CNN architectures and visualization

AlexNet

- Won the ImageNet competition 2012
- Start of deep learning revolution in computer vision

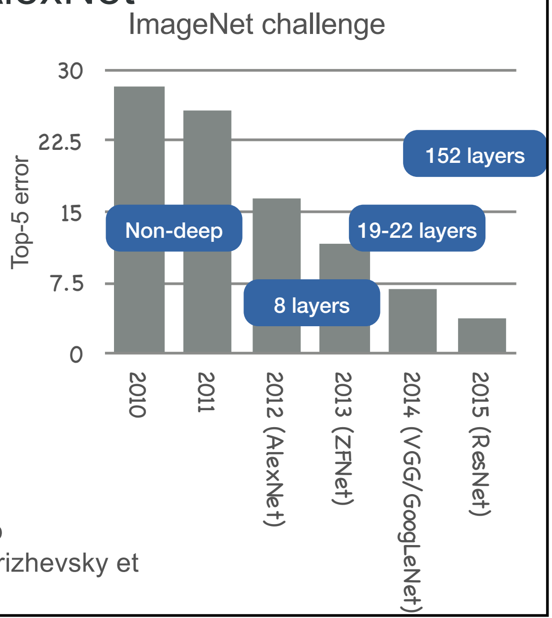

ImageNet Classification with Deep Convolutional Neural Networks, Krizhevsky et al., NIPS 2012

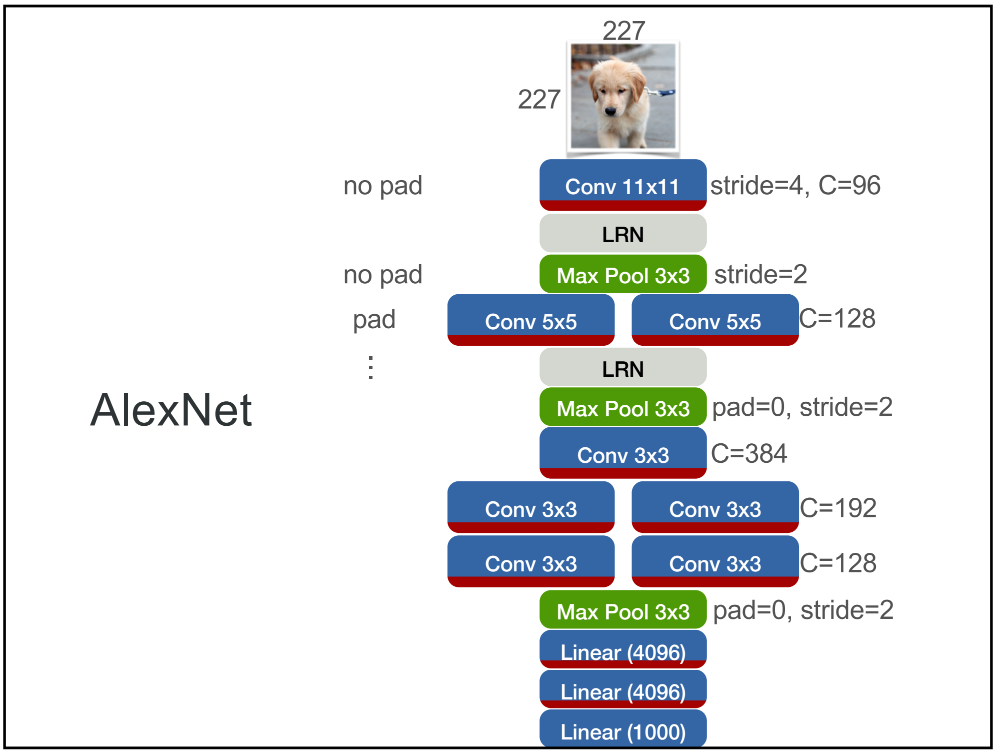
AlexNet
227
227
no pad
Conv 11x11
stride=4, C=96
LRN
no pad
Max Pool 3x3
stride=2
pad
Conv 5x5
Conv 5x5
C=128
⋮
LRN
Max Pool 3x3
pad=0, stride=2
Conv 3x3
C=384
Conv 3x3
Conv 3x3
C=192
Conv 3x3
Conv 3x3
C=128
Max Pool 3x3
pad=0, stride=2
Linear (4096)
Linear (4096)
Linear (1000)

