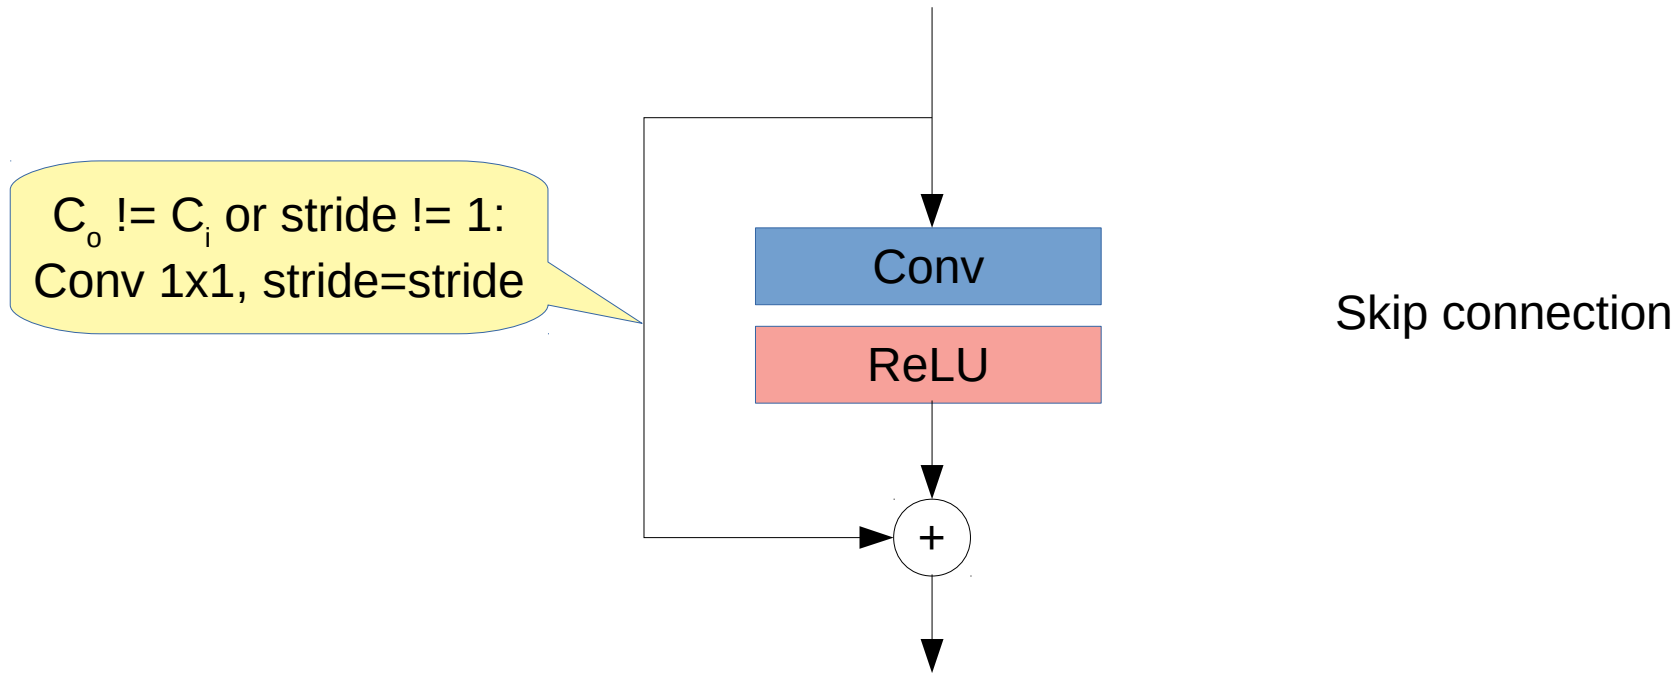
- We can train maybe 10-15 layers with the techniques we have seen so far.
- 20-30 with some other techniques we will see next week.



Deep Networks are Hard to Train




Residual Connections




Residual Connections – Gradient

Forward Pass



Backward Pass



Training networks
with up to 1000
layers

Upconvolution – Residual Connections

