

CMP784


DEEP LEARNING

Lecture #05 – Convolutional Neural Networks (CNNs)



Previously on CMP784

- data preprocessing and normalization
- weight initializations
- ways to improve generalization
- babysitting the learning process
- hyperparameter selection
- optimization



Lecture Overview

- convolution layer
- design guidelines for CNNs
- CNN architectures
- transfer learning
- semantic segmentation networks
- object detection networks

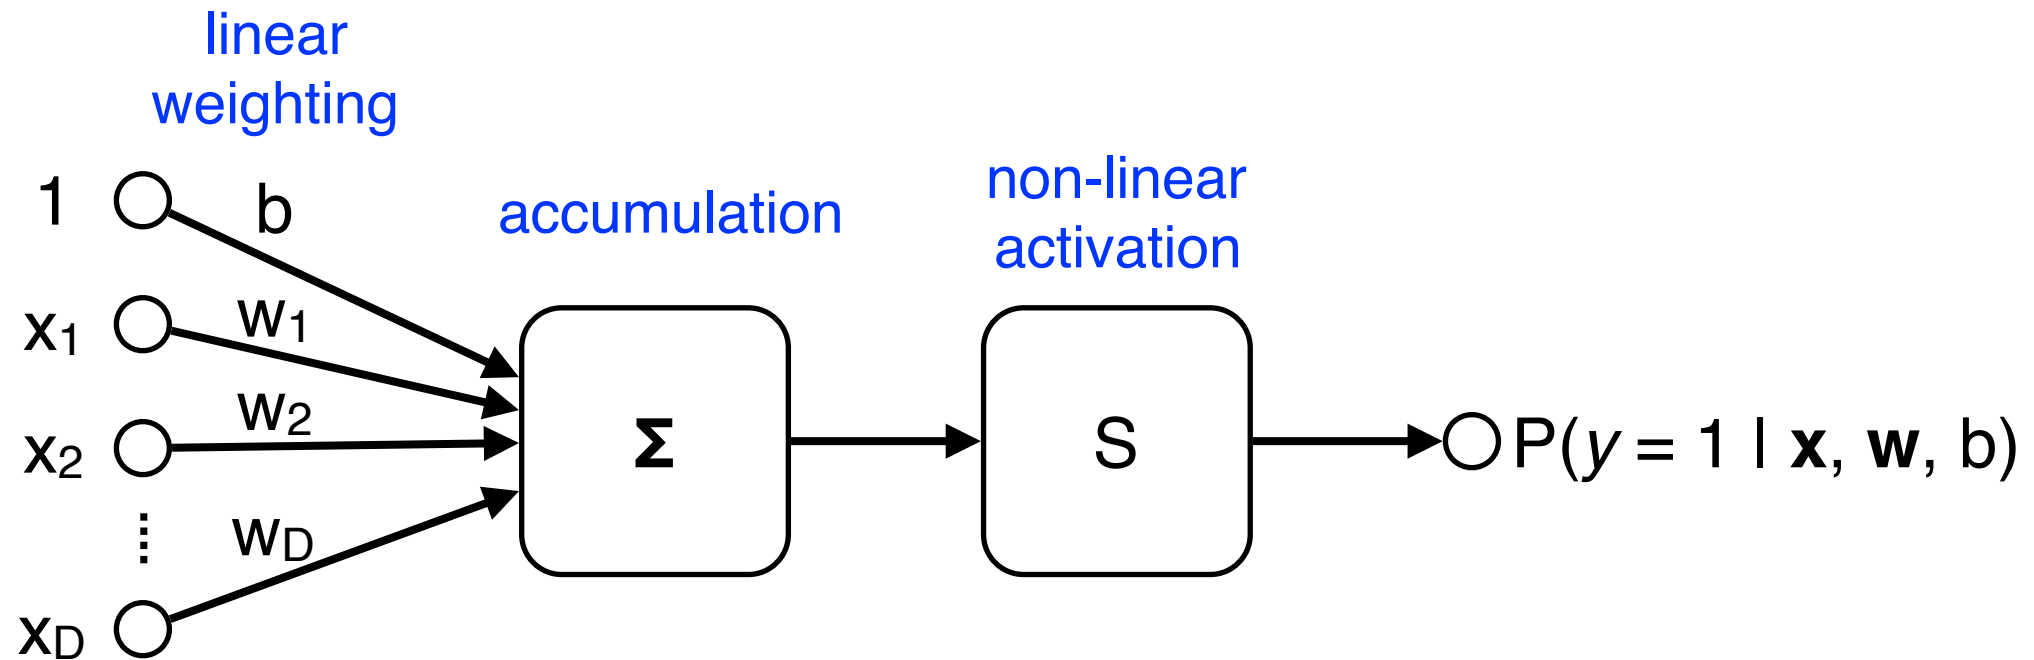
Disclaimer: Much of the material and slides for this lecture were borrowed from

- Andrea Vedaldi's tutorial on Convolutional Networks for Computer Vision Applications
- Kaiming He's ICML 2016 tutorial on Deep Residual Networks: Deep Learning Gets Way Deeper
- Ross Girshick's talk on The Past, Present, and Future of Object Detection
- Fei-Fei Li, Andrej Karpathy and Justin Johnson's CS231n class
- Justin Johnson's EECS 498/598 class

Perceptron





[Rosenblatt 57]

- The goal is estimating the posterior probability of the binary label y of a vector \mathbf{x} :




Discovery of oriented cells in the visual cortex

[Hubel and Wiesel 59]




oriented filter





Convolution



- Convolution = Spatial filtering

$$(a \star b)[i, j] = \sum_{i', j'} a[i', j'] b[i - i', j - j']$$

- Different filters (weights) reveal a different characteristics of the input.




$\star^{1/8}$

0	1	0
1	4	1
0	1	0



Convolution



- Convolution = Spatial filtering

$$(a \star b)[i, j] = \sum_{i', j'} a[i', j'] b[i - i', j - j']$$

- Different filters (weights) reveal a different characteristics of the input.




*

0	-1	0
-1	4	-1
0	-1	0



Convolution



- Convolution = Spatial filtering

$$(a \star b)[i, j] = \sum_{i', j'} a[i', j'] b[i - i', j - j']$$

- Different filters (weights) reveal a different characteristics of the input.



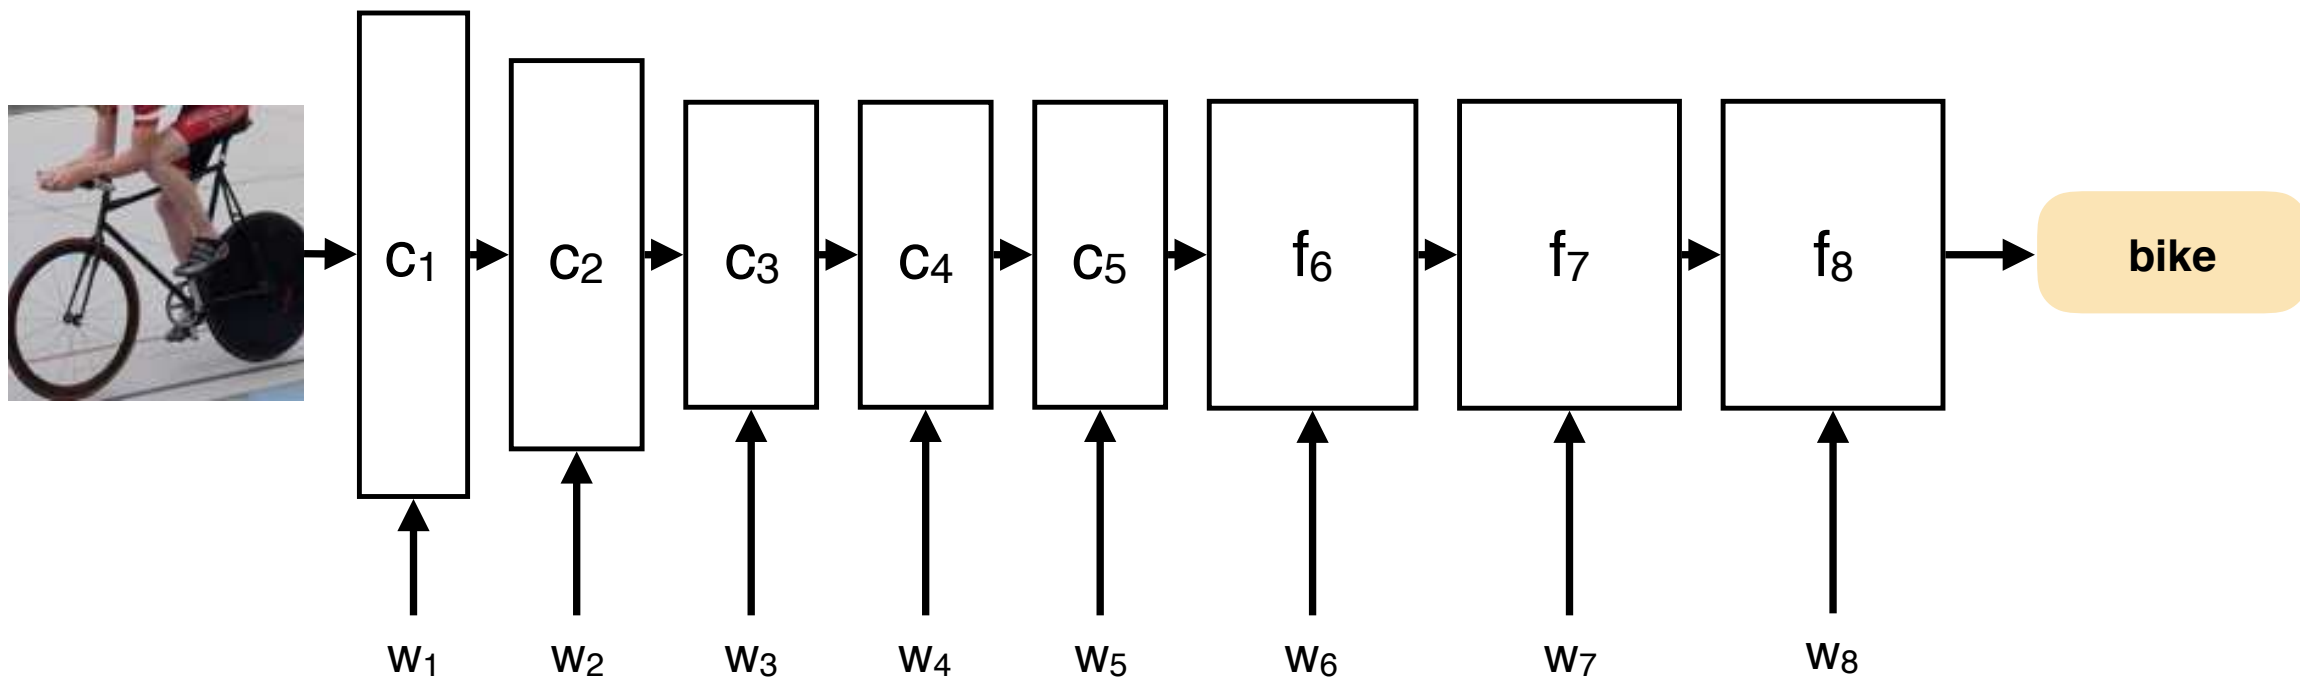
*

1	0	-1
2	0	-2
1	0	-1




Convolutional Neural Networks in a Nutshell

- A neural network model that consists of a sequence of local & translation invariant layers
 - Many identical copies of the same neuron: Weight/parameter sharing
 - Hierarchical feature learning




A bit of history

- Neocognitron model by Fukushima (1980)
- The first convolutional neural network (CNN) model
- so-called “sandwich” architecture
 - simple cells act like filters
 - complex cells perform pooling
- Difficult to train
 - No backpropagation yet



A bit of history


- LeNet-5 model




Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. **Gradient-based learning applied to document recognition.** Proceedings of the IEEE. **86** (11): 2278–2324, 1998.

A bit of history

- AlexNet model



Convolutional Neural Network

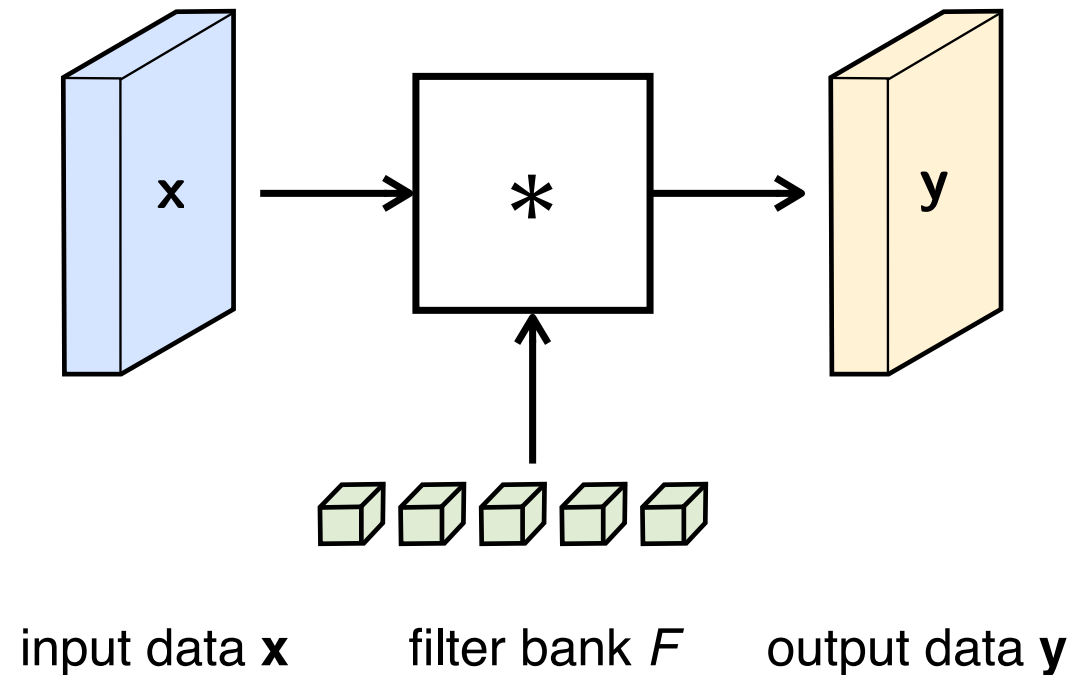


A. Krizhevsky, I. Sutskever, and G. E. Hinton. **Imagenet classification with deep convolutional neural networks**. In NIPS 2012.

Convolutional layer



- Learn a filter bank (a set of filters) once
- Use them over the input data to extract features

$$y = F * x + b$$







Data = 3D Tensors

- There is a vector of feature channels (e.g. RGB) at each spatial location (pixel).




Convolutions with 3D Filters

- Each filter acts on multiple input channels
- **Local**
Filters look locally
- **Translation invariant**
Filters act the same everywhere



Convolutional Layer

32x32x3 input




5x5x3 filter




Convolve the filter with the input
i.e. "slide over the image spatially,
computing dot products"

Convolutional Layer




Convolutional Layer




Convolutional Layer

consider a second, **green** filter




Convolutional Layer

- Multiple filters produce multiple output channels
- For example, if we had 6 5x5 filters, we'll get 6 separate activation maps:



We stack these up to get an output of size 28x28x6.

Spatial Arrangement of Output Volume



- **Depth:** number of filters
- **Stride:** filter step size (when we "slide" it)
- **Padding:** zero-pad the input

Input Volume (+pad 1) (7x7x3)

x[:, :, 0]						
0	0	0	0	0	0	0
0	1	2	0	0	0	0
0	2	2	0	0	0	0
0	2	2	2	1	0	0
0	1	0	0	2	0	0
0	1	1	0	1	0	0
0	0	0	0	0	0	0
x[:, :, 1]						
0	0	0	0	0	0	0
0	1	0	0	1	0	0
0	1	0	0	0	1	0
0	0	0	1	2	0	0
0	0	1	0	1	0	0
0	1	1	1	1	2	0
0	0	0	0	0	0	0
x[:, :, 2]						
0	0	0	0	0	0	0
0	0	0	2	2	2	0
0	1	2	2	2	0	0
0	1	1	1	1	0	0
0	1	2	0	0	0	0
0	2	2	2	1	0	0
0	0	0	0	0	0	0

Filter W0 (3x3x3)

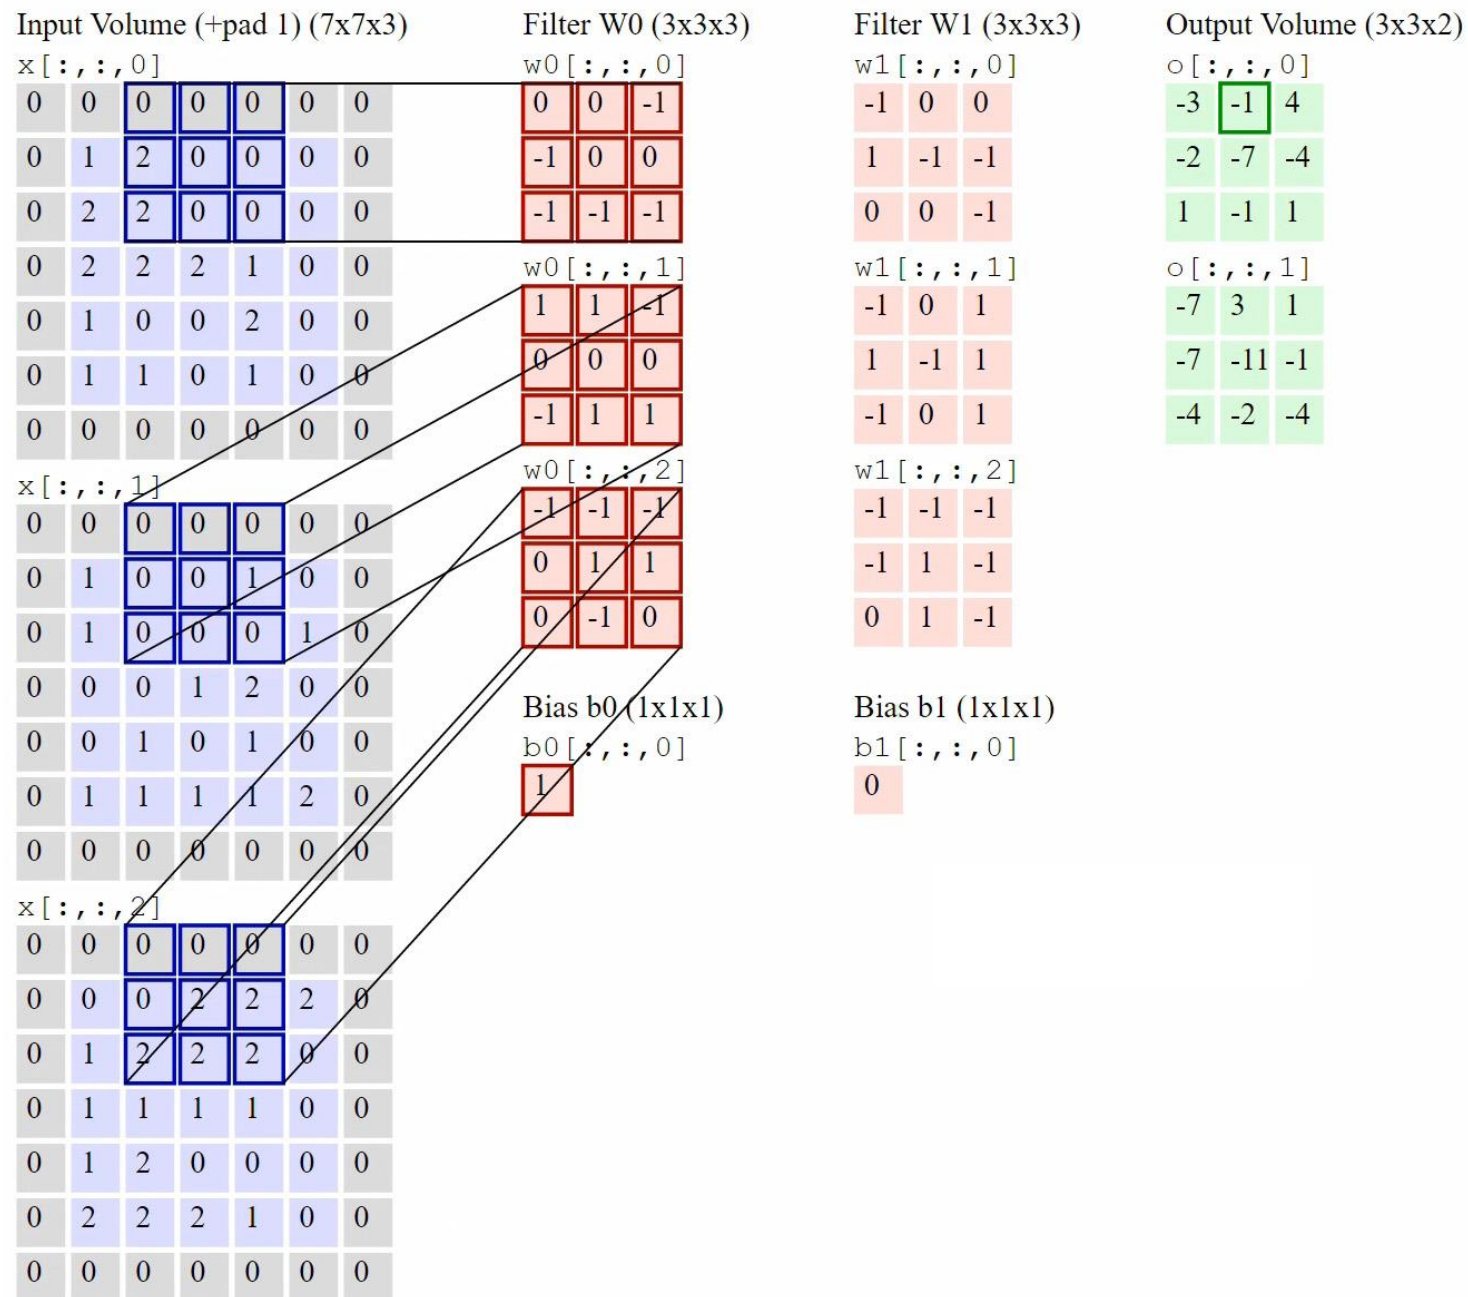
w0[:, :, 0]		
0	0	-1
-1	0	0
-1	-1	-1
w0[:, :, 1]		
1	1	-1
0	0	0
-1	1	1
w0[:, :, 2]		
-1	-1	-1
0	1	1
0	-1	0
Bias b0 (1x1x1)		
b0[:, :, 0]		
1		


Filter W1 (3x3x3)


w1[:, :, 0]		
-1	0	0
1	-1	-1
0	0	-1
w1[:, :, 1]		
-1	0	1
1	-1	1
-1	0	1
w1[:, :, 2]		
-1	-1	-1
-1	1	-1
0	1	-1
Bias b1 (1x1x1)		
b1[:, :, 0]		
0		