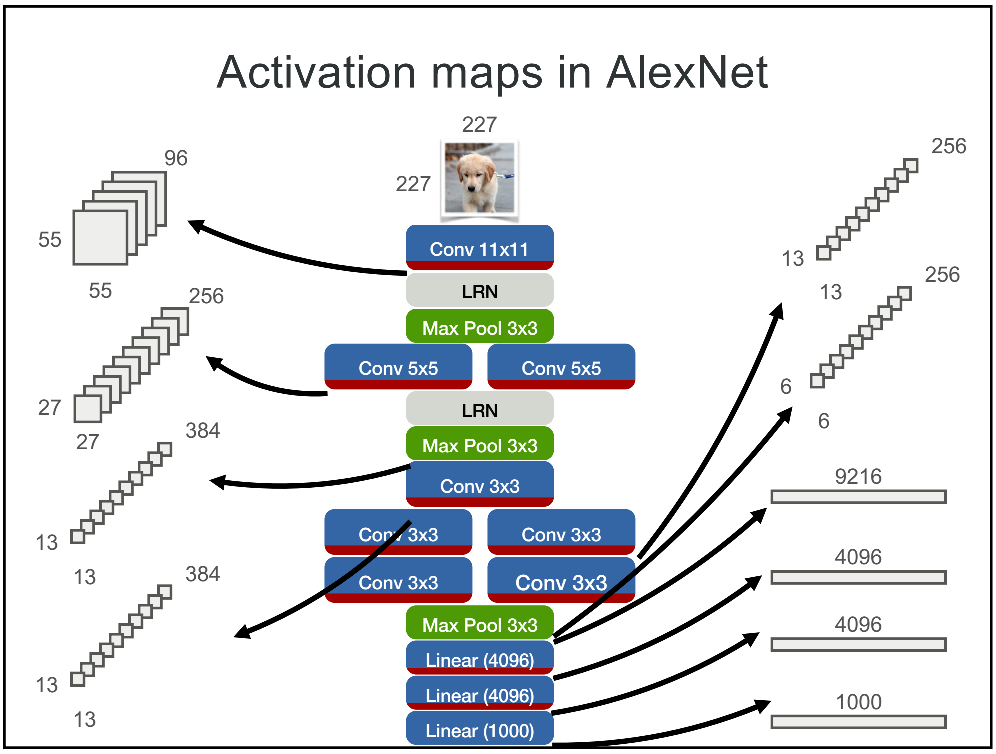
Activation maps in AlexNet
227
227
Conv 11x11
LRN
Max Pool 3x3
Conv 5x5
Conv 5x5
LRN
Max Pool 3x3
Conv 3x3
Conv 3x3
Conv 3x3
Conv 3x3
Conv 3x3
Max Pool 3x3
Linear (4096)
Linear (4096)
Linear (1000)
96
55
55
256
27
27
384
13
13
384
13
13
256
13
13
256
6
6
9216
4096
4096
1000

Activation maps in AlexNet

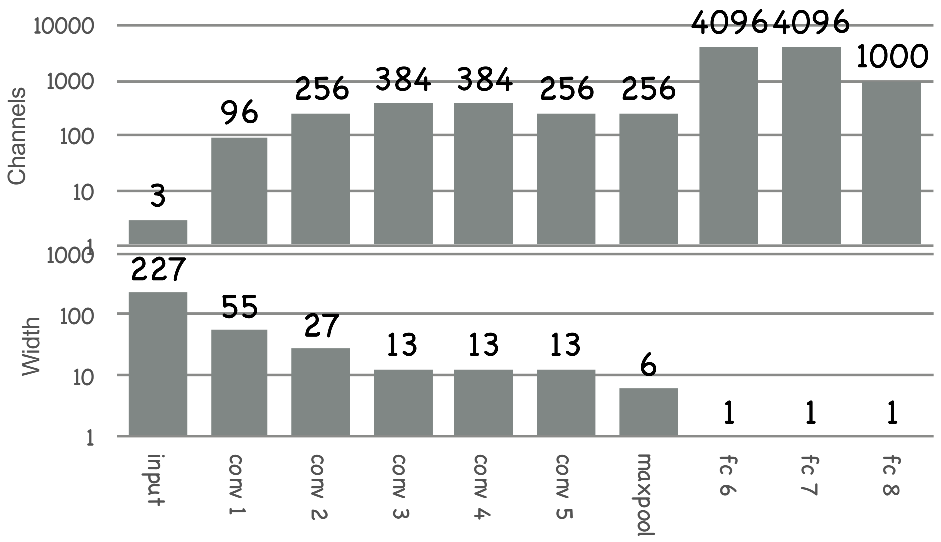

Parameters and computation

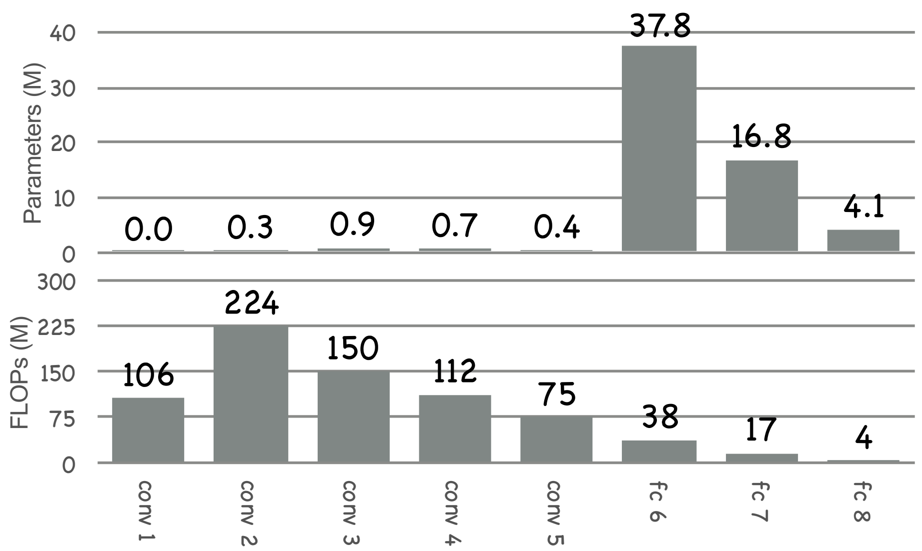

What's happening inside the black box?

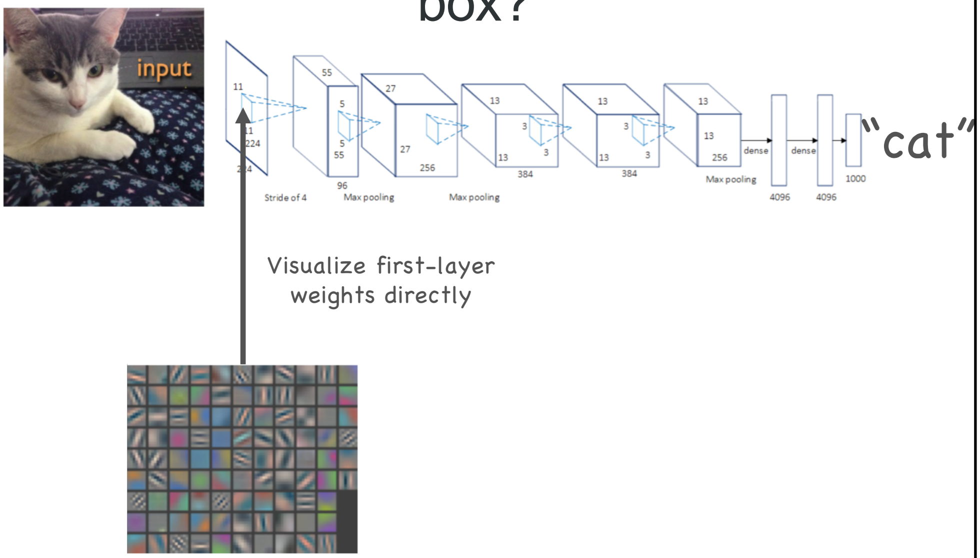

What's happening inside the black box?

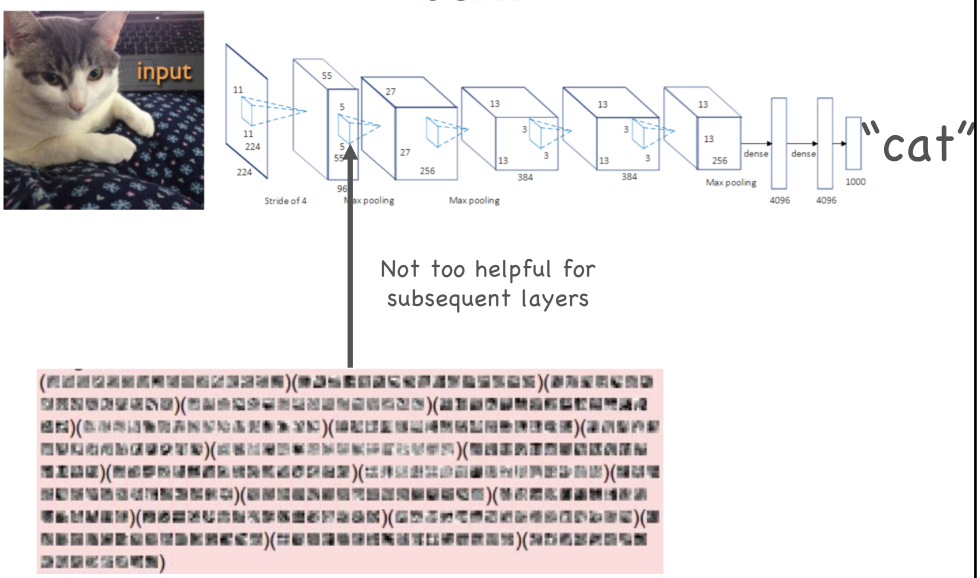

Features from a CIFAR10 network, via [Stanford CS231n]

What's happening inside the black box?