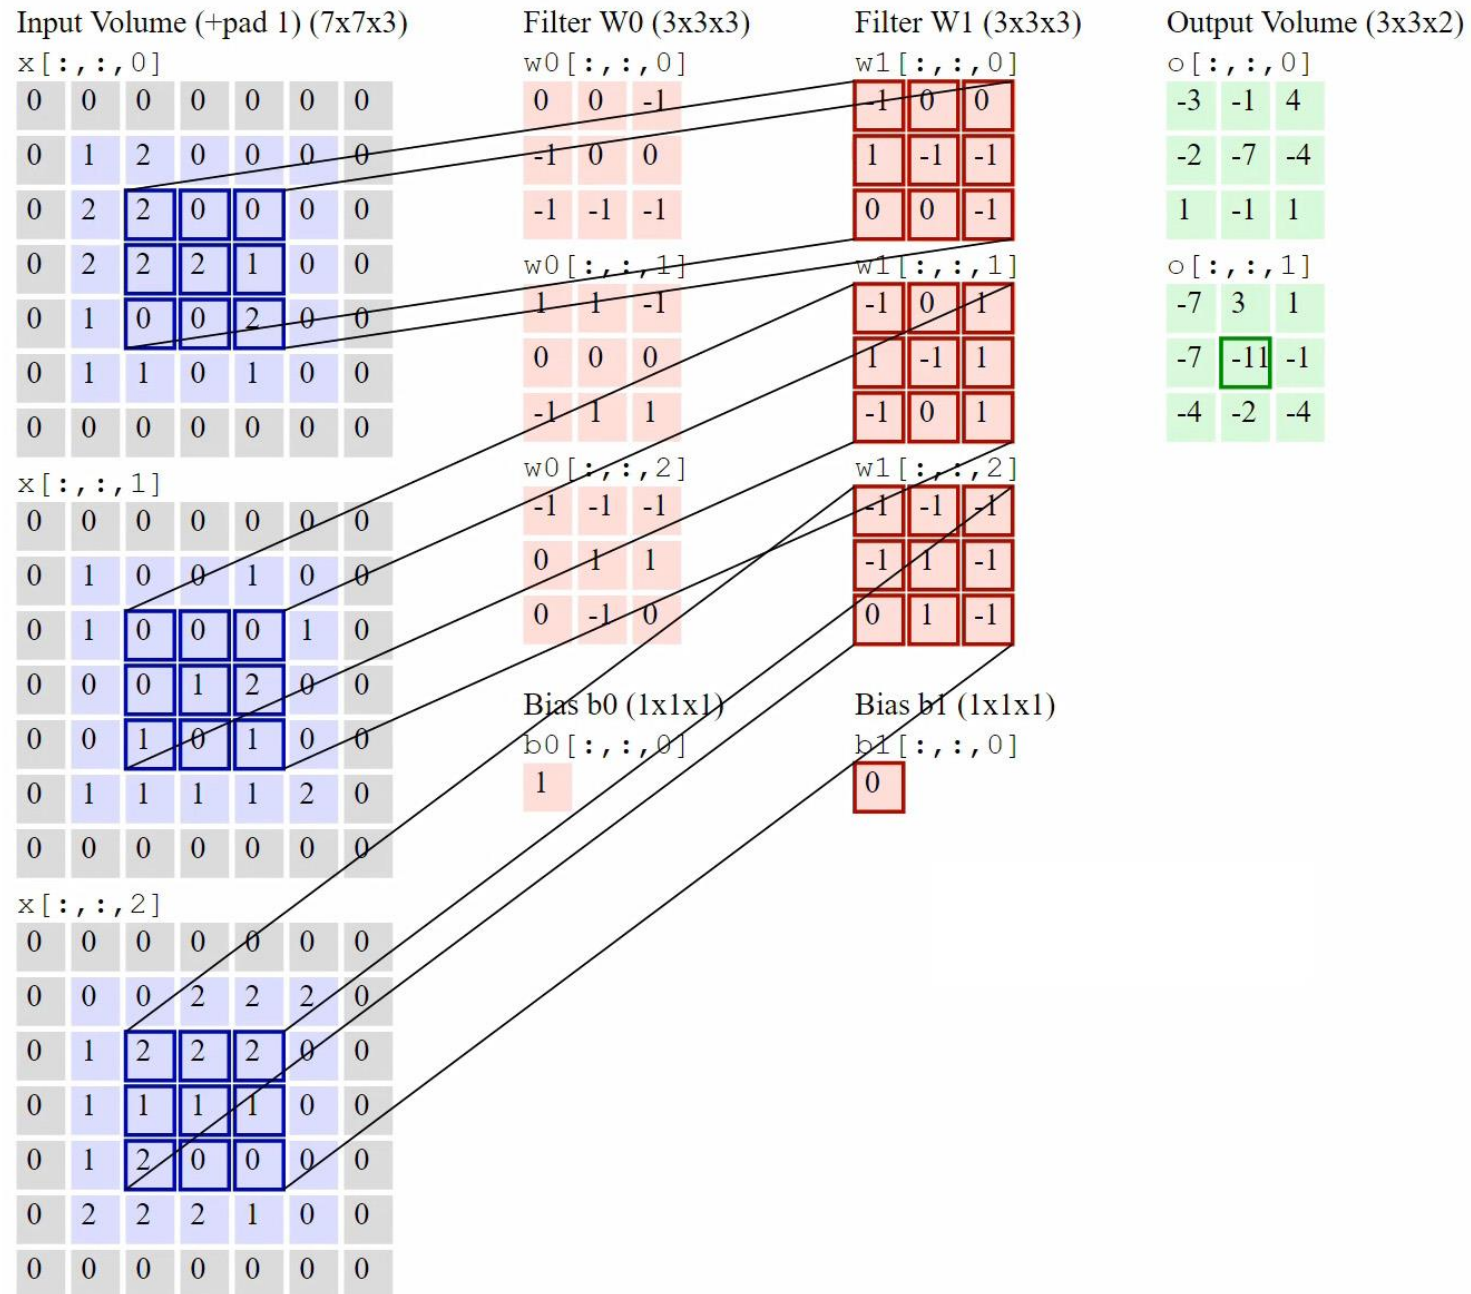
Output Volume (3x3x2)

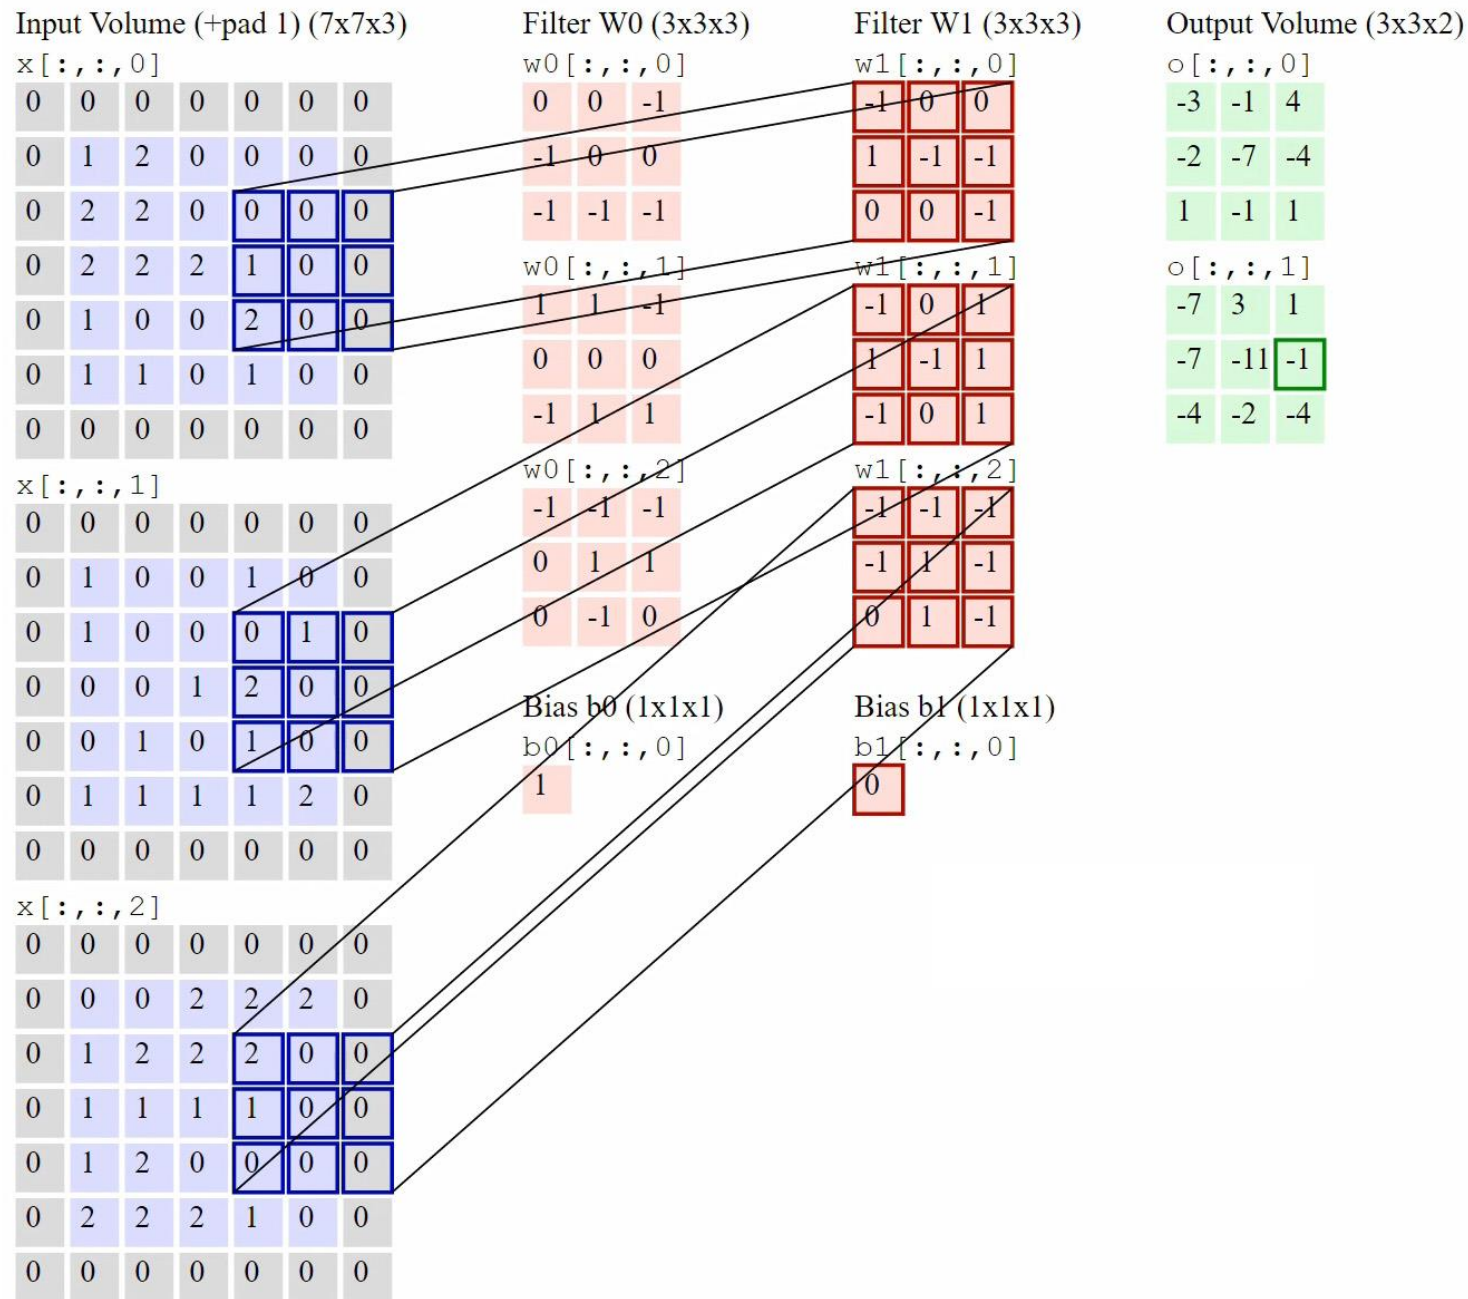
o[:, :, 0]		
-3	-1	4
-2	-7	-4
1	-1	1
o[:, :, 1]		
-7	3	1
-7	-11	-1
-4	-2	-4

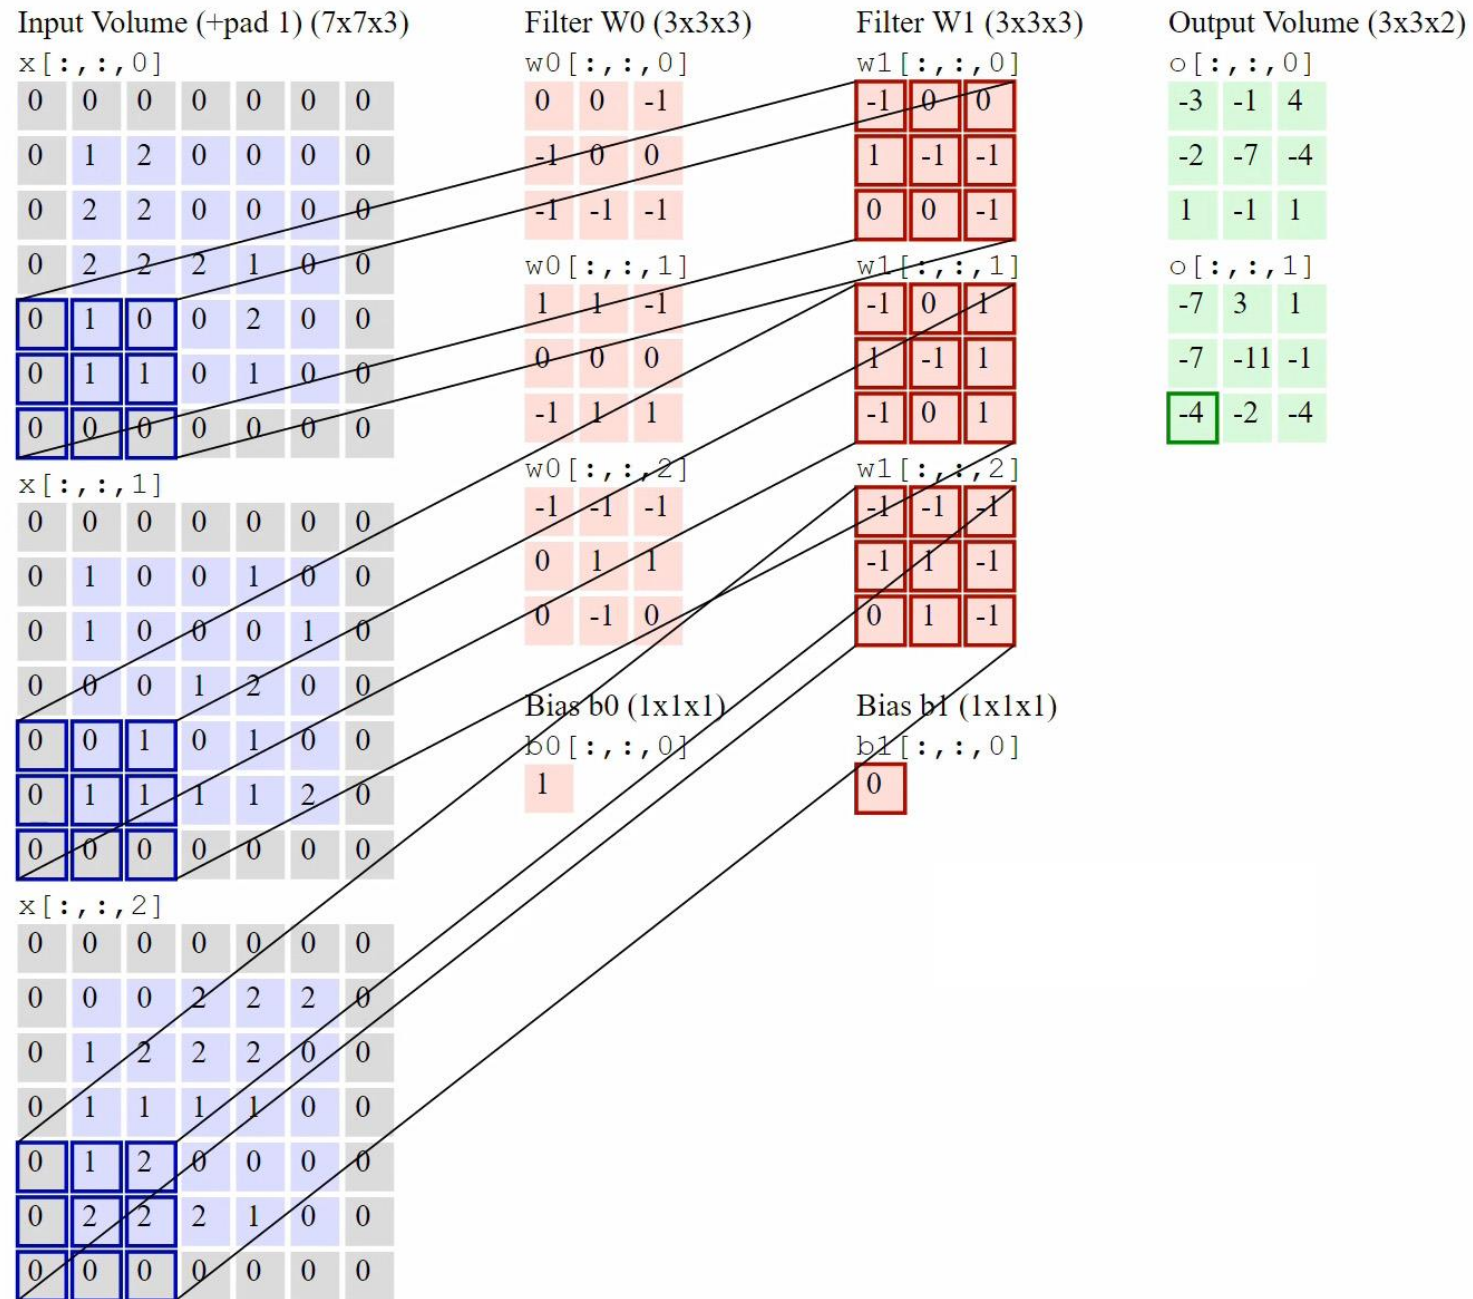


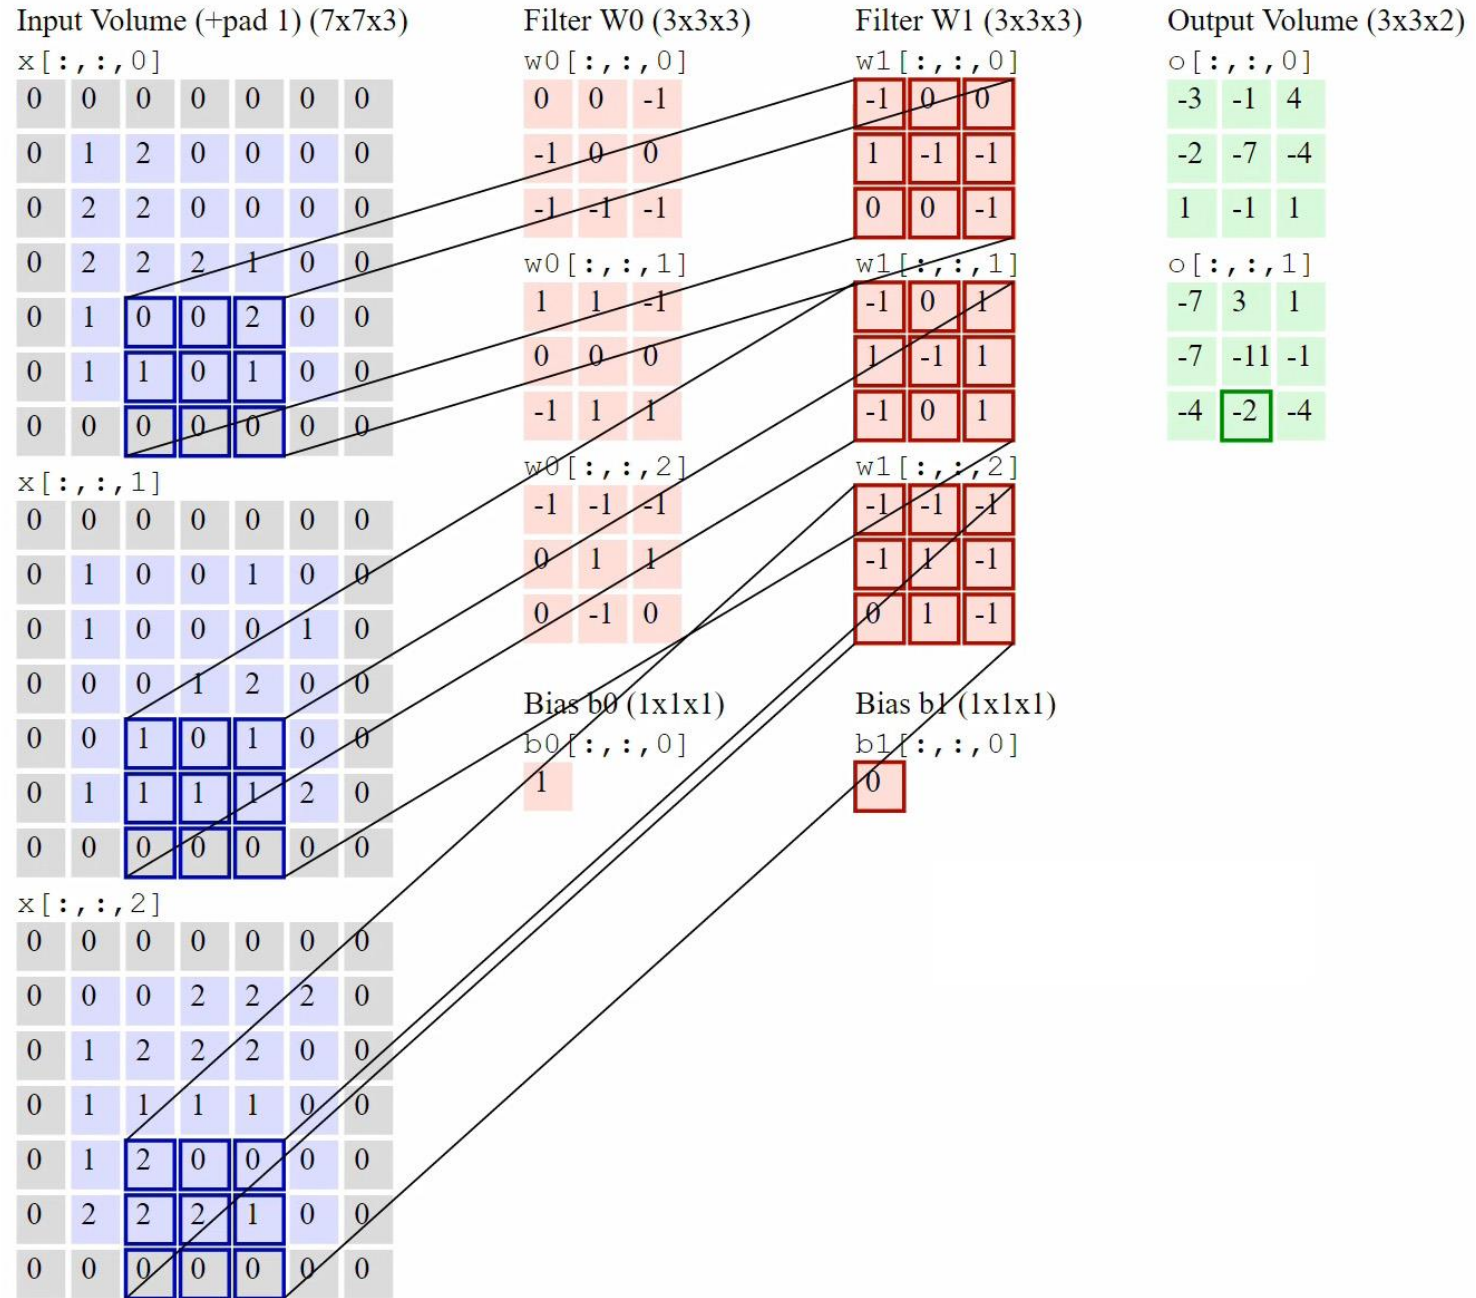
















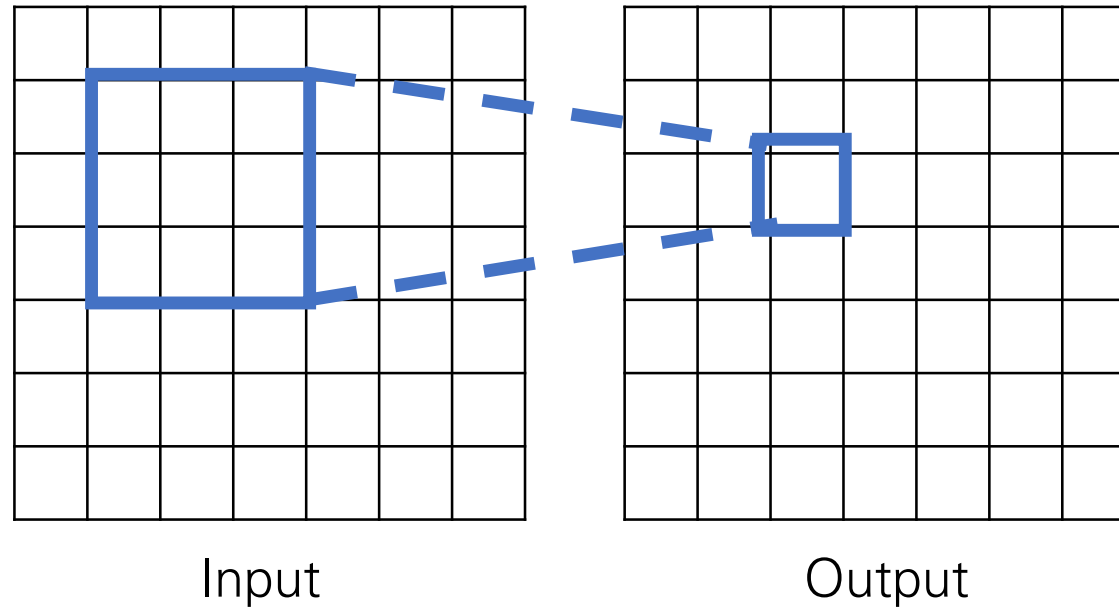
Convolutional layers

- Local receptive field
- Each column of hidden units looks at a different input patch




Receptive Fields

- For convolution with kernel size K , each element in the output depends on a $K \times K$ **receptive field** in the input




Receptive Fields

- Each successive convolution adds $K - 1$ to the receptive field size. With L layers the receptive field size is $1 + L * (K - 1)$




Problem: For large images we need many layers for each output to "see" the whole image

1x1 Convolution




1x1 CONV
with 32 filters




(each filter has size 64x1x1,
and performs a 64-
dimensional dot product)

Stacking 1x1 conv layers
gives MLP operating on
each input position




Other types of convolution

So far: 2D Convolution




1D Convolution

Input: $C_{in} \times W$
Weights: $C_{out} \times C_{in} \times K$




3D Convolution

Input: $C_{in} \times H \times W \times D$
Weights: $C_{out} \times C_{in} \times K \times K \times K$




C_{in} -dim vector at each point in the volume

Convolutional layers

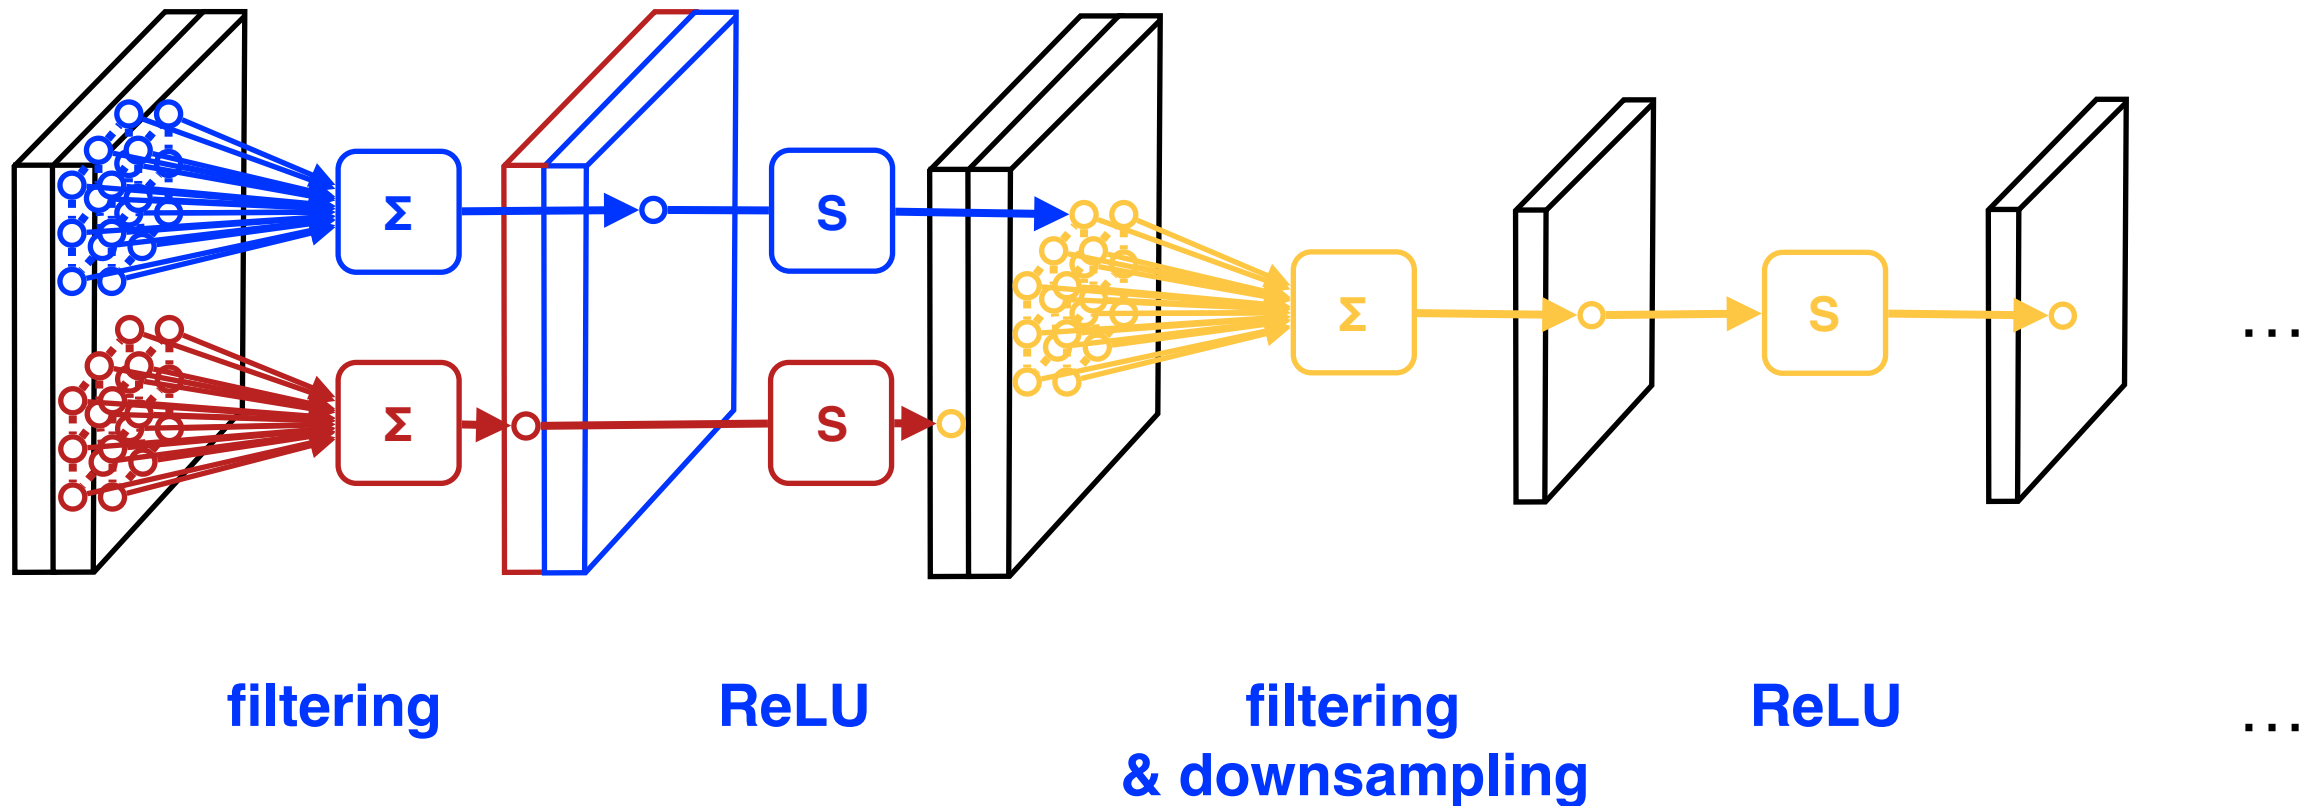


Repeat linear / non-linear operators




Linear/Non-linear Chains

- The basic blueprint of most architectures
- Stack multiple layers of convolutions



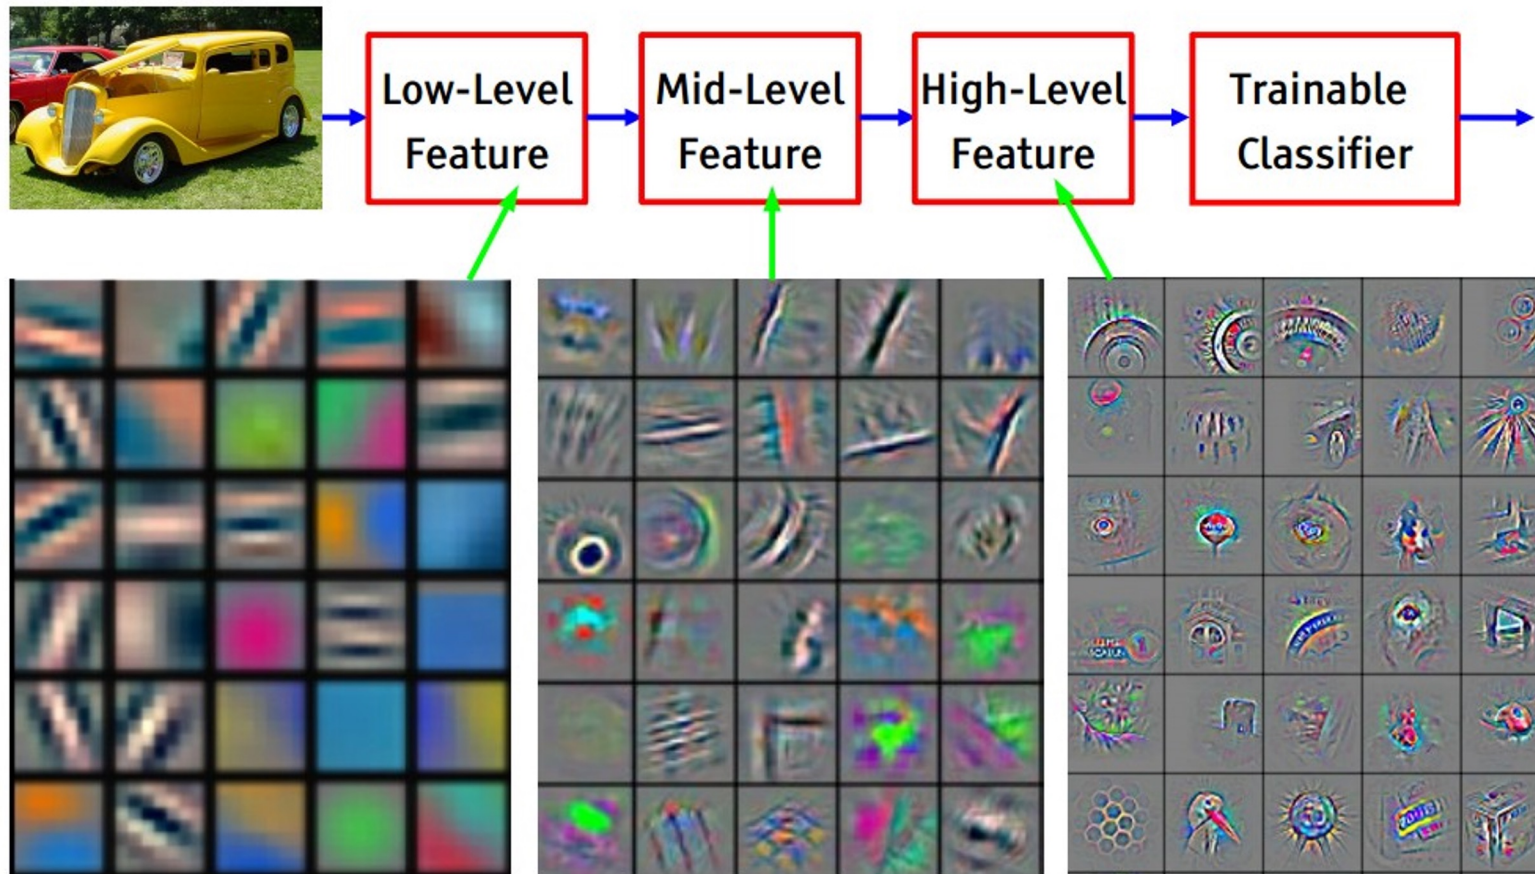
Feature Learning

- Hierarchical layer structure allows to learn hierarchical filters (features).



Feature Learning

- Hierarchical layer structure allows to learn hierarchical filters (features).



Feature visualization of convolutional net trained on ImageNet from [Zeiler & Fergus 2013]

Pooling layer


- makes the representations smaller and more manageable
- operates over each activation map independently:
- Max pooling, average pooling, etc.

Single depth slice

1	1	2	4
5	6	7	8
3	2	1	0
1	2	3	4

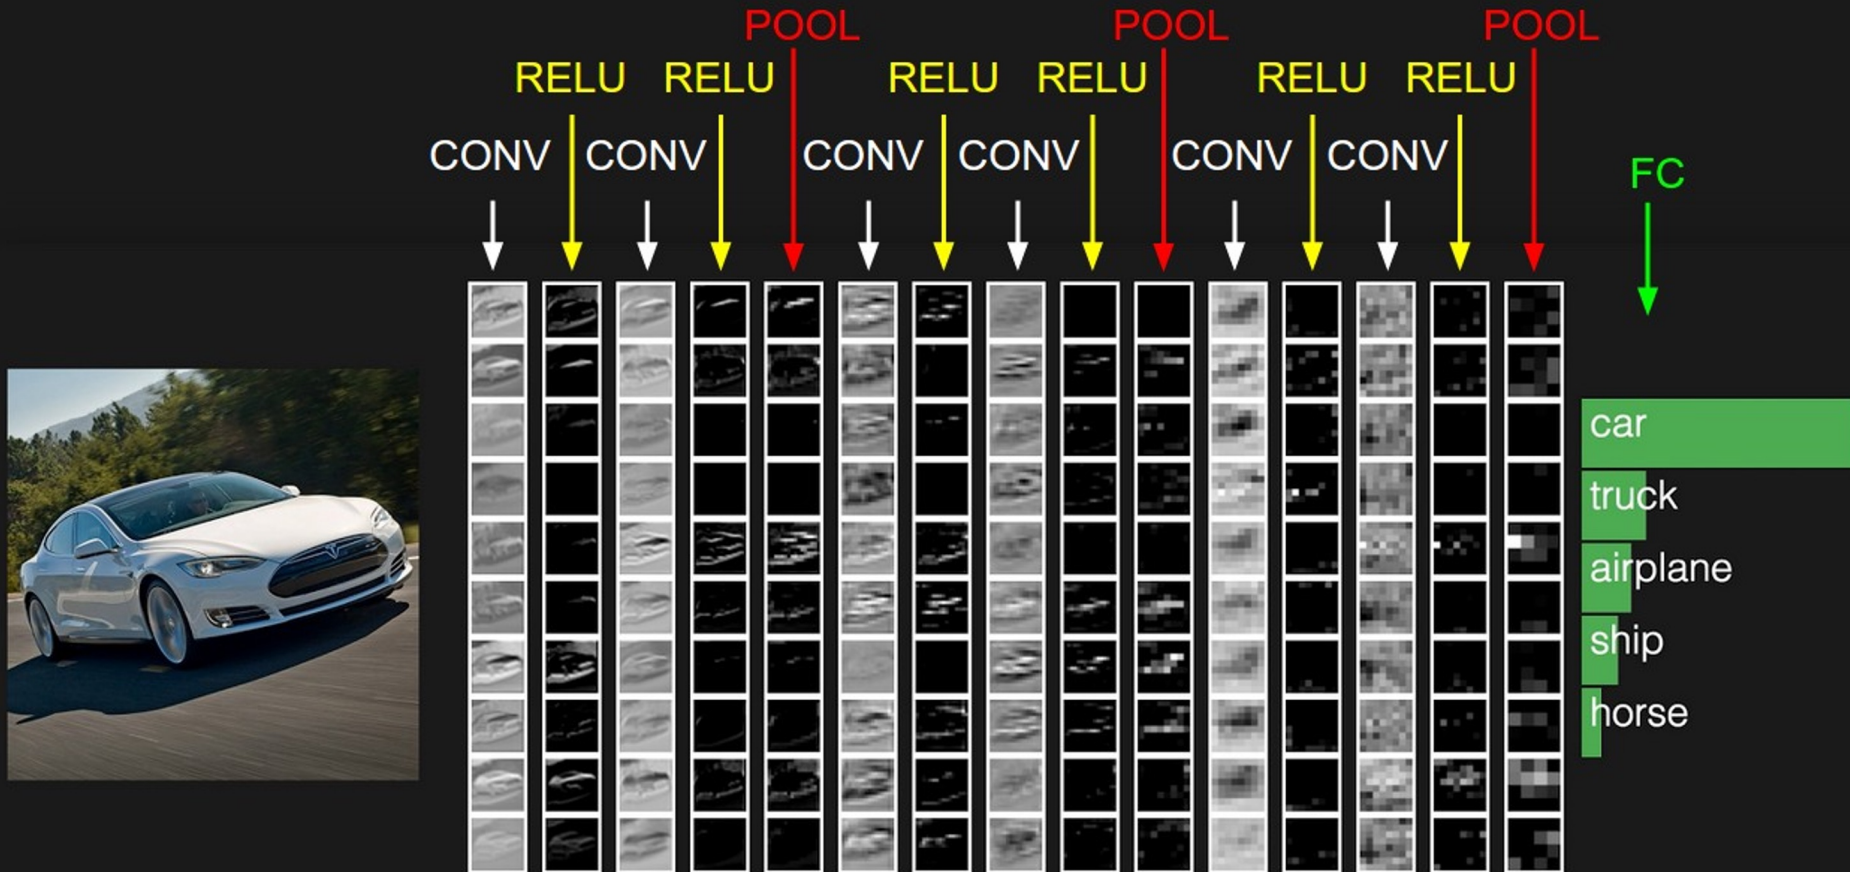
max pool with 2x2 filters and stride 2

6	8
3	4




Fully connected layer

- contains neurons that connect to the entire input volume, as in ordinary Neural Networks



Design Guidelines

Design Guidelines



Guideline 1: Avoid tight bottlenecks

- **From bottom to top**

- The spatial resolution $H \times W$ decreases
- The number of channels C increases


- **Guideline**

- Avoid tight information bottleneck
- Decrease the data volume $H \times W \times C$ slowly

K. Simonyan and A. Zisserman. **Very deep convolutional networks for large-scale image recognition**. In ICLR 2015.

C. Szegedy, V. Vanhoucke, S. Ioffe, and J. Shlens. **Rethinking the inception architecture for computer vision**. In CVPR 2016.

Receptive Field



Must be large enough

- **Receptive field of a neuron**

- The image region influencing a neuron
- Anything happening outside is invisible to the neuron

- **Importance**

- Large image structures cannot be detected by neurons with small receptive fields


- **Enlarging the receptive field**

- Large filters
- Chains of small filters


Design Guidelines

Guideline 2: Prefer small filter chains


One big filter bank



5 × 5 filters
+ ReLU



Two smaller filter banks



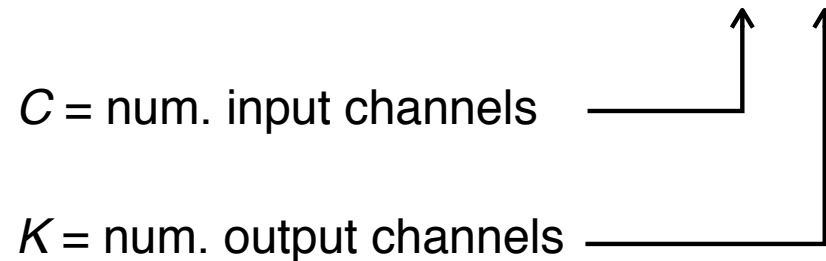

3 × 3 filters + ReLU 3 × 3 filters + ReLU

- **Remark:** 101 ResNet layers same size/speed as 16 VGG-VD layers
- **Reason:** Far fewer feature channels (quadratic speed/space gain)
- **Moral:** Optimize your architecture

Design Guidelines

Guideline 3:

Keep
the number
of channels
at bay



Num. of operations

$$\frac{H \times H_f}{\text{stride}} \times \frac{W \times W_f}{\text{stride}} \times \underline{C \times K}$$

Num. of parameters

$$H_f \times W_f \times \underline{C \times K}$$



complexity $\propto \underline{C \times K}$

Design Guidelines


Guideline 4:

Less computations with filter groups

M filters




G groups of M/G filters



Did we see this before?

$$\text{complexity} \propto (C \times K) / G$$


AlexNet




Design Guidelines

Guideline 4:

Less computations with filter groups



complexity: $C \times K$




complexity: $C \times K / G$

Groups = filters, seen as a matrix, have a “block” structure

Design Guidelines


Guideline 5:

Low-rank decompositions




filter bank
 $3 \times 3 \times C \times K$

decompose spatially




vertical
 $1 \times 3 \times C \times K$

*



horizontal
 $3 \times 1 \times K \times K$


*



vertical
 $1 \times 3 \times K \times K$




decompose channels



groups
 $3 \times 3 \times C/G \times K/G$

*




“network in network”
 $1 \times 1 \times K \times K$


Make sure to mix the information

Design Guidelines

Guideline 6: Dilated Convolutions



49 coefficients
18 degrees of freedom




Exponential expansion of the receptive field without loss of resolution

Convolutional Neural Network Demo

- ConvNetJS demo: training on CIFAR-10
- <http://cs.stanford.edu/people/karpathy/convnetjs/demo/cifar10.html>


CNN Architectures

ImageNet Classification Challenge



AlexNet

[Krizhevsky et al. 2012]



Full (simplified) AlexNet architecture:

[227x227x3] INPUT

[55x55x96] **CONV1**: 96 11x11 filters at stride 4, pad 0

[27x27x96] **MAX POOL1**: 3x3 filters at stride 2

[27x27x96] **NORM1**: Normalization layer

[27x27x256] **CONV2**: 256 5x5 filters at stride 1, pad 2

[13x13x256] **MAX POOL2**: 3x3 filters at stride 2

[13x13x256] **NORM2**: Normalization layer

[13x13x384] **CONV3**: 384 3x3 filters at stride 1, pad 1

[13x13x384] **CONV4**: 384 3x3 filters at stride 1, pad 1

[13x13x256] **CONV5**: 256 3x3 filters at stride 1, pad 1

[6x6x256] **MAX POOL3**: 3x3 filters at stride 2

[4096] **FC6**: 4096 neurons


[4096] **FC7**: 4096 neurons

[1000] **FC8**: 1000 neurons (class scores)

Details/Retrospectives:


- first use of ReLU
- used Norm layers (not common anymore)
- heavy data augmentation
- dropout 0.5
- batch size 128
- SGD Momentum 0.9
- Learning rate 1e-2, reduced by 10 manually when val accuracy plateaus
- L2 weight decay 5e-4
- 7 CNN ensemble: 18.2% -> 15.4%

AlexNet




Most of the **memory usage** is in the early convolution layers

Memory (KB)




Nearly all **parameters** are in the fully-connected layers

Params (K)




Most **floating-point ops** occur in the convolution layers

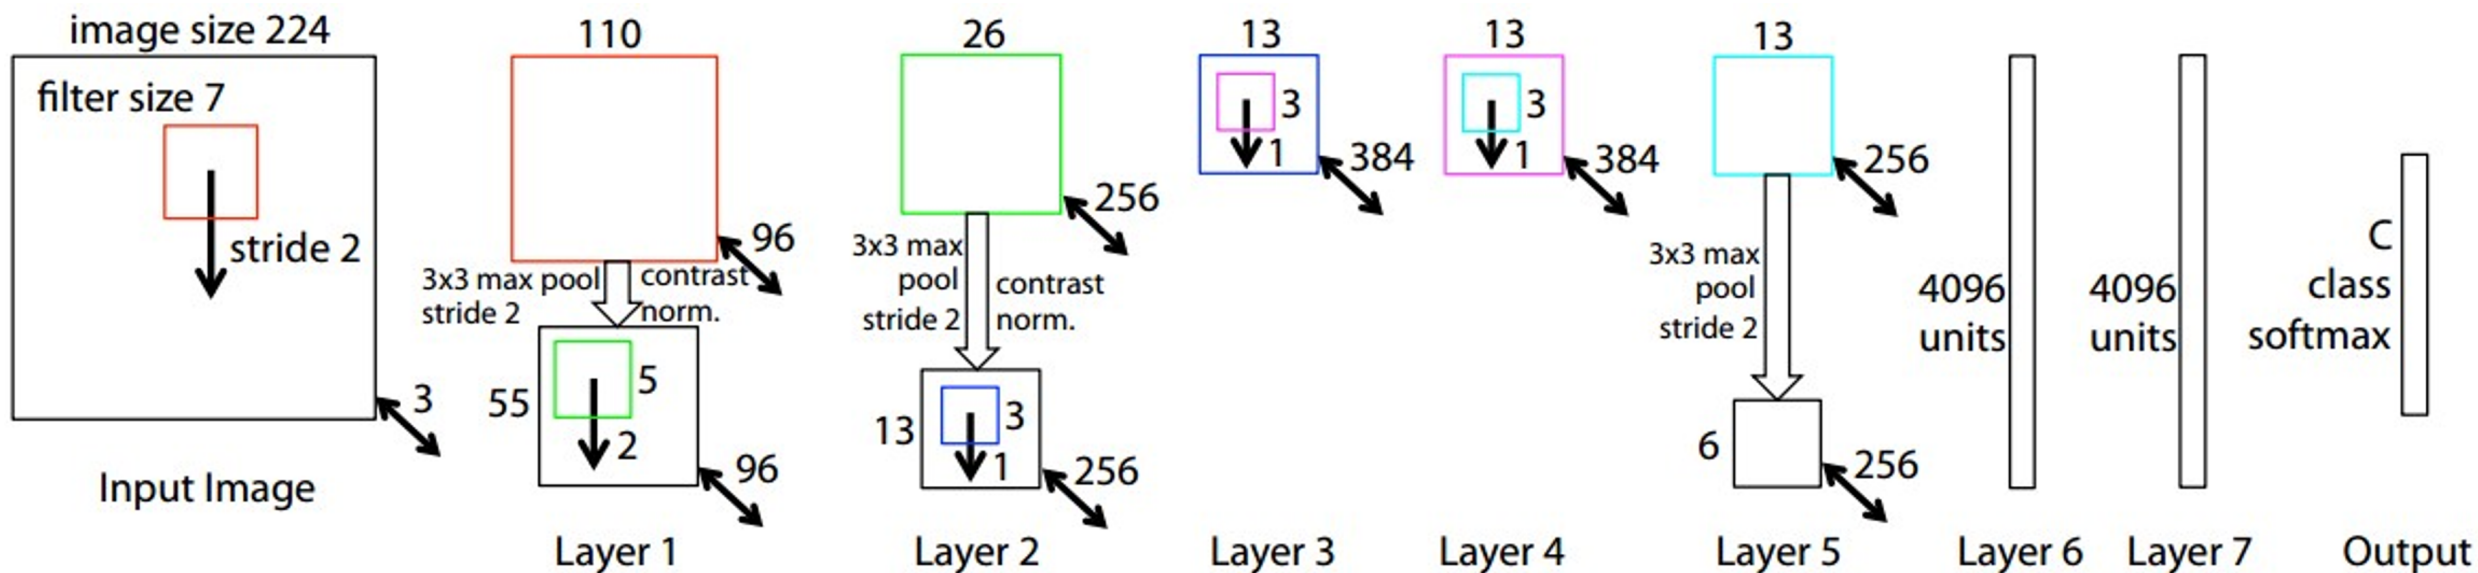
MFLOP



ImageNet Classification Challenge



ZFNet: A Bigger AlexNet




AlexNet but:

CONV1: change from (11x11 stride 4) to (7x7 stride 2)

CONV3,4,5: instead of 384, 384, 256 filters use 512, 1024, 512

More trial and error

ImageNet Classification Challenge



INPUT: [224x224x3] memory: $224*224*3=150K$ params: 0 (not counting biases)

CONV3-64: [224x224x64] memory: $224*224*64=3.2M$ params: $(3*3*3)*64 = 1,728$

CONV3-64: [224x224x64] memory: $224*224*64=3.2M$ params: $(3*3*64)*64 = 36,864$

POOL2: [112x112x64] memory: $112*112*64=800K$ params: 0

CONV3-128: [112x112x128] memory: $112*112*128=1.6M$ params: $(3*3*64)*128 = 73,728$

CONV3-128: [112x112x128] memory: $112*112*128=1.6M$ params: $(3*3*128)*128 = 147,456$

POOL2: [56x56x128] memory: $56*56*128=400K$ params: 0

CONV3-256: [56x56x256] memory: $56*56*256=800K$ params: $(3*3*128)*256 = 294,912$

CONV3-256: [56x56x256] memory: $56*56*256=800K$ params: $(3*3*256)*256 = 589,824$

CONV3-256: [56x56x256] memory: $56*56*256=800K$ params: $(3*3*256)*256 = 589,824$

POOL2: [28x28x256] memory: $28*28*256=200K$ params: 0

CONV3-512: [28x28x512] memory: $28*28*512=400K$ params: $(3*3*256)*512 = 1,179,648$

CONV3-512: [28x28x512] memory: $28*28*512=400K$ params: $(3*3*512)*512 = 2,359,296$

CONV3-512: [28x28x512] memory: $28*28*512=400K$ params: $(3*3*512)*512 = 2,359,296$

POOL2: [14x14x512] memory: $14*14*512=100K$ params: 0

CONV3-512: [14x14x512] memory: $14*14*512=100K$ params: $(3*3*512)*512 = 2,359,296$

CONV3-512: [14x14x512] memory: $14*14*512=100K$ params: $(3*3*512)*512 = 2,359,296$

CONV3-512: [14x14x512] memory: $14*14*512=100K$ params: $(3*3*512)*512 = 2,359,296$

POOL2: [7x7x512] memory: $7*7*512=25K$ params: 0

FC: [1x1x4096] memory: 4096 params: $7*7*512*4096 = 102,760,448$

FC: [1x1x4096] memory: 4096 params: $4096*4096 = 16,777,216$

FC: [1x1x1000] memory: 1000 params: $4096*1000 = 4,096,000$

TOTAL memory: 24M * 4 bytes \approx 93MB / image

(only forward! \sim *2 for bwd)

TOTAL params: 138M parameters

ConvNet Configuration			
B	C	D	
13 weight layers	16 weight layers	16 weight layers	19
put (224 x 224 RGB image)			
conv3-64	conv3-64	conv3-64	cc
conv3-64	conv3-64	conv3-64	cc
maxpool			
conv3-128	conv3-128	conv3-128	co
conv3-128	conv3-128	conv3-128	co
maxpool			
conv3-256	conv3-256	conv3-256	co
conv3-256	conv3-256	conv3-256	co
	conv1-256	conv3-256	co
			co
maxpool			
conv3-512	conv3-512	conv3-512	co
conv3-512	conv3-512	conv3-512	co
	conv1-512	conv3-512	co
			co
maxpool			
conv3-512	conv3-512	conv3-512	co
conv3-512	conv3-512	conv3-512	co
	conv1-512	conv3-512	co
			co
maxpool			
FC-4096			
FC-4096			
FC-1000			
soft-max			

VGG-16 Net

VGG: Deeper Networks, Regular Design

VGG Design rules:

All conv are 3x3 stride 1 pad 1

All max pool are 2x2 stride 2

After pool, double #channels

Network has 5 convolutional stages:

Stage 1: conv-conv-pool

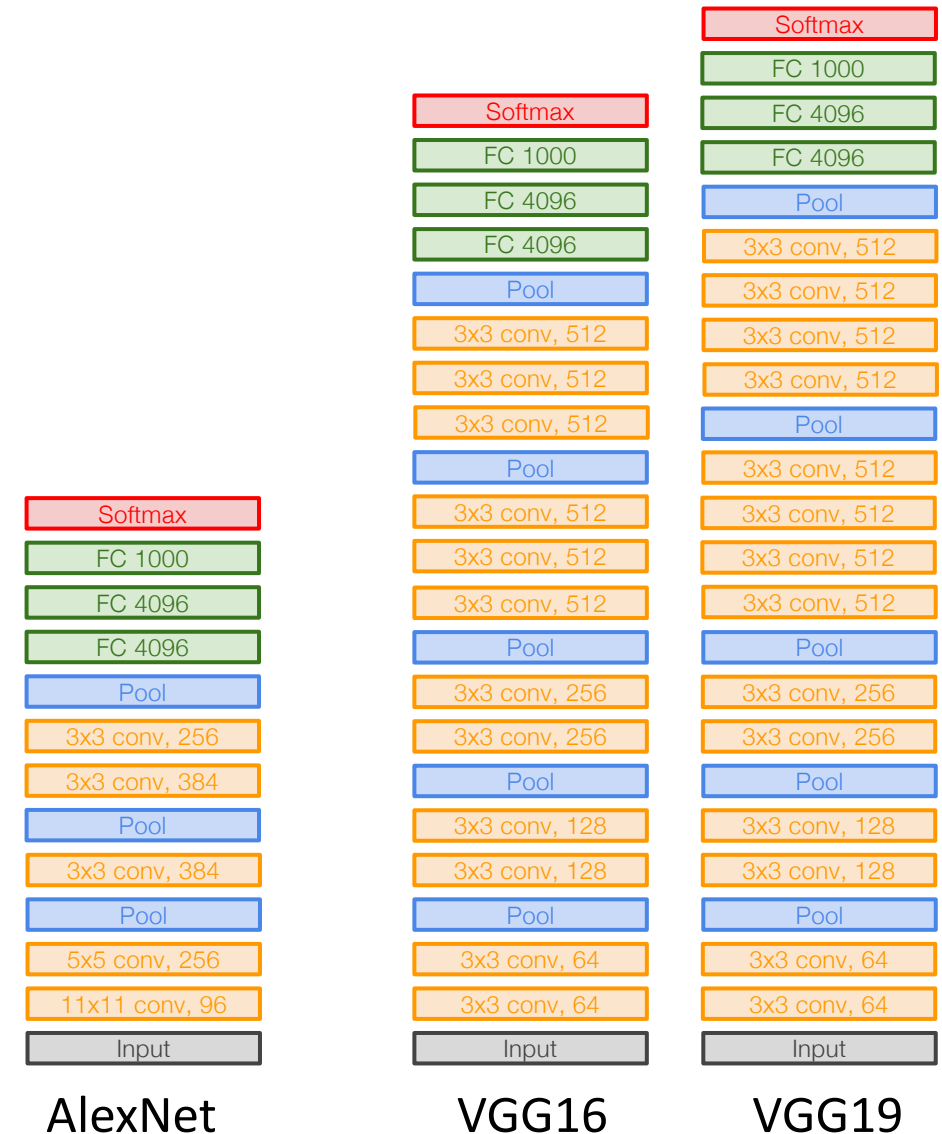
Stage 2: conv-conv-pool

Stage 3: conv-conv-pool

Stage 4: conv-conv-conv-[conv]-pool

Stage 5: conv-conv-conv-[conv]-pool

(VGG-19 has 4 conv in stages 4 and 5)



VGG: Deeper Networks, Regular Design

VGG Design rules:

All conv are 3x3 stride 1 pad 1

All max pool are 2x2 stride 2

After pool, double #channels

Network has 5 convolutional stages:

Stage 1: conv-conv-pool

Stage 2: conv-conv-pool

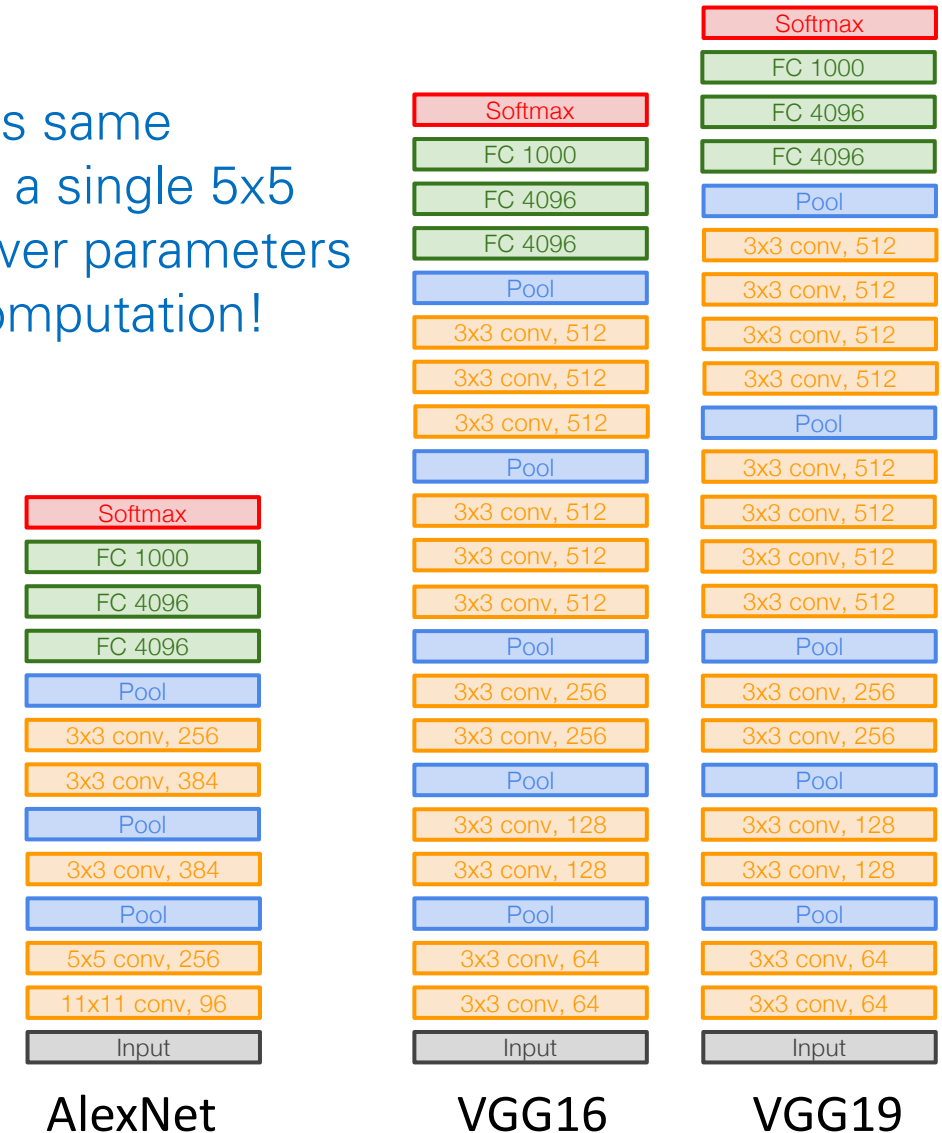
Stage 3: conv-conv-pool

Stage 4: conv-conv-conv-[conv]-pool

Stage 5: conv-conv-conv-[conv]-pool

(VGG-19 has 4 conv in stages 4 and 5)

Two 3x3 conv has same receptive field as a single 5x5 conv, but has fewer parameters and takes less computation!



VGG: Deeper Networks, Regular Design

VGG Design rules:

All conv are 3x3 stride 1 pad 1

All max pool are 2x2 stride 2

After pool, double #channels

Network has 5 convolutional stages:

Stage 1: conv-conv-pool

Stage 2: conv-conv-pool

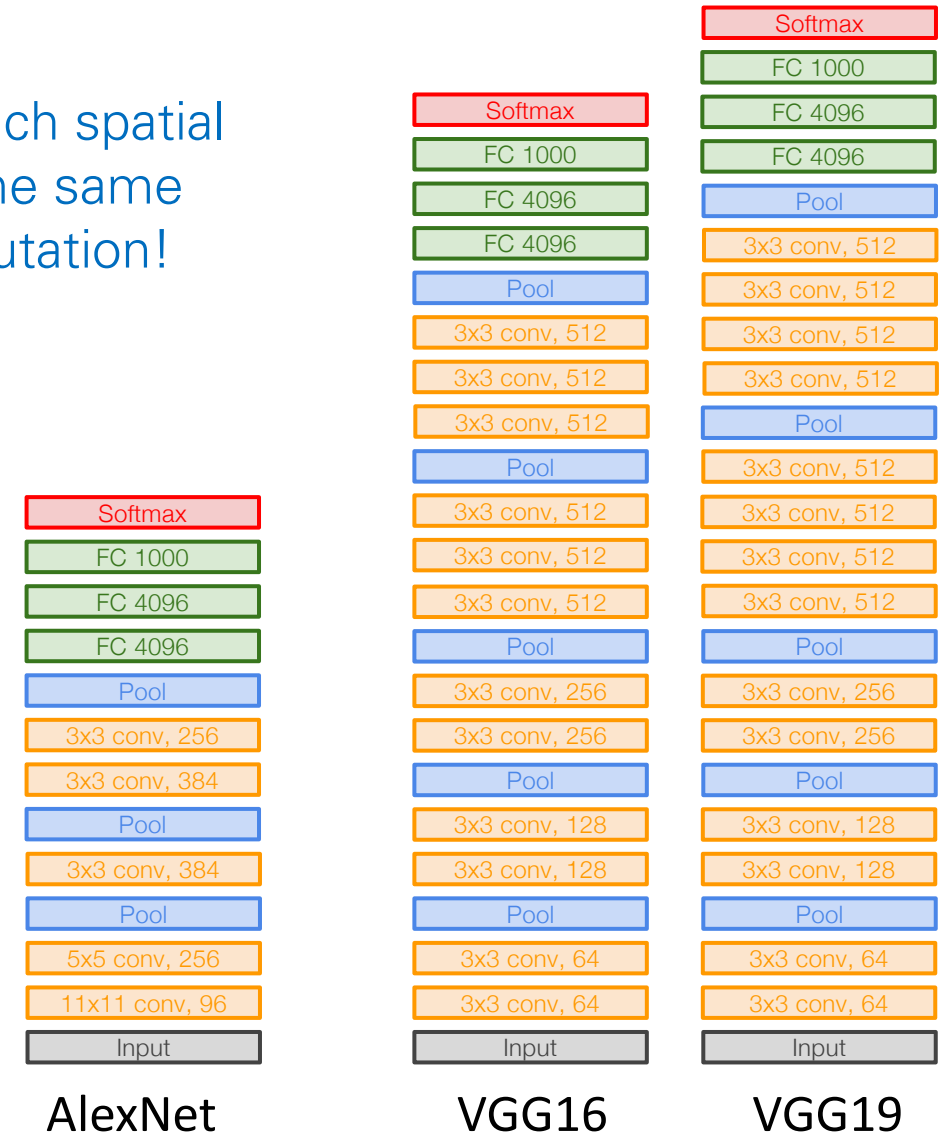
Stage 3: conv-conv-pool

Stage 4: conv-conv-conv-[conv]-pool


Stage 5: conv-conv-conv-[conv]-pool

(VGG-19 has 4 conv in stages 4 and 5)

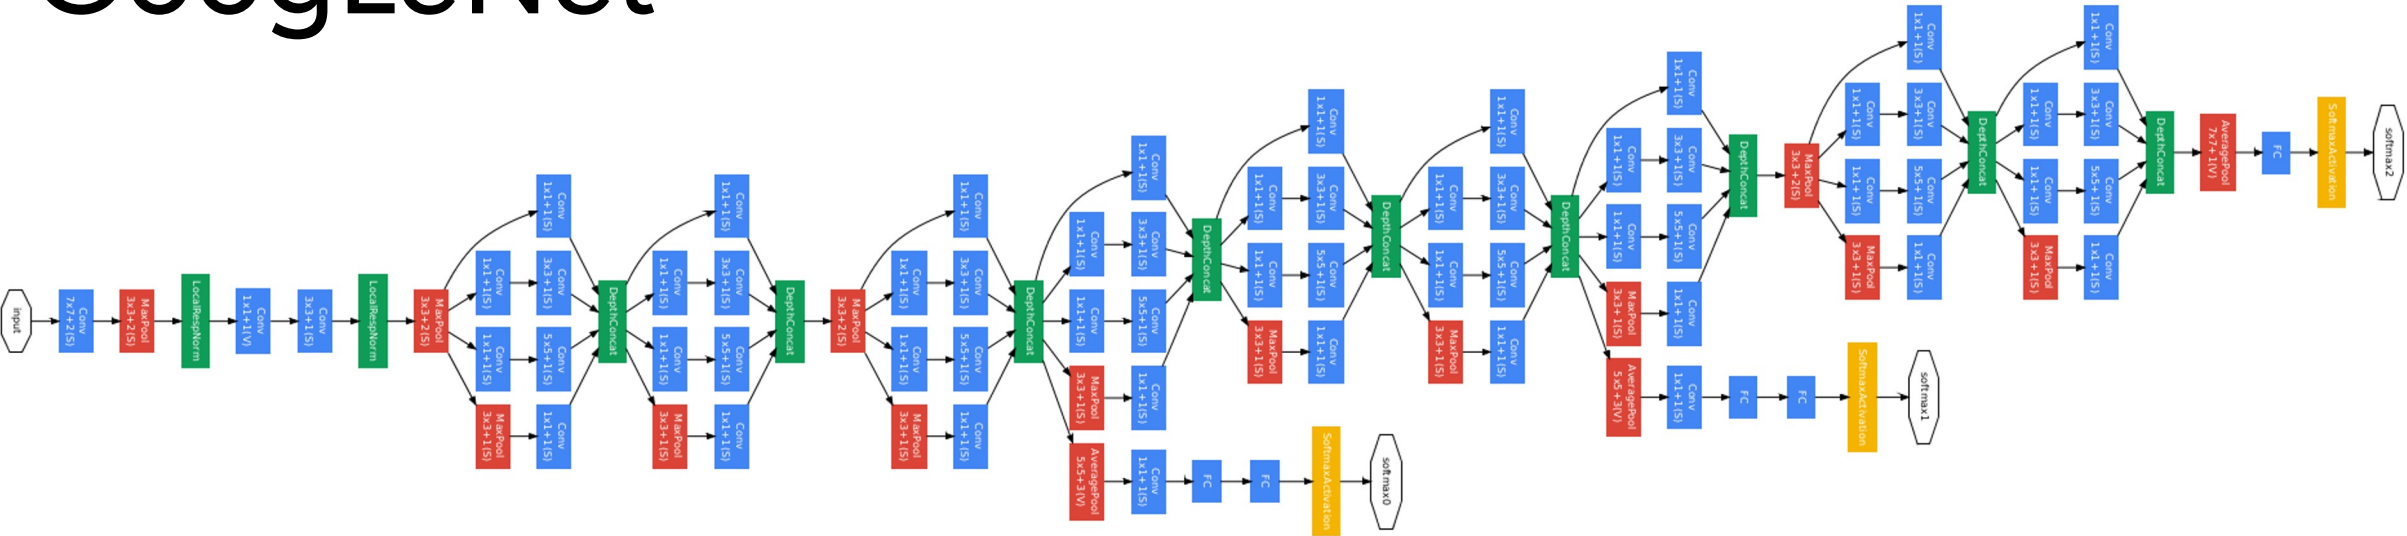
Conv layers at each spatial resolution take the same amount of computation!



ImageNet Classification Challenge

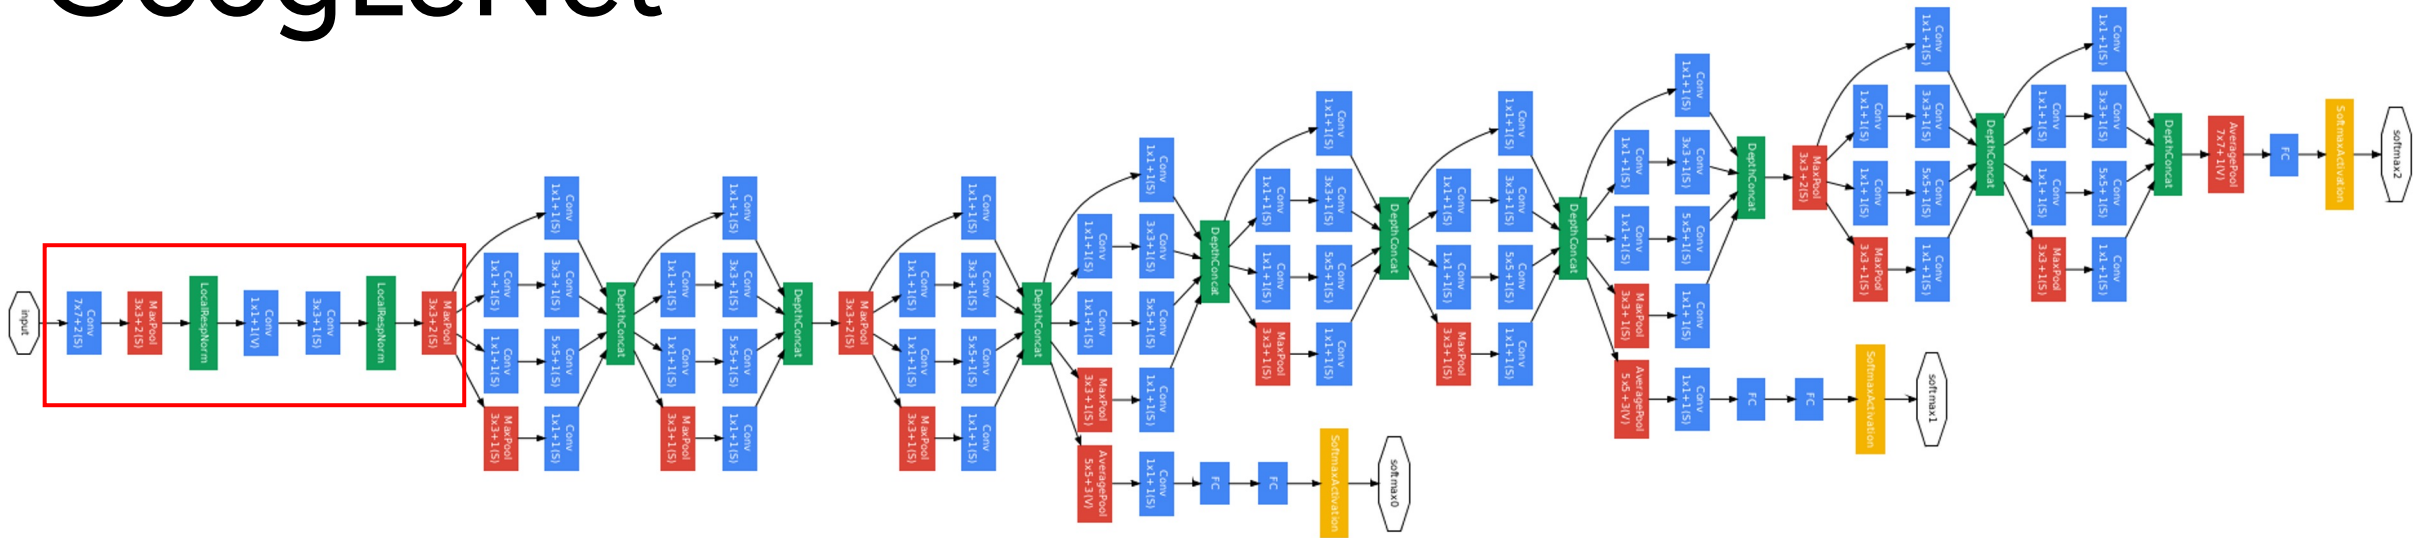


GoogLeNet



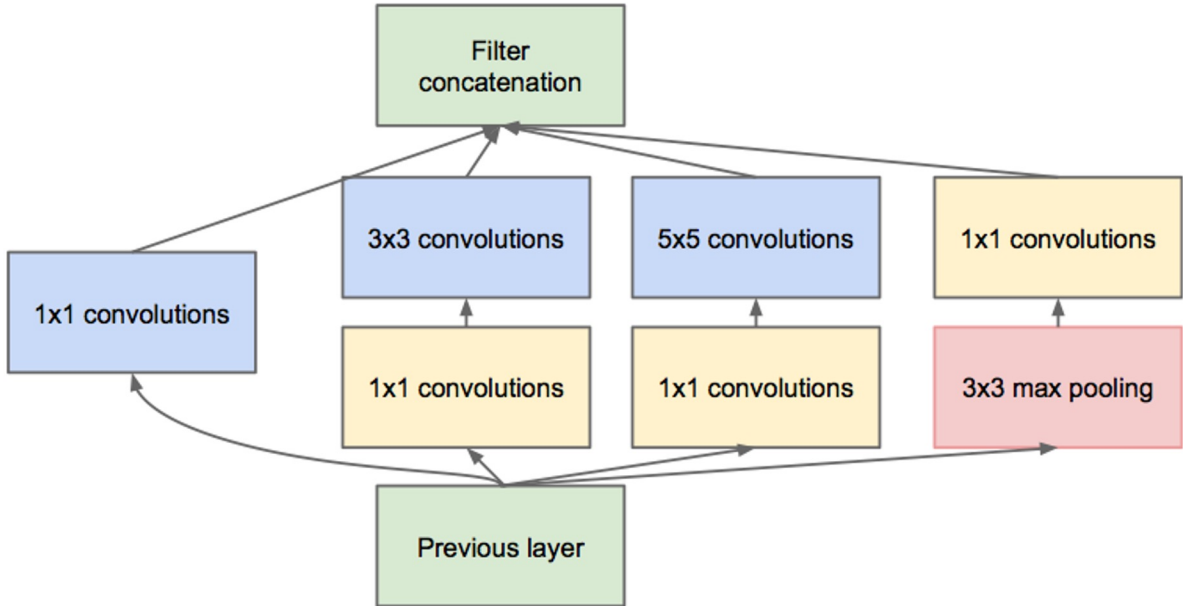

Many innovations for efficiency: reduce parameter count, memory usage, and computation