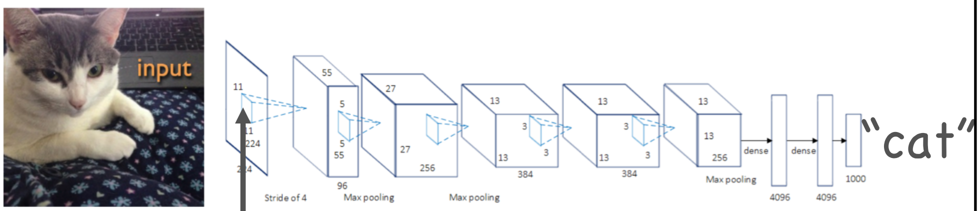

Visualize maximally activating patches:
pick a unit; run many images through the
network; visualize patches that produce
the highest output values

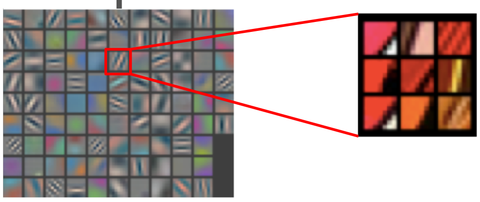

What's happening inside the black box?

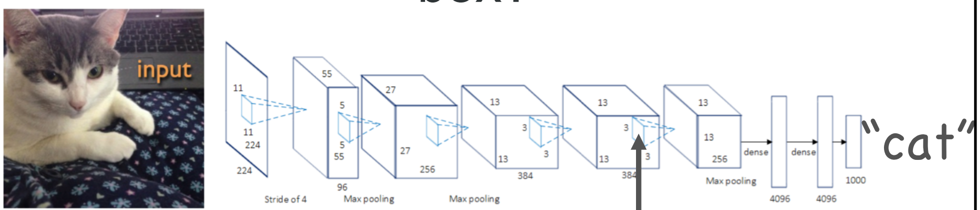

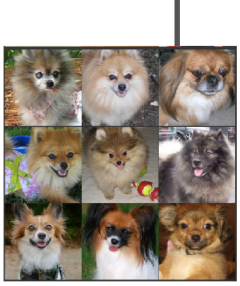

Visualize maximally
activating patches

What's happening inside the black box?

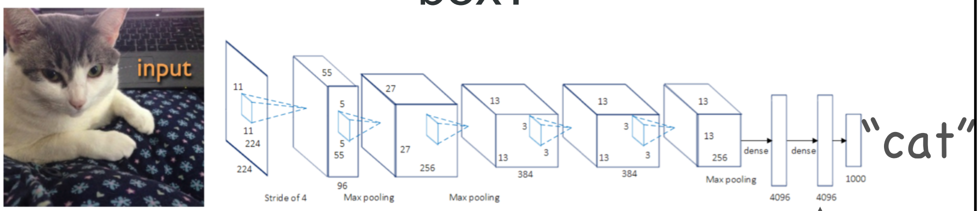

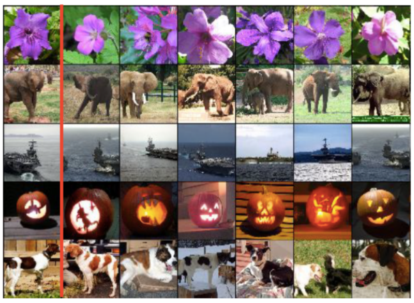

What about FC layers?

Visualize nearest neighbor images according to activation vectors

Source: Stanford CS231n

What's happening inside the black box?

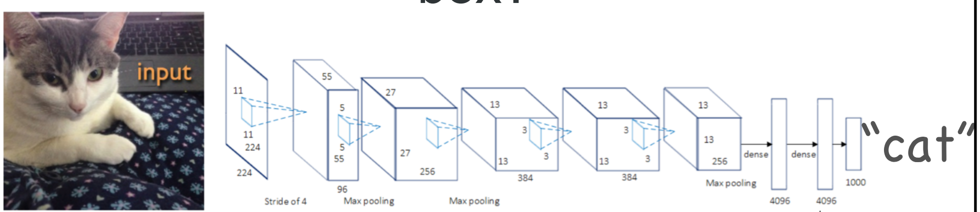

What about FC layers?

Fancy dimensionality reduction, e.g., t-SNE

Source: Andrej Karpathy

What's happening inside the black box?

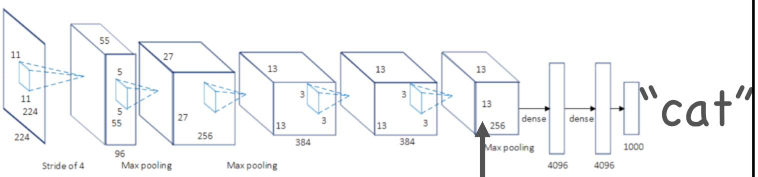

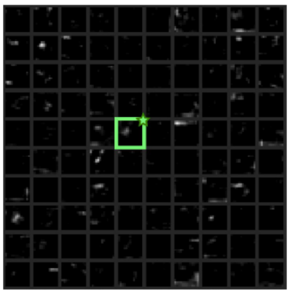

Source

Visualize activations for a particular image

What's happening inside the black box?

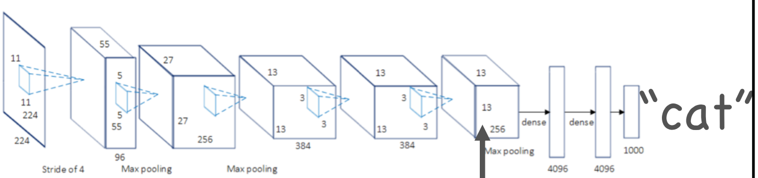

Visualize pixel values responsible for the activation

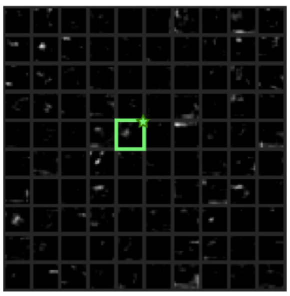

Source

Visualize activations for a particular image

1. Unpool
2. Rectify
3. Deconvolve

Deep visualization toolbox

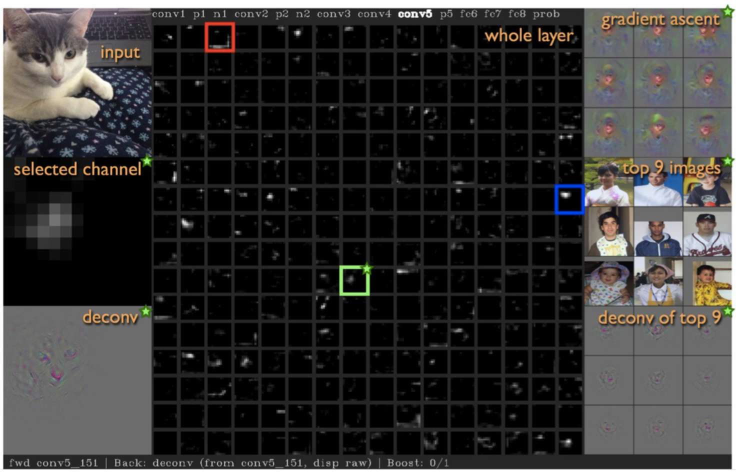

[YouTube video]

J. Yosinski, J. Clune, A. Nguyen, T. Fuchs, and H. Lipson, [Understanding neural networks through deep visualization], ICML DL workshop, 2015

Additional visualization techniques

K. Simonyan, A. Vedaldi, and A. Zisserman, [Deep Inside Convolutional Networks: Visualising Image Classification Models and Saliency Maps], ICLR 2014

J. Springenberg, A. Dosovitskiy, T. Brox, M. Riedmiller, [Striving for simplicity: The all convolutional net], ICLR workshop, 2015

A. Nguyen, J. Yosinski, J. Clune, [Multifaceted Feature Visualization: Uncovering the Different Types of Features Learned By Each Neuron in Deep Neural Networks], ICML workshop, 2016

ZFNet

- Won the ImageNet competition 2013
- Nice analysis and visualization of AlexNet
  - Introduced up-conv