GoogLeNet



Stem network at the start aggressively downsamples input (Recall in VGG-16: Most of the compute was at the start)


GoogLeNet



Inception module

Local unit with parallel branches
 Local structure repeated many times throughout the network
 Uses 1x1 "Bottleneck" layers to reduce channel dimension before expensive conv

GoogLeNet




Auxiliary Classifiers

Training using loss at the end of the network didn't work well: Network is too deep, gradients don't propagate cleanly

As a hack, attach "auxiliary classifiers" at several intermediate points in the network that also try to classify the image and receive loss

GoogLeNet was before batch normalization! With BatchNorm no longer need to use this trick

ImageNet Classification Challenge

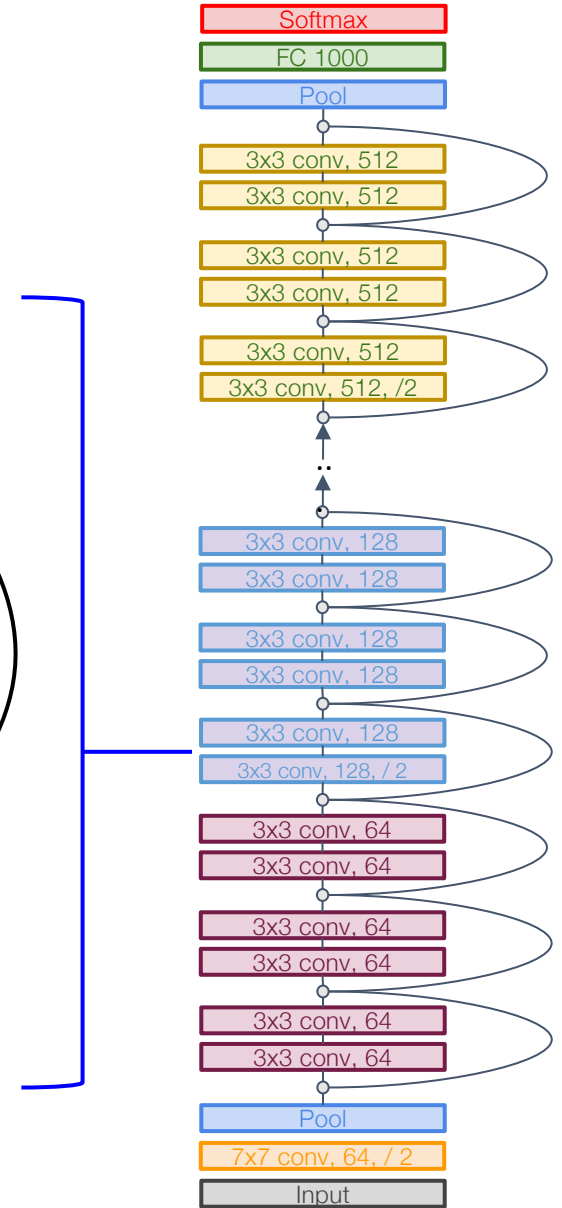



Residual Net (ResNet)



A residual network is a stack of many residual blocks

Regular design, like VGG: each residual block has two 3x3 conv

Network is divided into **stages**: the first block of each stage halves the resolution (with stride-2 conv) and doubles the number of channels



Residual Net (ResNet)





Residual Learning

Fixed identity
// learned residual


$$\mathbf{x}_{n+5} = \mathbf{x}_n + \underbrace{(\phi_{\text{ReLU}} \circ \phi_* \circ \phi_{\text{ReLU}} \circ \phi_*)}_{\text{residual}}(\mathbf{x}_n)$$

↑
identity

K. He, X. Zhang, S. Ren, and J. Sun.
Deep residual learning for image recognition. In CVPR 2016.



Residual Learning




- “Overly deep” plain nets have **higher training error**
- A general phenomenon, observed in many datasets
- This is optimization issue, deeper models are harder to optimize


Residual Learning

- Richer solution space
- A deeper model should not have **higher training error**
- A solution by construction:
 - original layers: copied from a learned shallower model
 - extra layers: set as identity
 - at least the same training error


a shallower model
(18 layers)



a deeper counterpart
(34 layers)




Residual Learning





- The loss surface of a 56-layer net using the CIFAR-10 dataset, both without (left) and with (right) residual connections.

ImageNet Classification Challenge





Comparing Complexity





Comparing Complexity



Inception-v4: Resnet + Inception!




Comparing Complexity




Comparing Complexity





Comparing Complexity




AlexNet: Low compute, lots of parameters



Comparing Complexity



ImageNet Classification Challenge





Today:
More recent CNN
architectures



Post-ResNet Architectures

ResNet made it possible to increase accuracy with larger, deeper models

Many followup architectures emphasize **efficiency**: can we improve accuracy while controlling for model “complexity”?



Measures of Model Complexity

Parameters: How many learnable parameters does the model have?



Floating Point Operations (FLOPs): How many arithmetic operations does it take to compute the forward pass of the model?

Watch out, lots of subtlety here:

- Many papers only count operations in conv layers (ignore ReLU, pooling, BatchNorm).
Most papers use "1 FLOP" = "1 multiply and 1 addition" so dot product of two N-dim vectors takes N FLOPs; some papers say MADD or MACC instead of FLOP
- Other sources (e.g. NVIDIA marketing material) count "1 multiply and one addition" = 2 FLOPs, so dot product of two N-dim vectors takes 2N FLOPs

Network Runtime: How long does a forward pass of the model take on real hardware?


Comparing Complexity




Key ingredient:
Grouped / Separable convolution

Recall: Convolution Layer


Each filter has the same number of channels as the input



Input: $C_{in} \times H \times W$



Weights: $C_{out} \times C_{in} \times K \times K$






Output: $C_{out} \times H' \times W'$

Recall: Convolution Layer

Each filter has the same number of channels as the input

Each plane of the output depends on the full input and one filter



Input: $C_{in} \times H \times W$


Weights: $C_{out} \times C_{in} \times K \times K$

Output: $C_{out} \times H' \times W'$


Recall: Convolution Layer

Each filter has the same number of channels as the input


Each plane of the output depends on the full input and one filter



Input: $C_{in} \times H \times W$



Weights: $C_{out} \times C_{in} \times K \times K$




Output: $C_{out} \times H' \times W'$


Recall: Convolution Layer

Each filter has the same number of channels as the input


Each plane of the output depends on the full input and one filter



Input: $C_{in} \times H \times W$



Weights: $C_{out} \times C_{in} \times K \times K$






Output: $C_{out} \times H' \times W'$

Recall: Convolution Layer

Each filter has the same number of channels as the input

Each plane of the output depends on the full input and one filter




Input: $C_{in} \times H \times W$

Weights: $C_{out} \times C_{in} \times K \times K$


Output: $C_{out} \times H' \times W'$

Grouped Convolution





Grouped Convolution

Divide channels of input into G groups with (C_{in}/G) channels each



Example:
 $G=2$



Input: $C_{in} \times H \times W$


Weights: $C_{out} \times C_{in} \times K \times K$

Output: $C_{out} \times H' \times W'$



Grouped Convolution

Divide filters into G groups; each group looks at a **subset** of input channels

Divide channels of input into G groups with (C_{in}/G) channels each



Example:
 $G=2$



Input: $C_{in} \times H \times W$

Weights: $C_{out} \times (C_{in}/G) \times K \times K$


Output: $C_{out} \times H' \times W'$

Grouped Convolution



Divide filters into G groups; each group looks at a **subset** of input channels

Each plane of the output depends on one filter and a **subset** of the input channels

Divide channels of input into G groups with (C_{in}/G) channels each



Example:
 $G = 2$



Input: $C_{in} \times H \times W$

Weights: $C_{out} \times (C_{in}/G) \times K \times K$


Output: $C_{out} \times H' \times W'$

Group Convolution



Divide channels of input into G groups with (C_{in}/G) channels each

Divide filters into G groups; each group looks at a **subset** of input channels

Each plane of the output depends on one filter and a **subset** of the input channels



Example:
 $G=2$



Input: $C_{in} \times H \times W$


Weights: $C_{out} \times (C_{in}/G) \times K \times K$ Output: $C_{out} \times H' \times W'$

Group Convolution



Divide channels of input into G groups with (C_{in}/G) channels each

Divide filters into G groups; each group looks at a **subset** of input channels

Each plane of the output depends on one filter and a **subset** of the input channels



Example:
 $G=2$



Input: $C_{in} \times H \times W$


Weights: $C_{out} \times (C_{in}/G) \times K \times K$ Output: $C_{out} \times H' \times W'$

Group Convolution



Divide channels of input into G groups with (C_{in}/G) channels each

Divide filters into G groups; each group looks at a **subset** of input channels

Each plane of the output depends on one filter and a **subset** of the input channels



Example:
 $G=2$



Input: $C_{in} \times H \times W$


Weights: $C_{out} \times (C_{in}/G) \times K \times K$ Output: $C_{out} \times H' \times W'$

Group Convolution


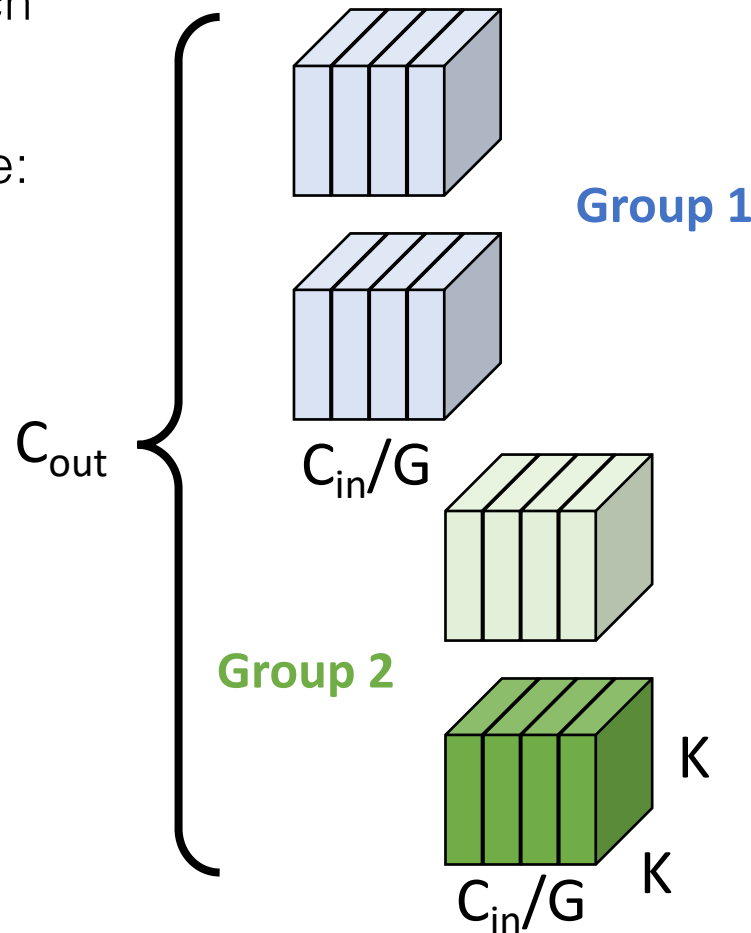
Divide channels of input into G groups with (C_{in}/G) channels each

Divide filters into G groups; each group looks at a **subset** of input channels

Each plane of the output depends on one filter and a **subset** of the input channels



Example:
 $G=2$



Input: $C_{in} \times H \times W$


Weights: $C_{out} \times (C_{in}/G) \times K \times K$ Output: $C_{out} \times H' \times W'$

Group Convolution



Divide channels of input into G groups with (C_{in}/G) channels each

Divide filters into G groups; each group looks at a **subset** of input channels

Each plane of the output depends on one filter and a **subset** of the input channels



Example:
 $G=2$



Input: $C_{in} \times H \times W$


Weights: $C_{out} \times (C_{in}/G) \times K \times K$ Output: $C_{out} \times H' \times W'$

Group Convolution



Divide channels of input into G groups with (C_{in}/G) channels each

Divide filters into G groups; each group looks at a **subset** of input channels

Each plane of the output depends on one filter and a **subset** of the input channels



Example:
 $G=4$




Input: $C_{in} \times H \times W$

Weights: $C_{out} \times (C_{in}/G) \times K \times K$ Output: $C_{out} \times H' \times W'$


Special Case: Depthwise Convolution

Number of groups equals
number of input channels




Input: $C_{in} \times H \times W$

Common to also set $C_{out} = G$



Weights: $C_{out} \times 1 \times K \times K$


Output only mixes **spatial**
information from input;
channel information not
mixed




Output: $C_{out} \times H' \times W'$

Special Case: Depthwise Convolution


Number of groups equals
number of input channels



Can still have multiple filters
per group (e.g. $C_{out} = 2C_{in}$)



Output only mixes **spatial**
information from input;
channel information not
mixed



Input: $C_{in} \times H \times W$

Weights: $C_{out} \times 1 \times K \times K$

Output: $C_{out} \times H' \times W'$

Grouped Convolution vs Standard Convolution

Grouped Convolution (G groups):

G parallel conv layers; each "sees"
 C_{in}/G input channels and produces
 C_{out}/G output channels

Input: $C_{in} \times H \times W$

Split to $G \times [(C_{in} / G) \times H \times W]$

Weight: $G \times (C_{out} / G) \times (C_{in} / G) \times K \times K$

G parallel convolutions

Output: $G \times [(C_{out} / G) \times H' \times W']$

Concat to $C_{out} \times H' \times W'$

FLOPs: $C_{out} C_{in} K^2 HW / G$

Standard Convolution (groups=1)

Input: $C_{in} \times H \times W$

Weight: $C_{out} \times C_{in} \times K \times K$


Output: $C_{out} \times H' \times W'$

FLOPs: $C_{out} C_{in} K^2 HW$

All convolutional kernels touch
all C_{in} channels of the input

Using G groups reduces
FLOPs by a factor of G!


Improving ResNets



“Bottleneck”
Residual block


Total FLOPs:
 $17HWC^2$

Improving ResNets: ResNeXt



“Bottleneck”
Residual block

Total FLOPs:
 $17HWC^2$

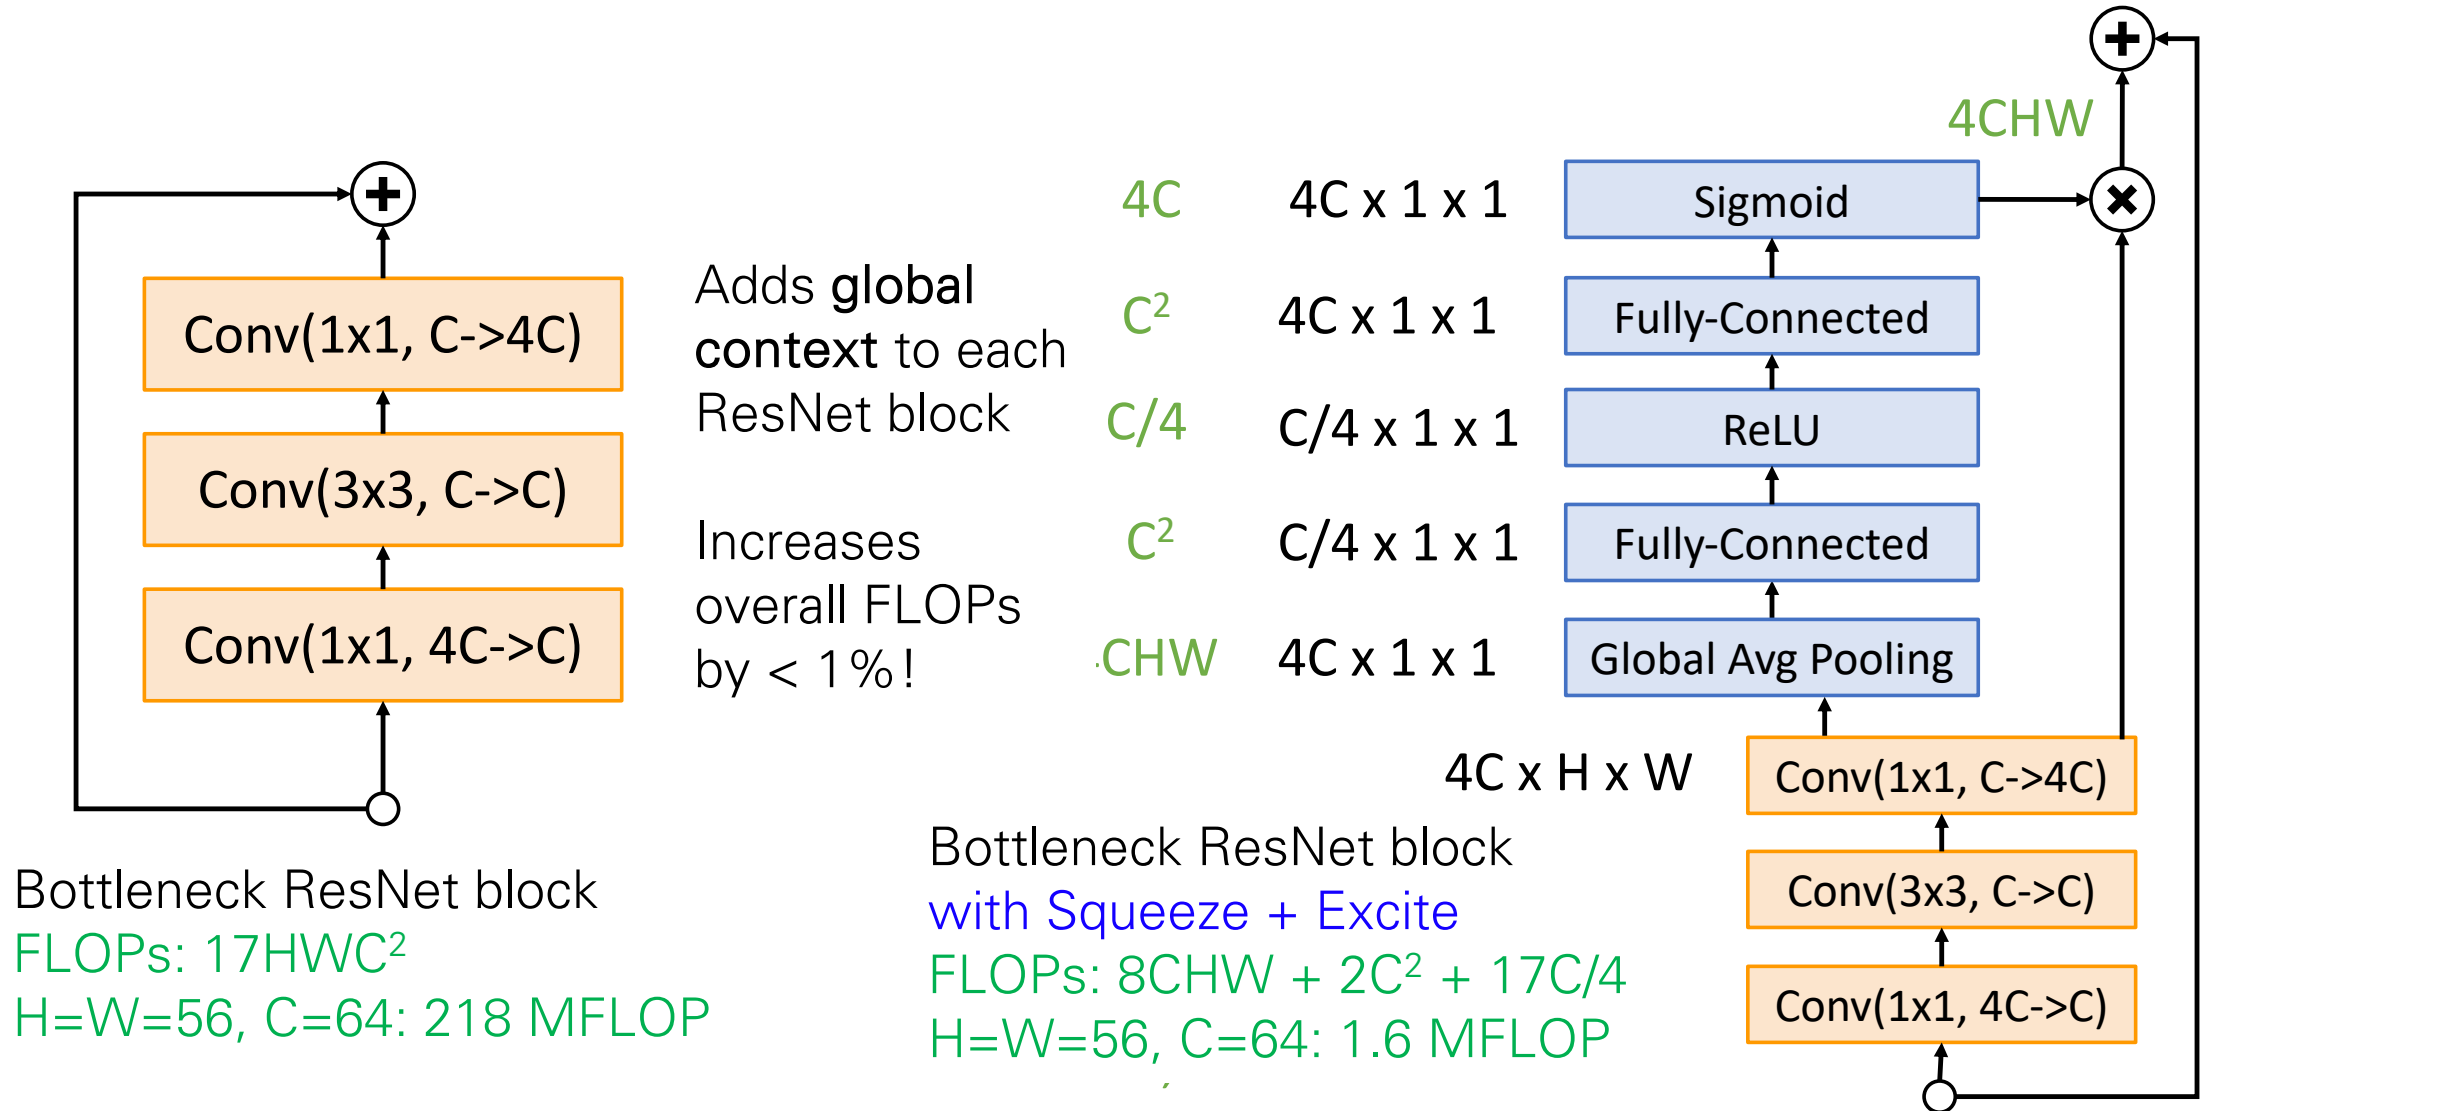


Same FLOPs when
 $9Gc^2 + 8GCc - 17C^2 = 0$

Total FLOPs:
 $(8Cc + 9c^2) * HWG$


Example: $C=64, G=4, c=24$; $C=64, G=32, c=4$

Squeeze-and-Excitation Networks (SENet)