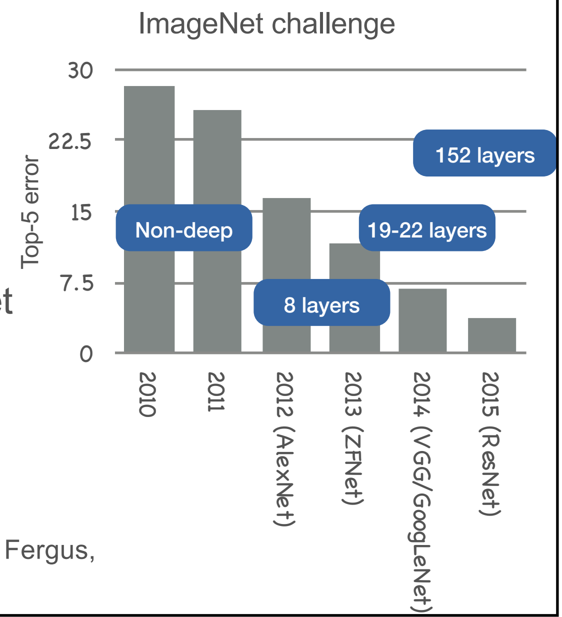

Visualizing and Understanding Convolutional Networks, Zeiler & Fergus, ECCV 2014

AlexNet to ZFNet

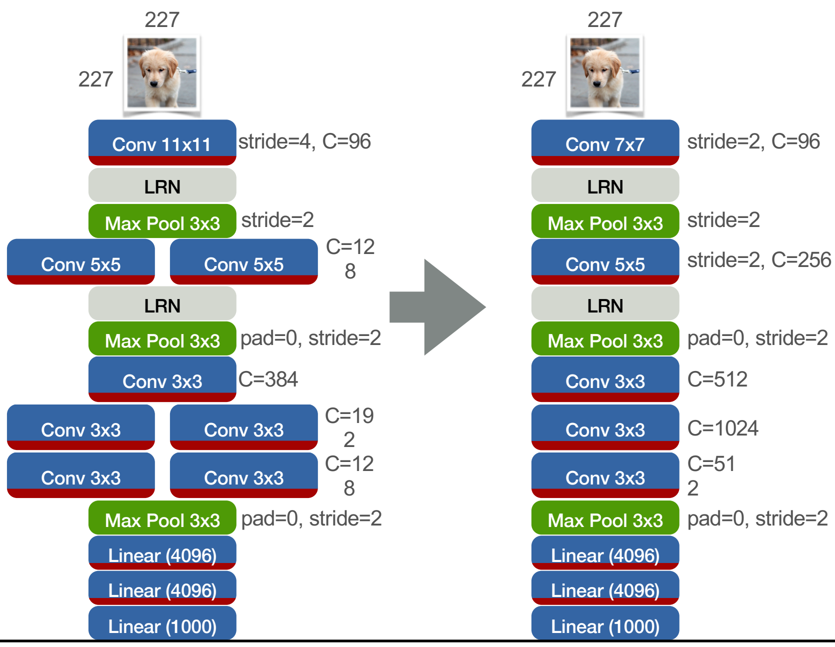

Activation maps in ZFNet

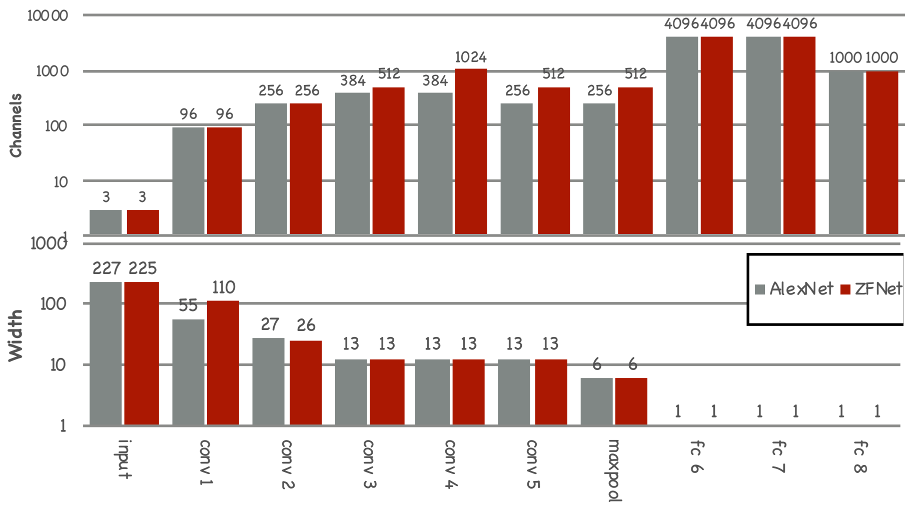

Parameters and computation

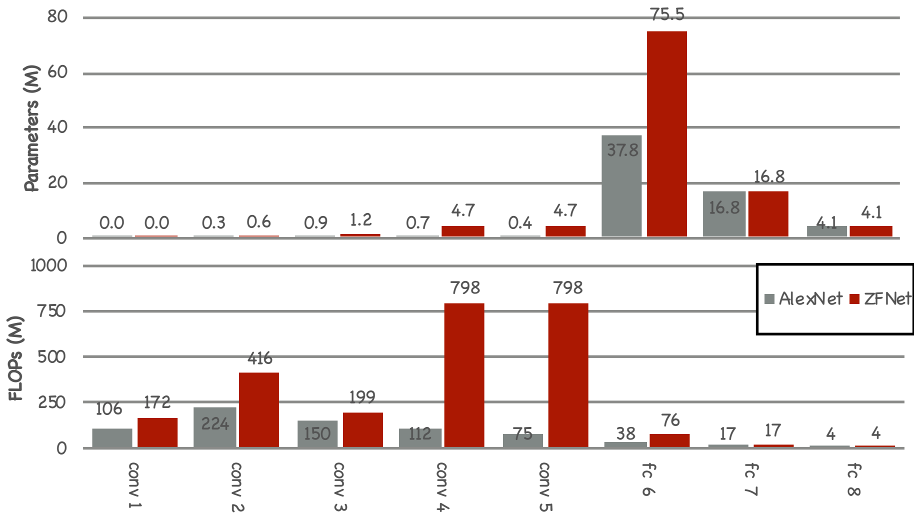

VGG

- Deeper AlexNet/ZFNet

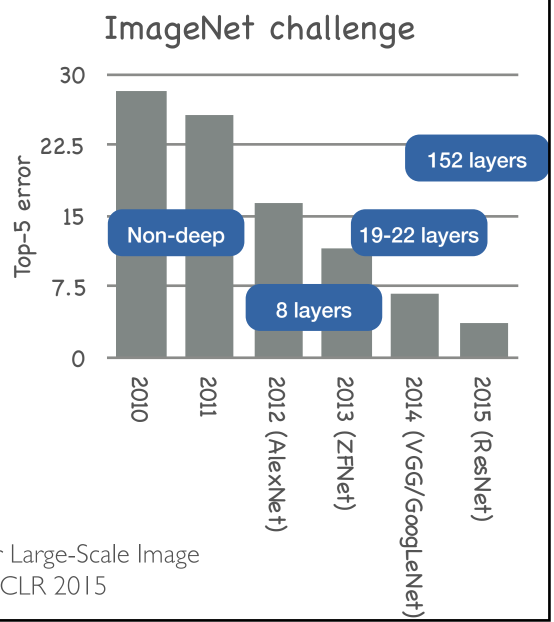

Very Deep Convolutional Networks for Large-Scale Image Recognition, Simonyan and Zisserman, ICLR 2015

ZFNet to VGG

Insights in VGG

- Why use smaller filters?
  - Factorization

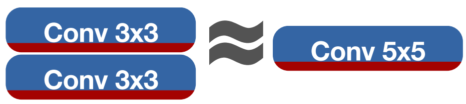

Training VGG

- Vanishing gradients

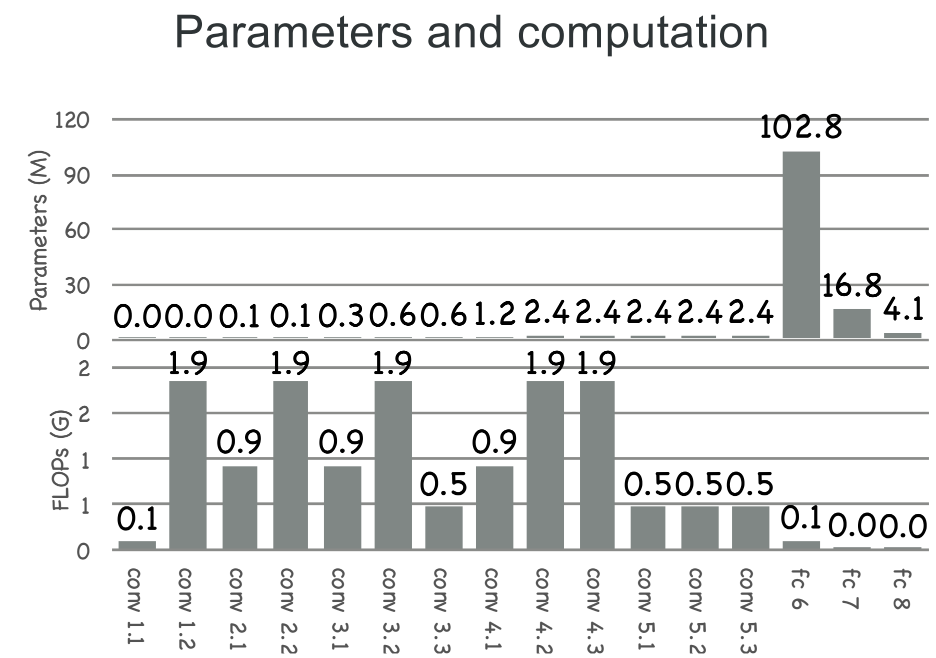

1x1 convolutions and factorization

- Convolutions are linear operators
  - Can we make them non-linear?