Squeeze-and-Excitation Networks (SENet)

ImageNet Top-1 Accuracy




Add SE to any architecture, enjoy 1-2% boost in accuracy

Recall: Convolution Layer

Instead of pushing for the largest network with biggest accuracy, consider tiny networks and accuracy / complexity tradeoff

Compare **families** of models:


One family is better than another if it moves the whole curve up and to the left



MobileNets: Tiny Networks (For Mobile Devices)

Standard Convolution Block

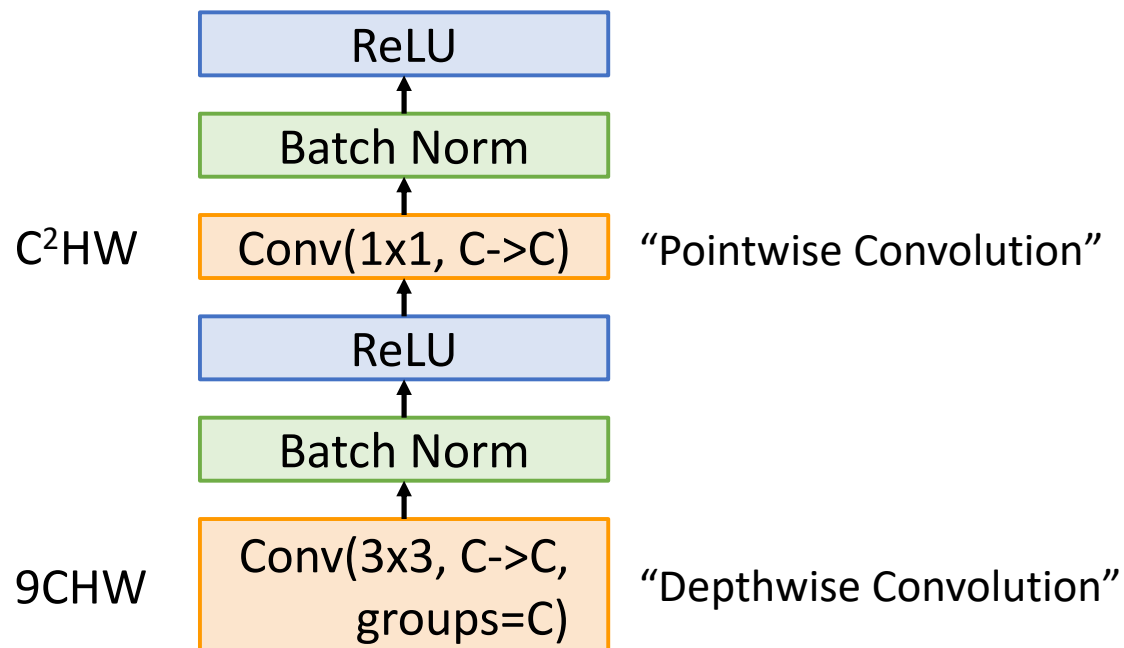
Total cost: $9C^2HW$



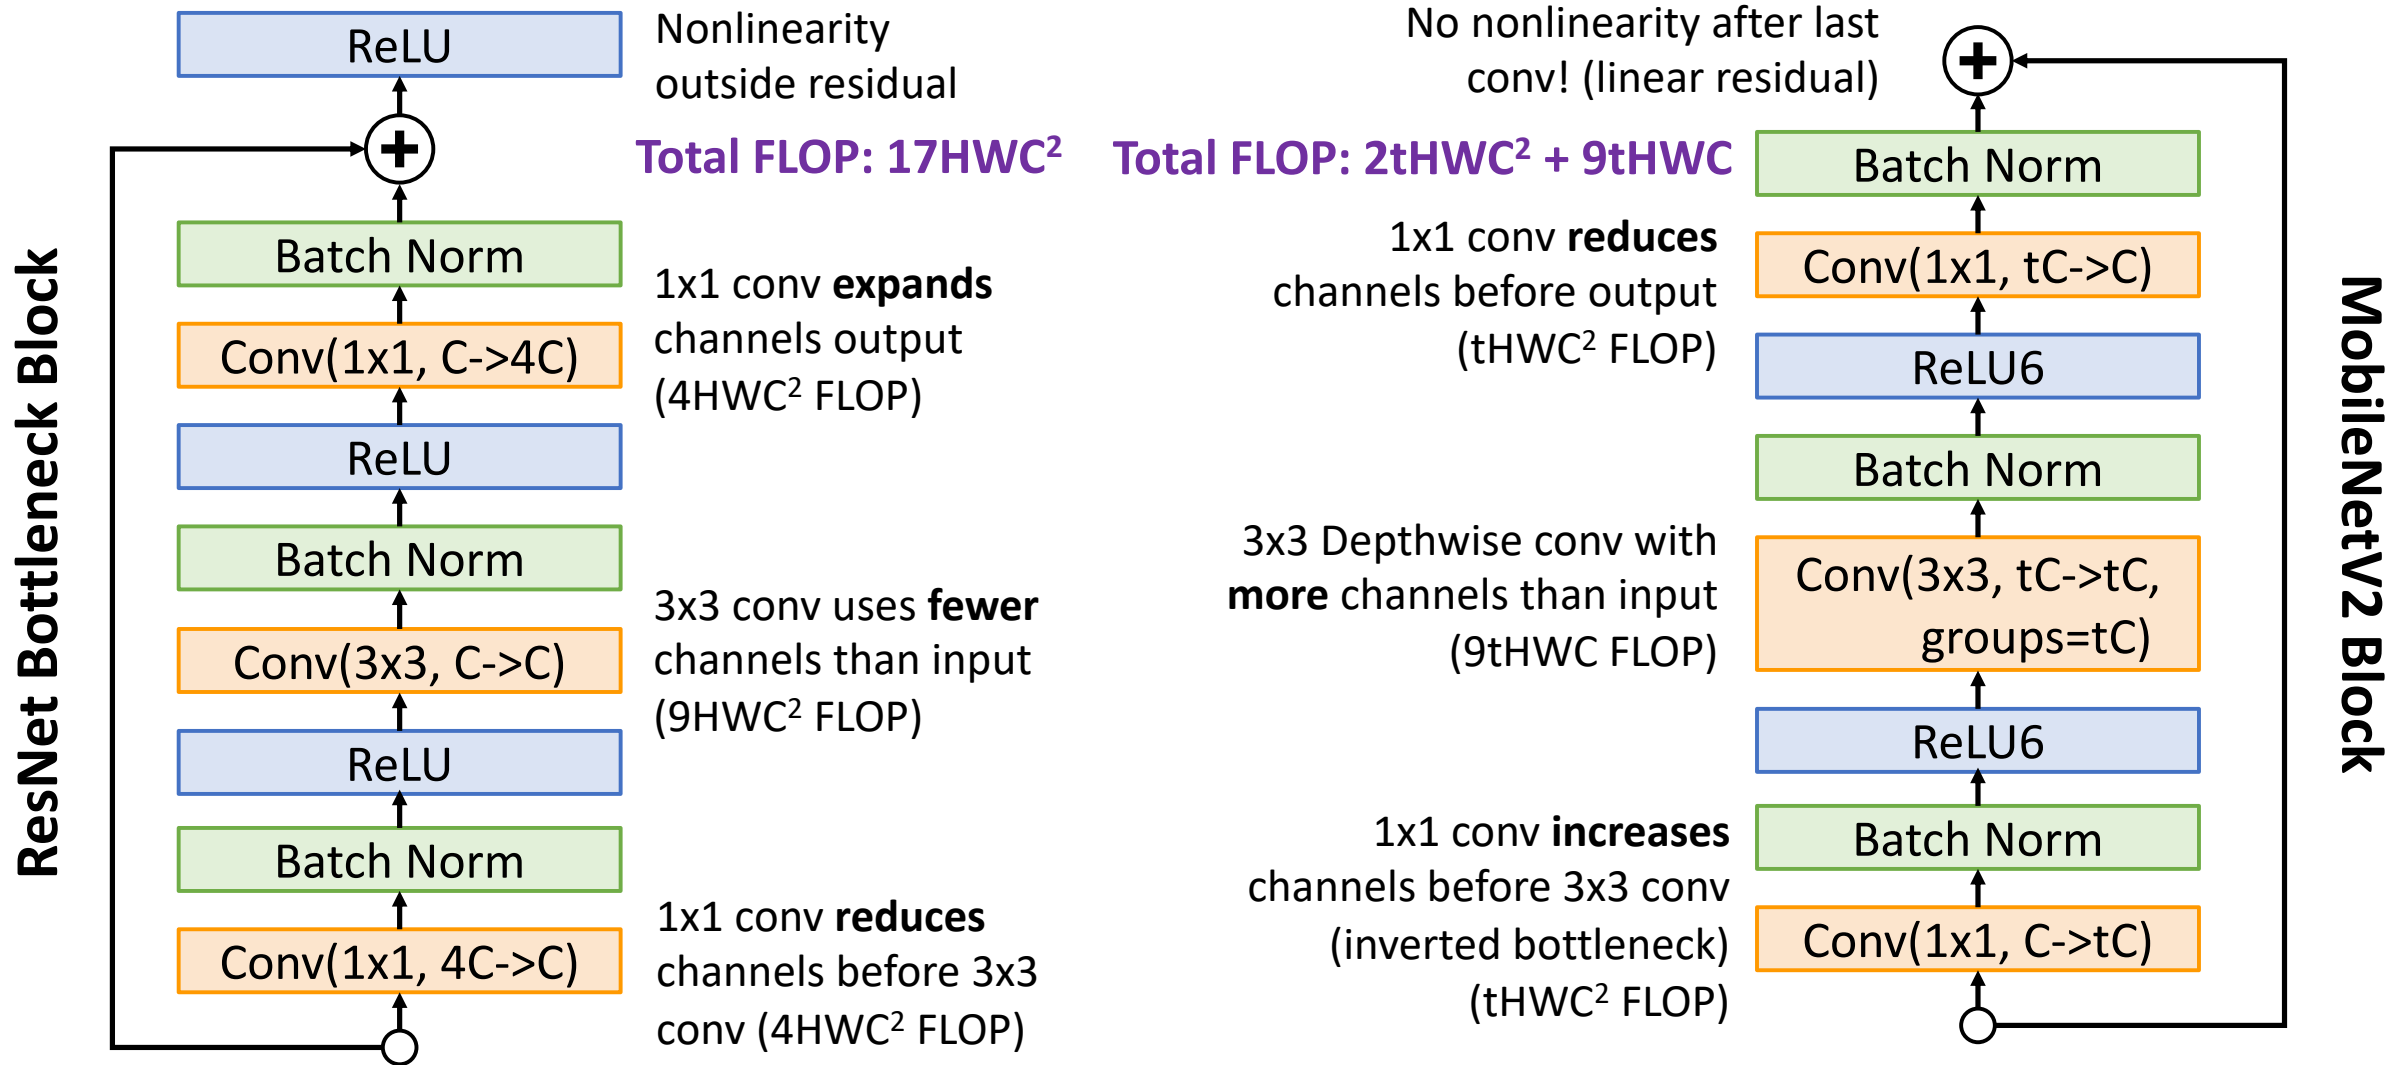
$$\begin{aligned}\text{Speedup} &= 9C^2 / (9C + C^2) \\ &= 9C / (9 + C) \\ &\Rightarrow 9 \text{ (as } C \rightarrow \text{inf)}\end{aligned}$$

Depthwise Separable Convolution


Total cost: $(9C + C^2)HW$



MobileNetV2: Inverted Bottleneck, Linear Residual

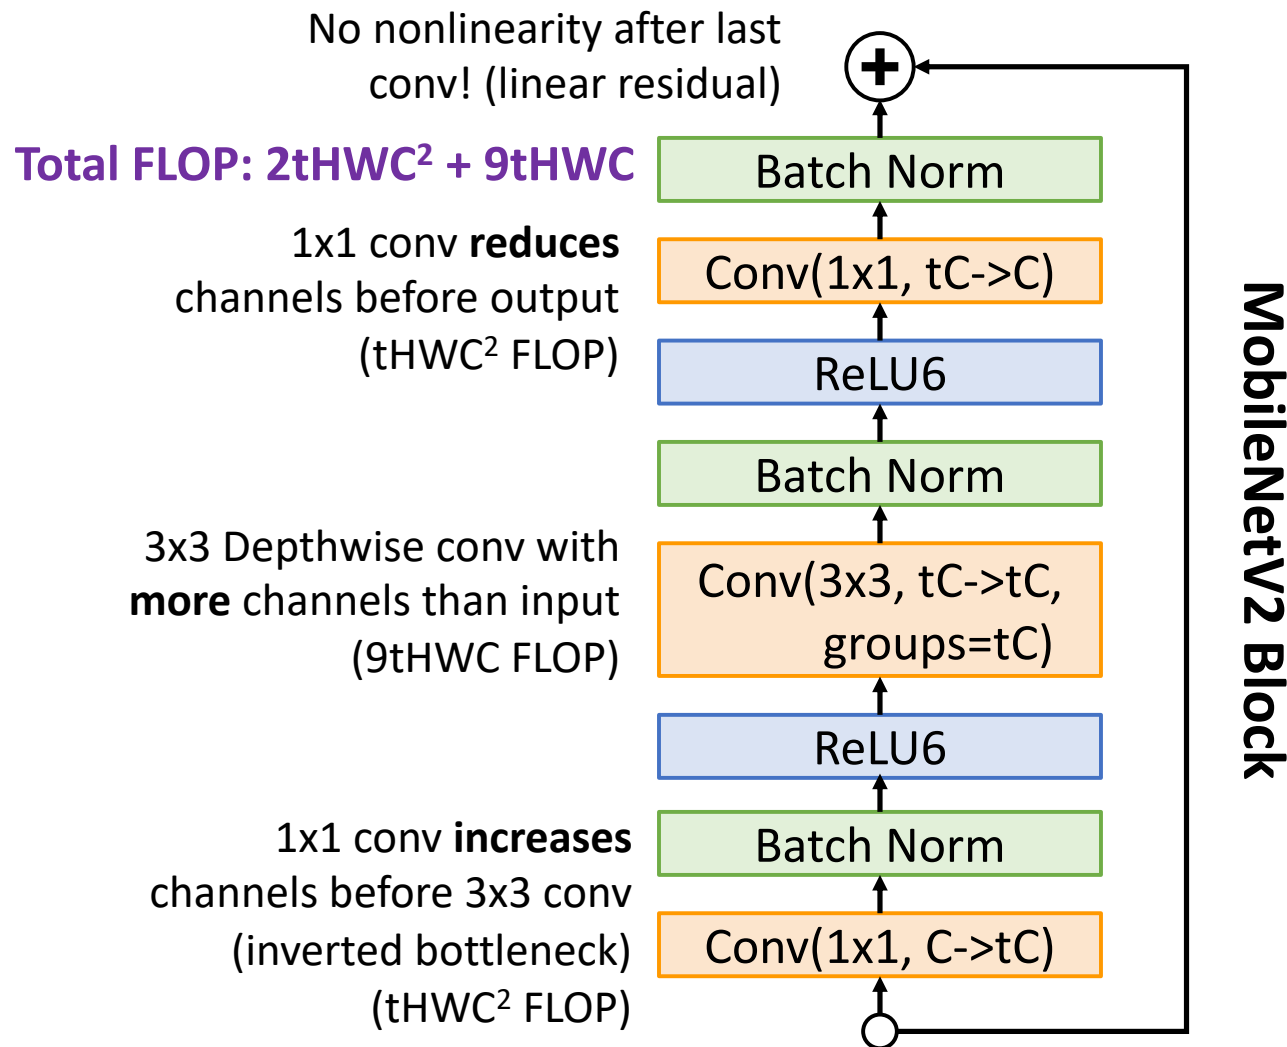


MobileNetV2: Inverted Bottleneck, Linear Residual


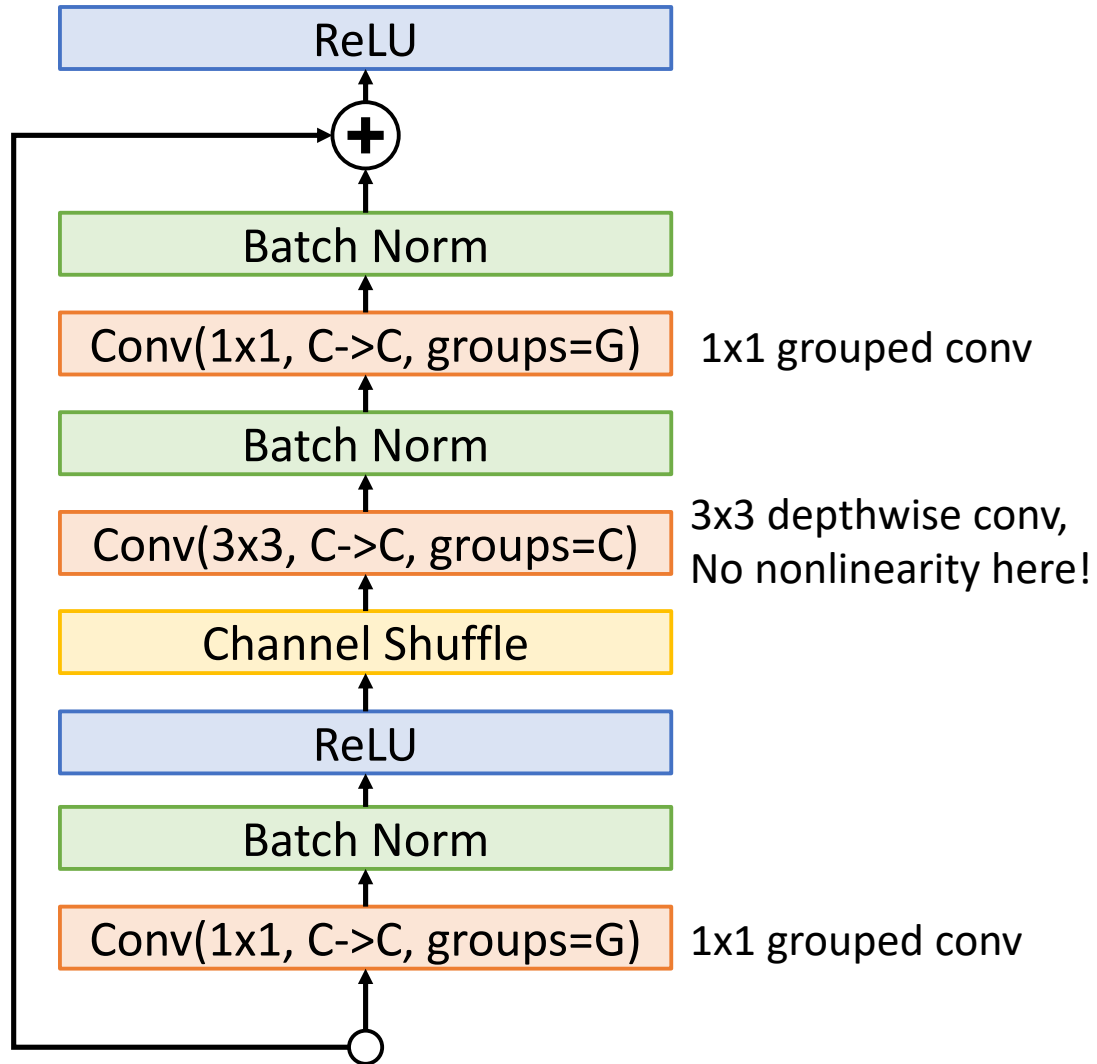


$$ReLU6(x) = \begin{cases} 0 & \text{if } x \leq 0 \\ x & \text{if } 0 < x < 6 \\ 6 & \text{if } x \geq 6 \end{cases}$$

Keeps activations in reasonable range when running inference in low precision



ShuffleNet



CNN Architectures Summary

- Early work (AlexNet->VGG->ResNet): **bigger networks work better**
- New focus on **efficiency**: Improve accuracy, control for network complexity
- Grouped and Depthwise Convolution appear in many modern architectures
- **Squeeze-and-Excite** adds accuracy boost to just about any architecture while only adding a tiny amount of FLOPs and runtime
- Tiny networks for **mobile devices** (MobileNet, ShuffleNet)
- **Neural Architecture Search(NAS)** promised to automate architecture design
- More recent work has moved towards **careful improvements to ResNet-like architectures**
- ResNet and ResNeXt are still surprisingly strong and popular architectures!


Transfer Learning with Convolutional Neural Networks

Beyond CNNs

- Do features extracted from the CNN generalize other tasks and datasets?
 - Donahue et al. (2013), Chatfield et al. (2014), Razavian et al. (2014), Yosinski et al. (2014), etc.
- CNN activations as deep features
- Finetuning CNNs


CNN activations as deep features

- CNNs discover effective representations. Why not to use them?




CNN activations as deep features

- CNNs discover effective representations. Why not to use them?



CNN activations as deep features


- CNNs discover effective representations. Why not to use them?



Zeiler et al., 2014

CNN activations as deep features

- CNNs discover effective representations. Why not




Zeiler et al., 2014

Nguyen et al., 2014

CNNs as deep features

- CNNs discover effective representations. Why not to use them?



Transfer Learning with CNNs

- A CNN trained on a (large enough) dataset generalizes to other visual tasks



Transfer Learning with CNNs

- Keep layers 1-7 of our ImageNet-trained model fixed
- Train a new softmax classifier on top using the training images of the new dataset.



1. Train on Imagenet



2. Small dataset: feature extractor

Freeze these

Train this



3. Medium dataset: finetuning

more data = retrain more of the network (or all of it)


Freeze these

tip: use only $\sim 1/10$ th of the original learning rate in finetuning top layer, and $\sim 1/100$ th on intermediate layers

Train this


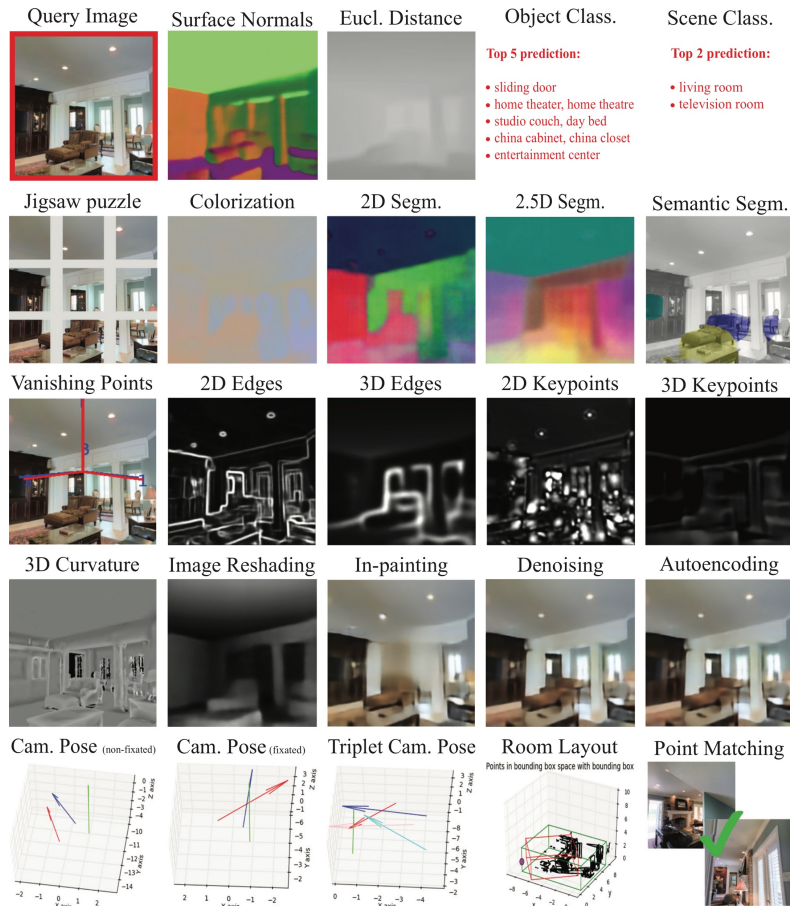
How transferable are features in CNN networks?

- Divide ImageNet into man-made objects A (449 classes) and natural objects B (551 classes)
- The transferability of features decreases as the distance between the base task and target task increases



How transferable are features in CNN networks?

- An open research problem



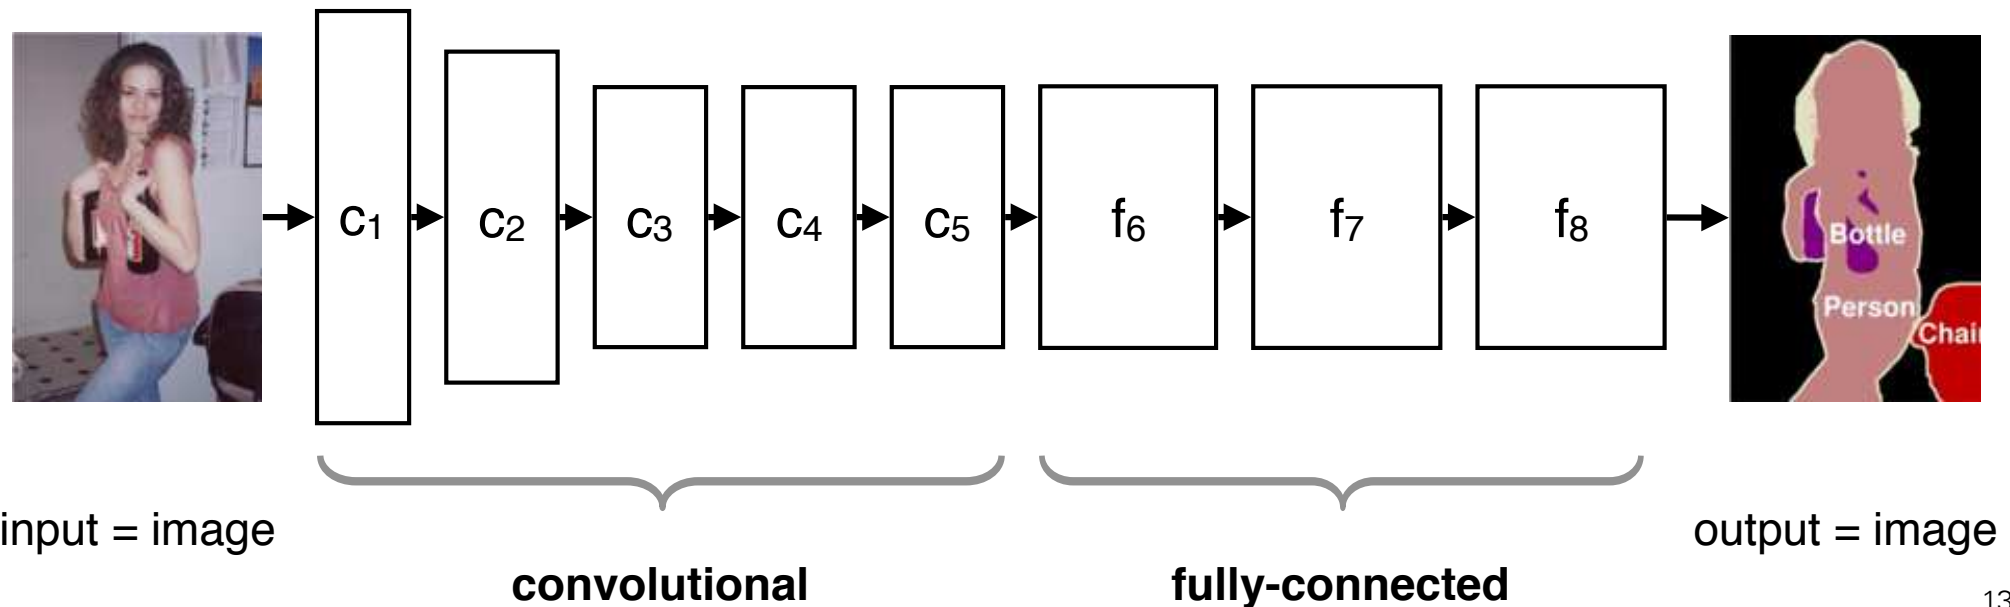

Slide credit: Amir Zamir

Semantic Segmentation




Semantic Image Segmentation

- Label individual pixels



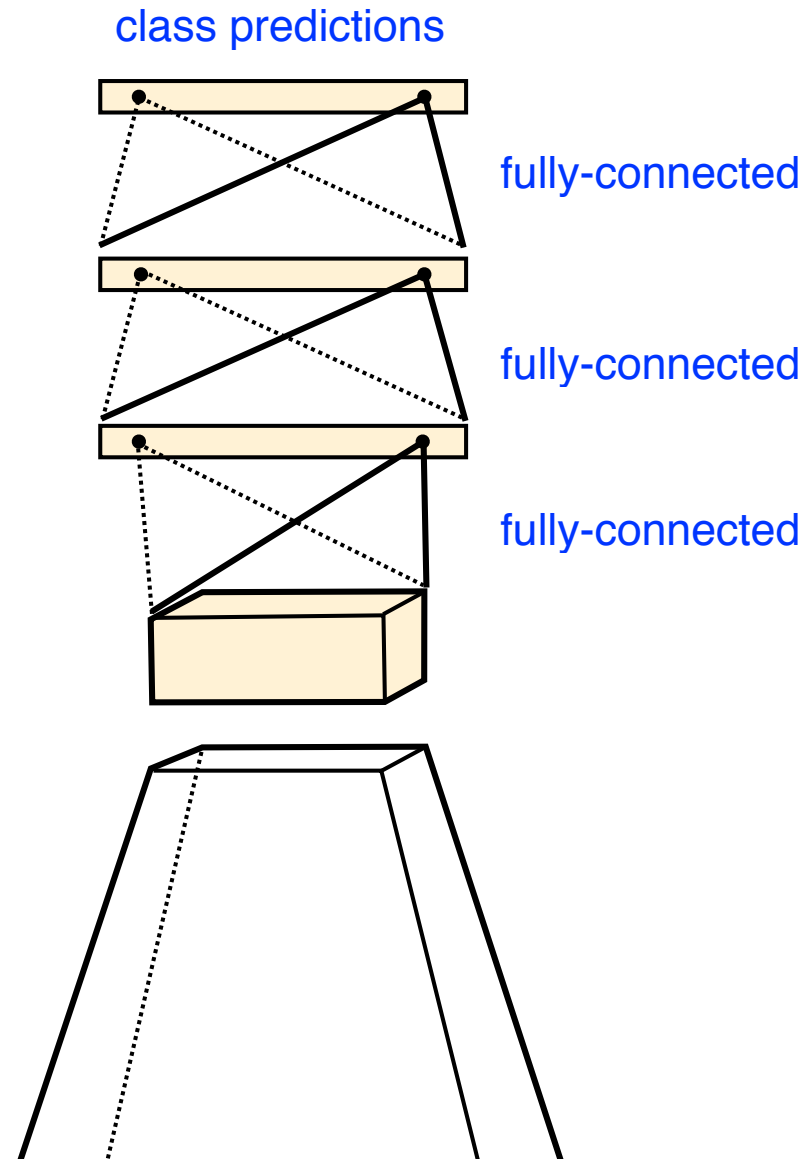
Convolutional Layers

- Local receptive field



Fully Connected Layers

- Global receptive field




Convolutional vs. Fully Connected

- Comparing the receptive fields


Downsampling filters

Responses are spatially selective, can be used to localize things.




Upsampling filters

Responses are global, do not characterize well position




Which one is more useful for pixel level labelling?


Fully-Connected Layer = Large Filter



Fully-Convolutional Neural Networks



Fully-Convolutional Neural Networks



- **Dense evaluation**

- Apply the whole network convolutional
- Estimates a vector of class probabilities at each pixel


- **Downsampling**

- In practice most network downsample the data fast
- The output is very low resolution (e.g. 1/32 of original)


Upsampling The Resolution

- Interpolating filter

Downsampling filters



Upsampling filters




Upsampling filters allow to increase the resolution of the output

Very useful to get full-resolution segmentation results

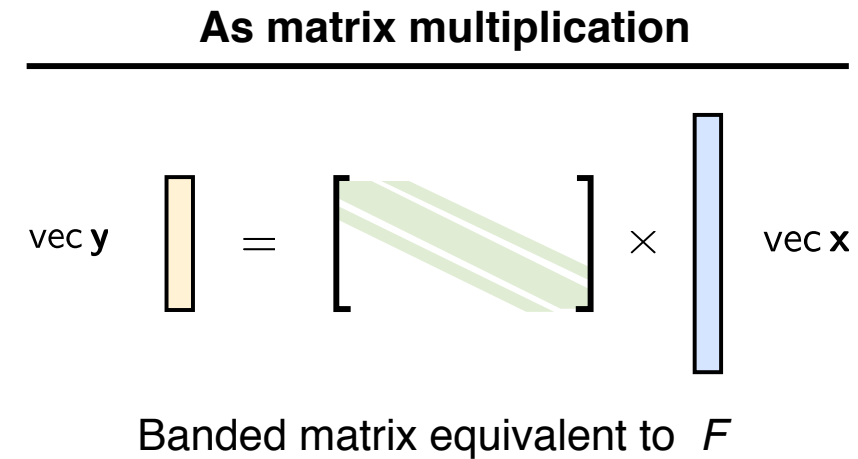

Deconvolution Layer

- Or convolution transpose




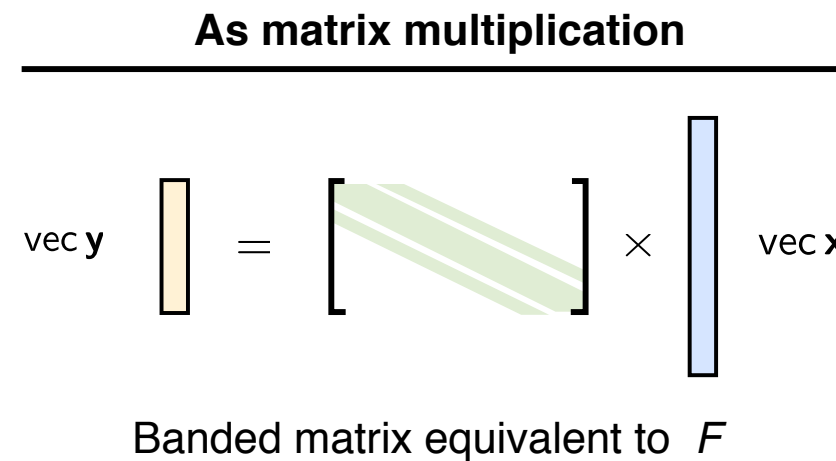


Deconvolution Layer

- Or convolution transpose




Deconvolution Layer

- Or convolution transpose



U-Architectures

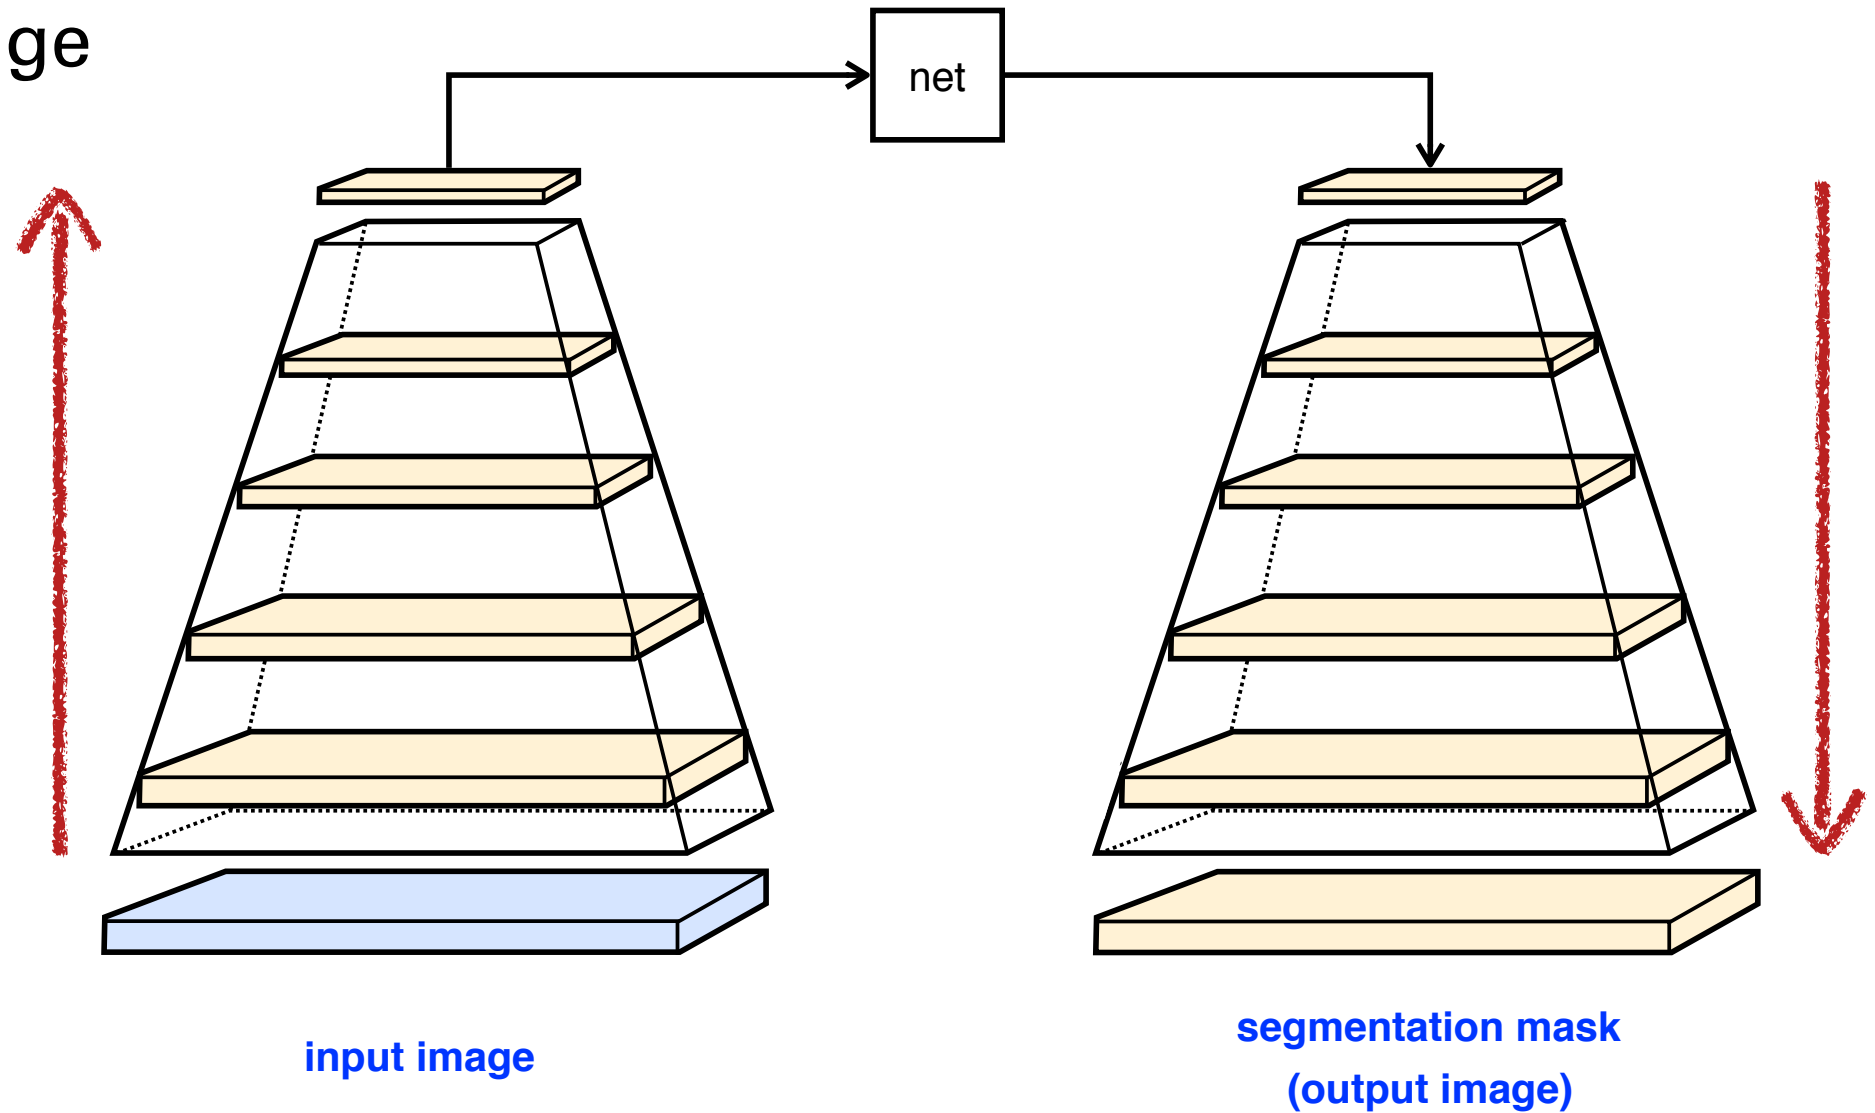
- Image to image



input image

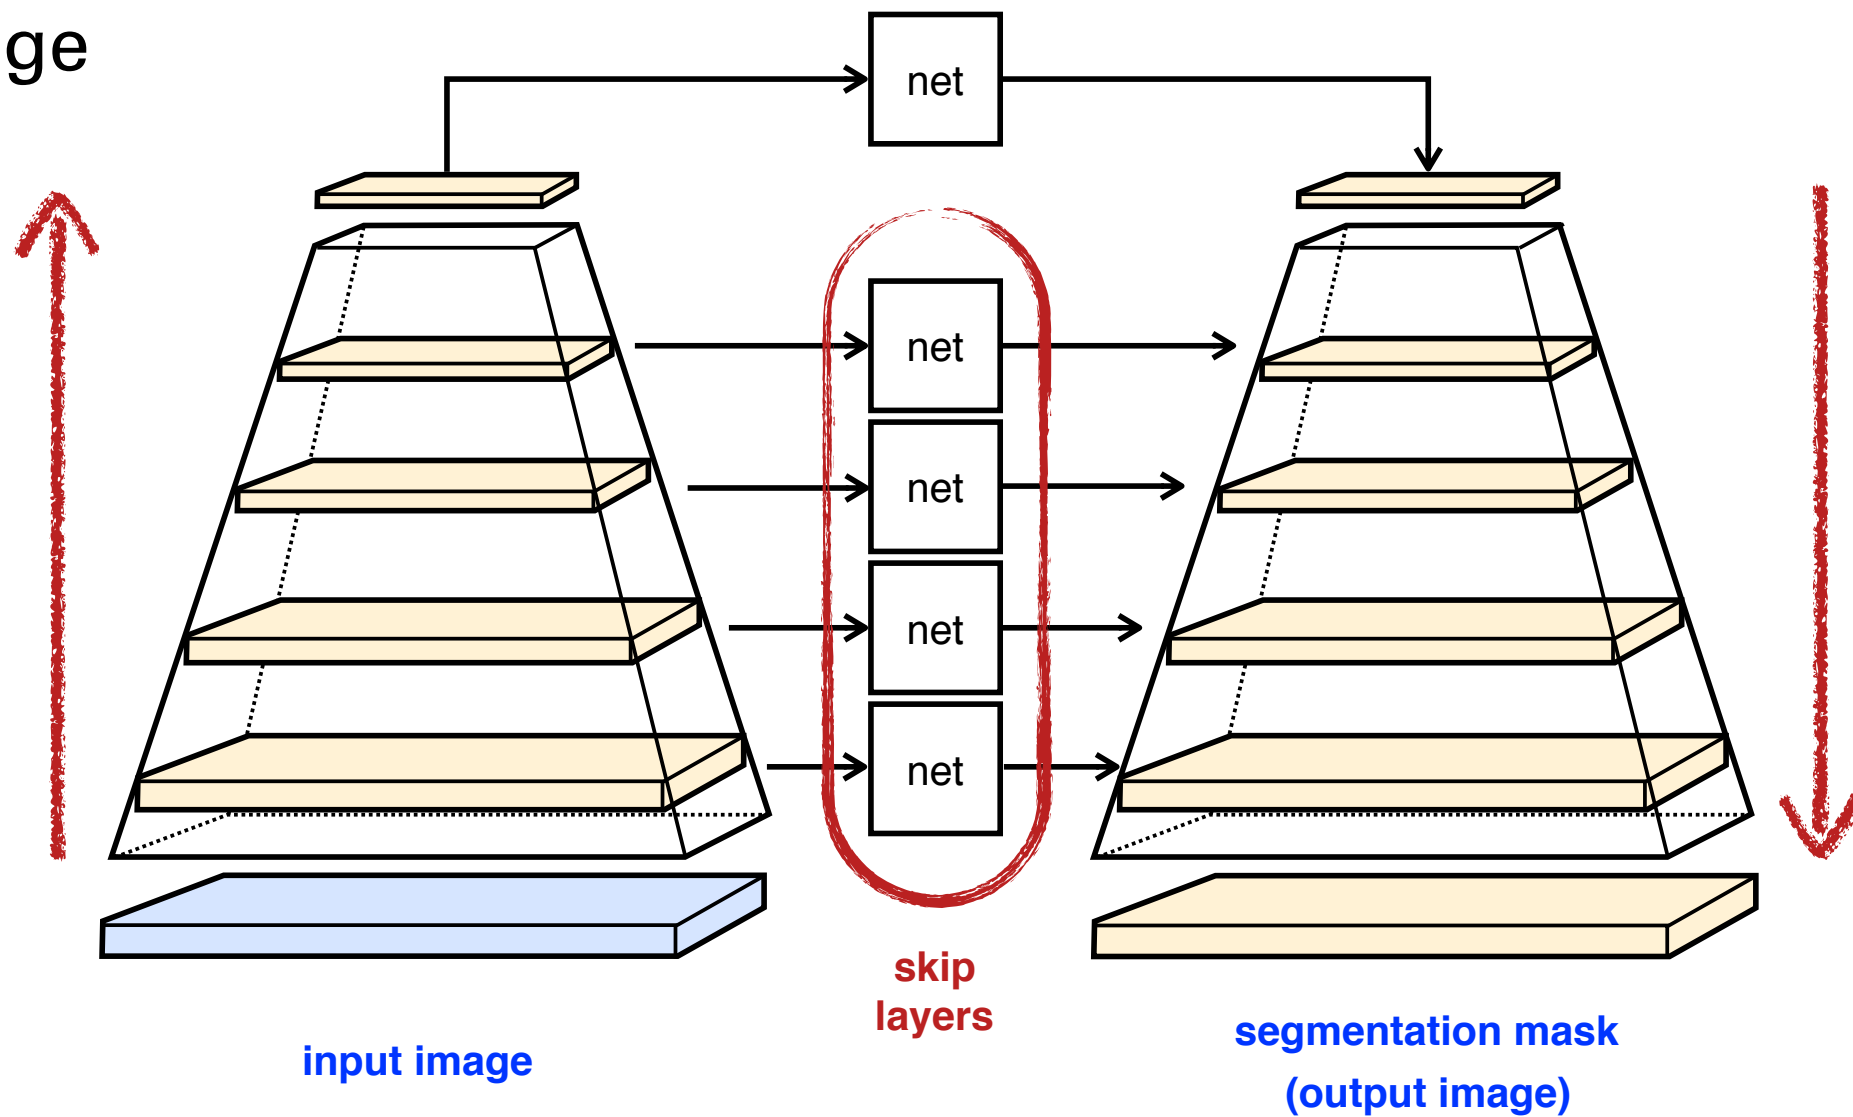
U-Architectures

- Image to image






U-Architectures

- Image to image



U-Architectures

- Several variants: FCN, U-arch, deconvolution, ...