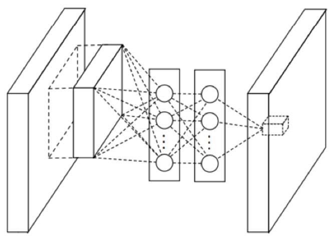

Figure source: "Network-in-Network", ICLR 2014 paper

Network-in-Network, Lin et al., ICLR 2014

Factorized convolution

- Implements "non-linear" convolution
- Less computation, fewer parameters

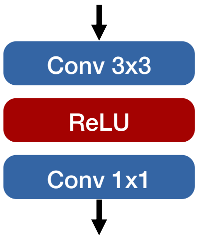

Inception architecture

- GoogLeNet
  - Evolution of Network-in-Network

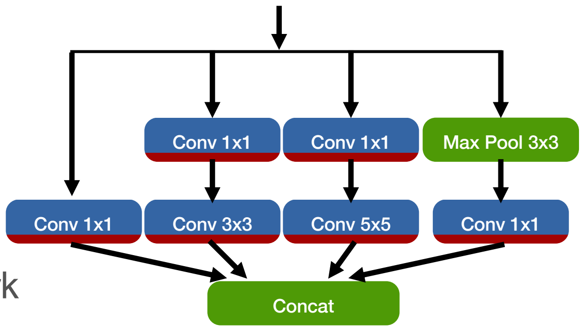

Going deeper with convolutions, Szegedy et al. CVPR 2015.

Inception architecture

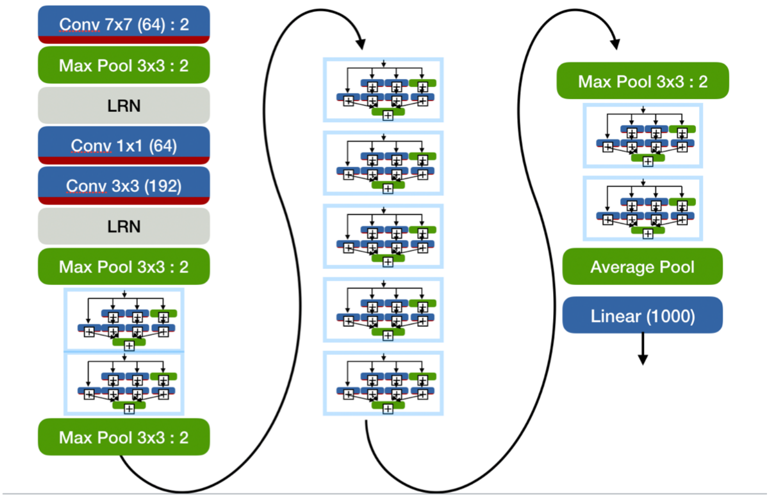

Residual Networks

- Uses residual connections to build deeper networks
    - Activation maps are additive

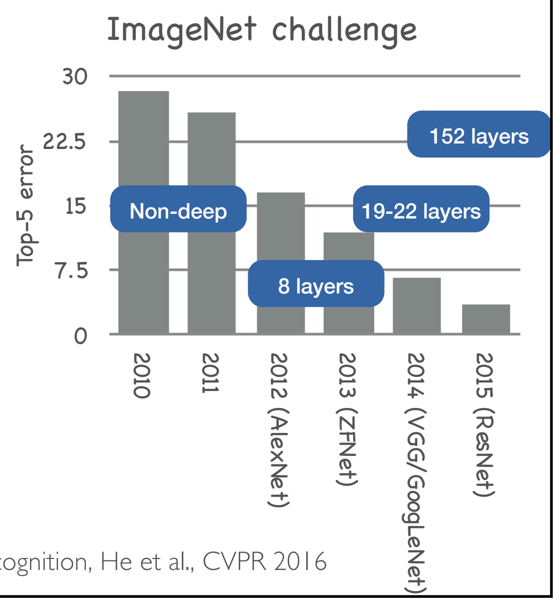

Deep Residual Learning for Image Recognition, He et al., CVPR 2016

Residual blocks

- Add shortcut connections for gradients
  - Identity
  - Strided 1x1 convolution

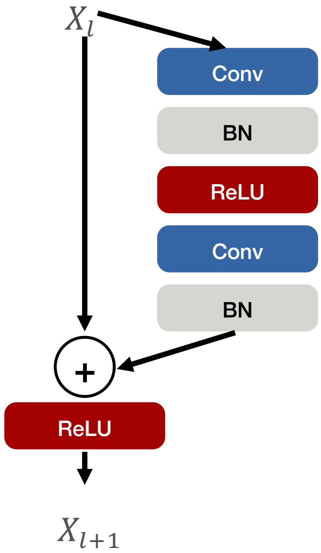

ResNet

- Multiple variants
  - ResNet-18, ResNet-34, ResNet-50, ResNet-101, ResNet-152, ResNet-1001