J. Long, E. Shelhamer, and T. Darrell. **Fully convolutional models for semantic segmentation.** In CVPR 2015

H. Noh, S. Hong, and B. Han. **Learning deconvolution network for semantic segmentation.** In ICCV 2015

O. Ronneberger, P. Fischer, and T. Brox. **U-net: Convolutional networks for biomedical image segmentation.** In MICCAI 2015

Object Detection



MS COCO Dataset Images




MS COCO Annotations

- 80 different categories



MS COCO


Dataset Images
+
Annotations



COCO Object Detection Average Precision (%)

- Area under a detector's precision-recall curve, averaged over...
 - Object categories
 - **True positive overlap requirement** (IoU from 0.5 to 0.95; see below)

boxes



masks


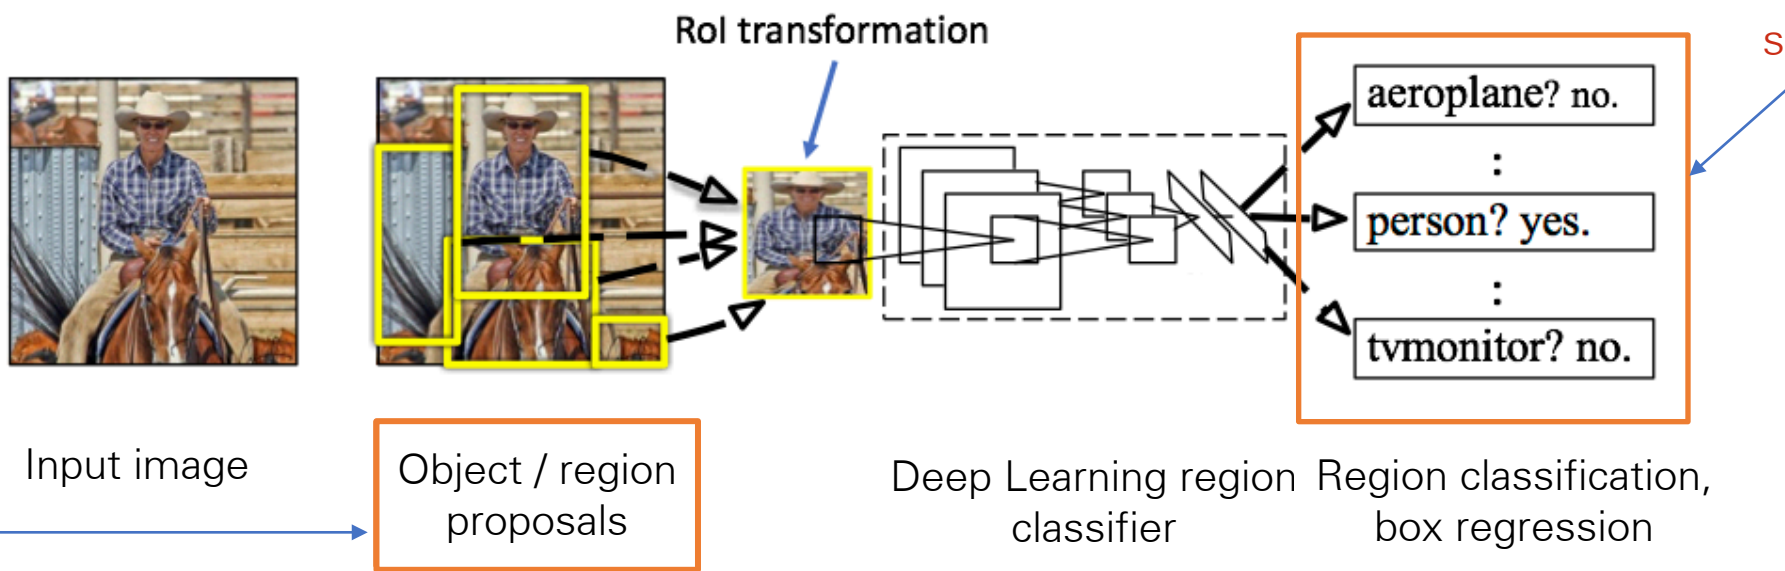


Figure credits: Dollár and Zitnick (top), Krähenbühl and Kulton (bottom)

More than one "stage" (≈proposal based; but doesn't require proposals)



classification of reduced output space elements

Cascade-like reduction in output space



One stage

Direct classification
Of all output space elements



"You only look once"
"Single shot"

COCO Object Detection Average Precision (%)


Past
(best circa
2012)

5




DPM
(Pre DL)

COCO Object Detection Average Precision (%)




COCO Object Detection Average Precision (%)



COCO Object Detection Average Precision (%)

Past
(best circa
2012)


Early
2015



COCO Object Detection Average Precision (%)

Past
(best circa
2012)


Early
2015



COCO Object Detection Average Precision (%)

Past
(best circa
2012)

Early
2015




COCO Object Detection Average Precision (%)


Past
(best circa
2012)

Early
2015

2017



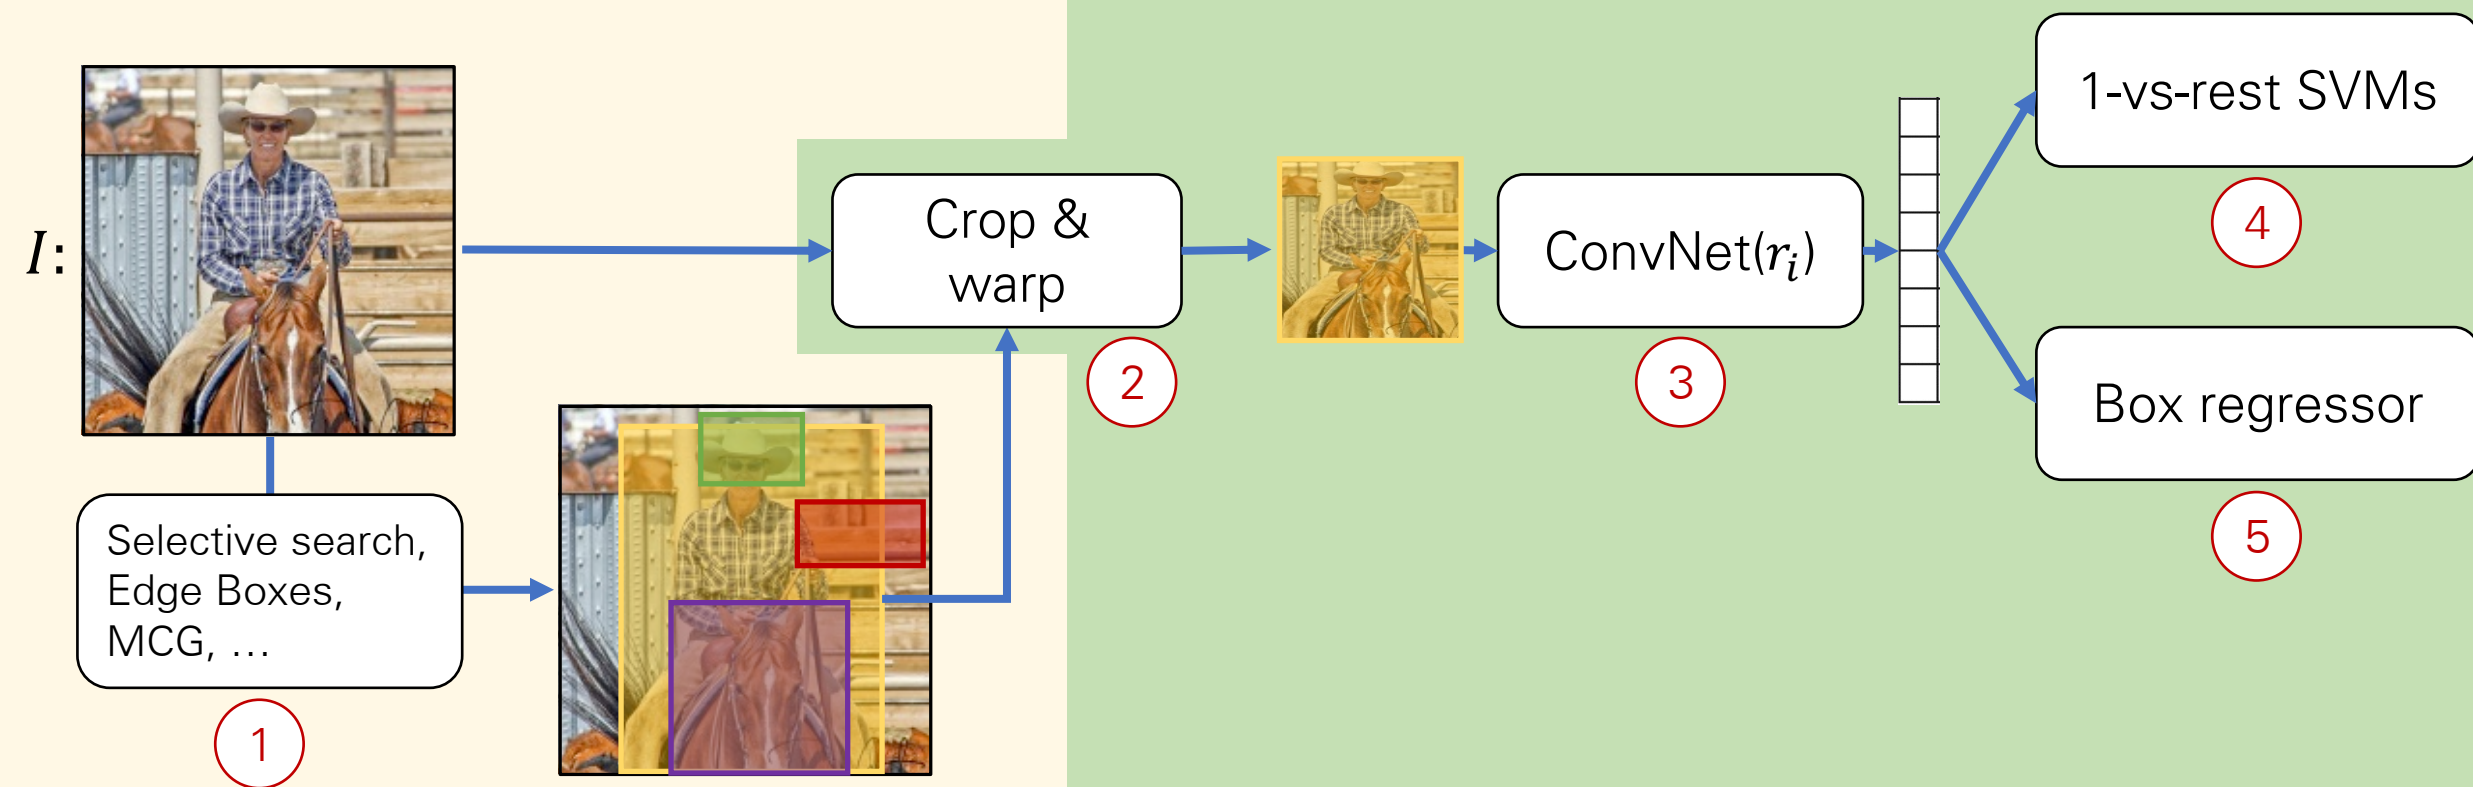
COCO Object Detection Average Precision (%)



"Slow" R-CNN

Per-image computation

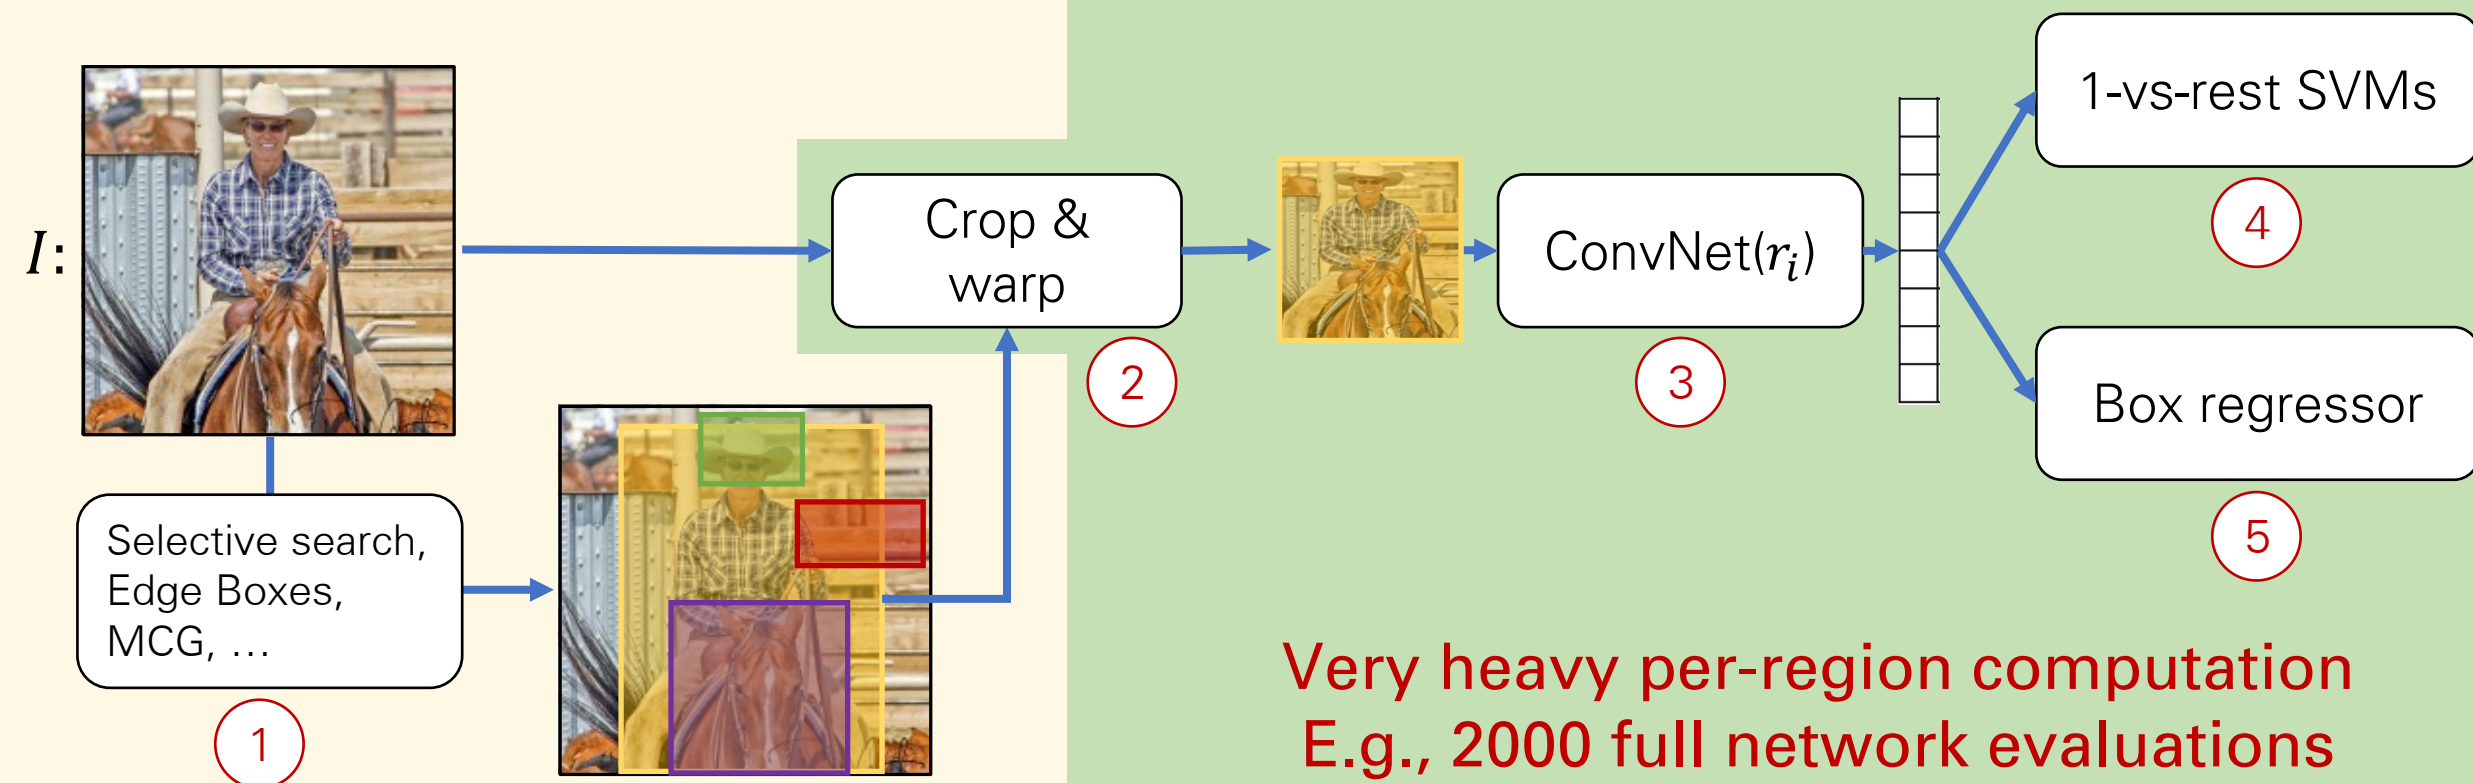
Per-region computation for each $r_i \in r(I)$



"Slow" R-CNN

Per-image computation

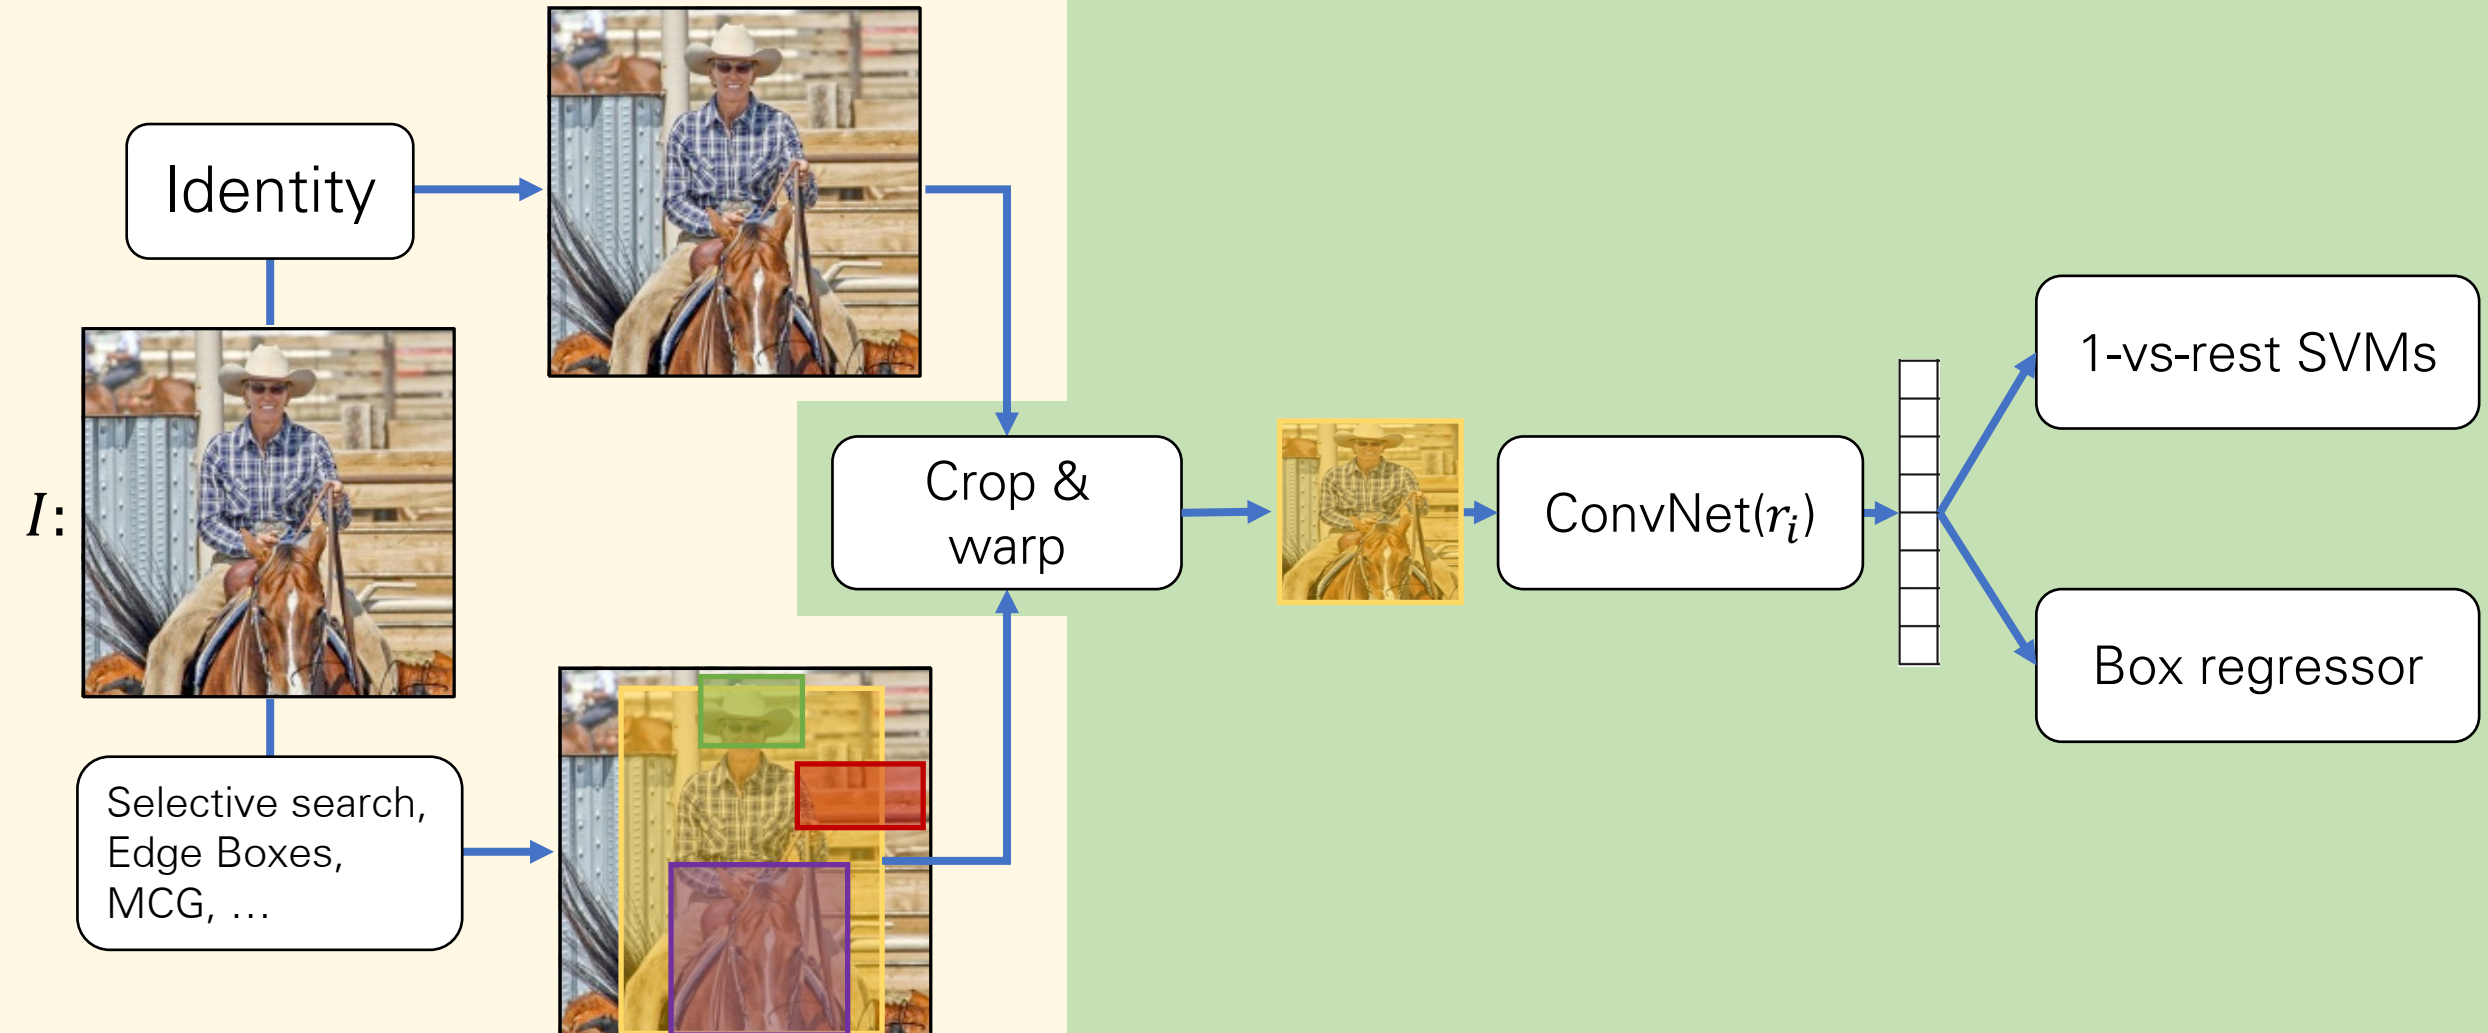
Per-region computation for each $r_i \in r(I)$



"Slow" R-CNN

Per-image computation


Per-region computation for each $r_i \in r(I)$



Generalized R-CNN Approach to Detection

Per-image computation


Per-region computation for each $r_i \in r(I)$



Fast R-CNN

Per-image computation


Per-region computation for each $r_i \in r(I)$



Fast R-CNN

Per-image computation




Per-region computation for each $r_i \in r(I)$



Whole-image FCN

- Use **any standard ConvNet** as the “**backbone architecture**”
 - AlexNet, VGG, ResNet, Inception, Inception-ResNet, ResNeXt, DenseNet, ...
 - Use the first N layers with spatial extent (e.g., up to “conv5”)

Example:
ResNet-34

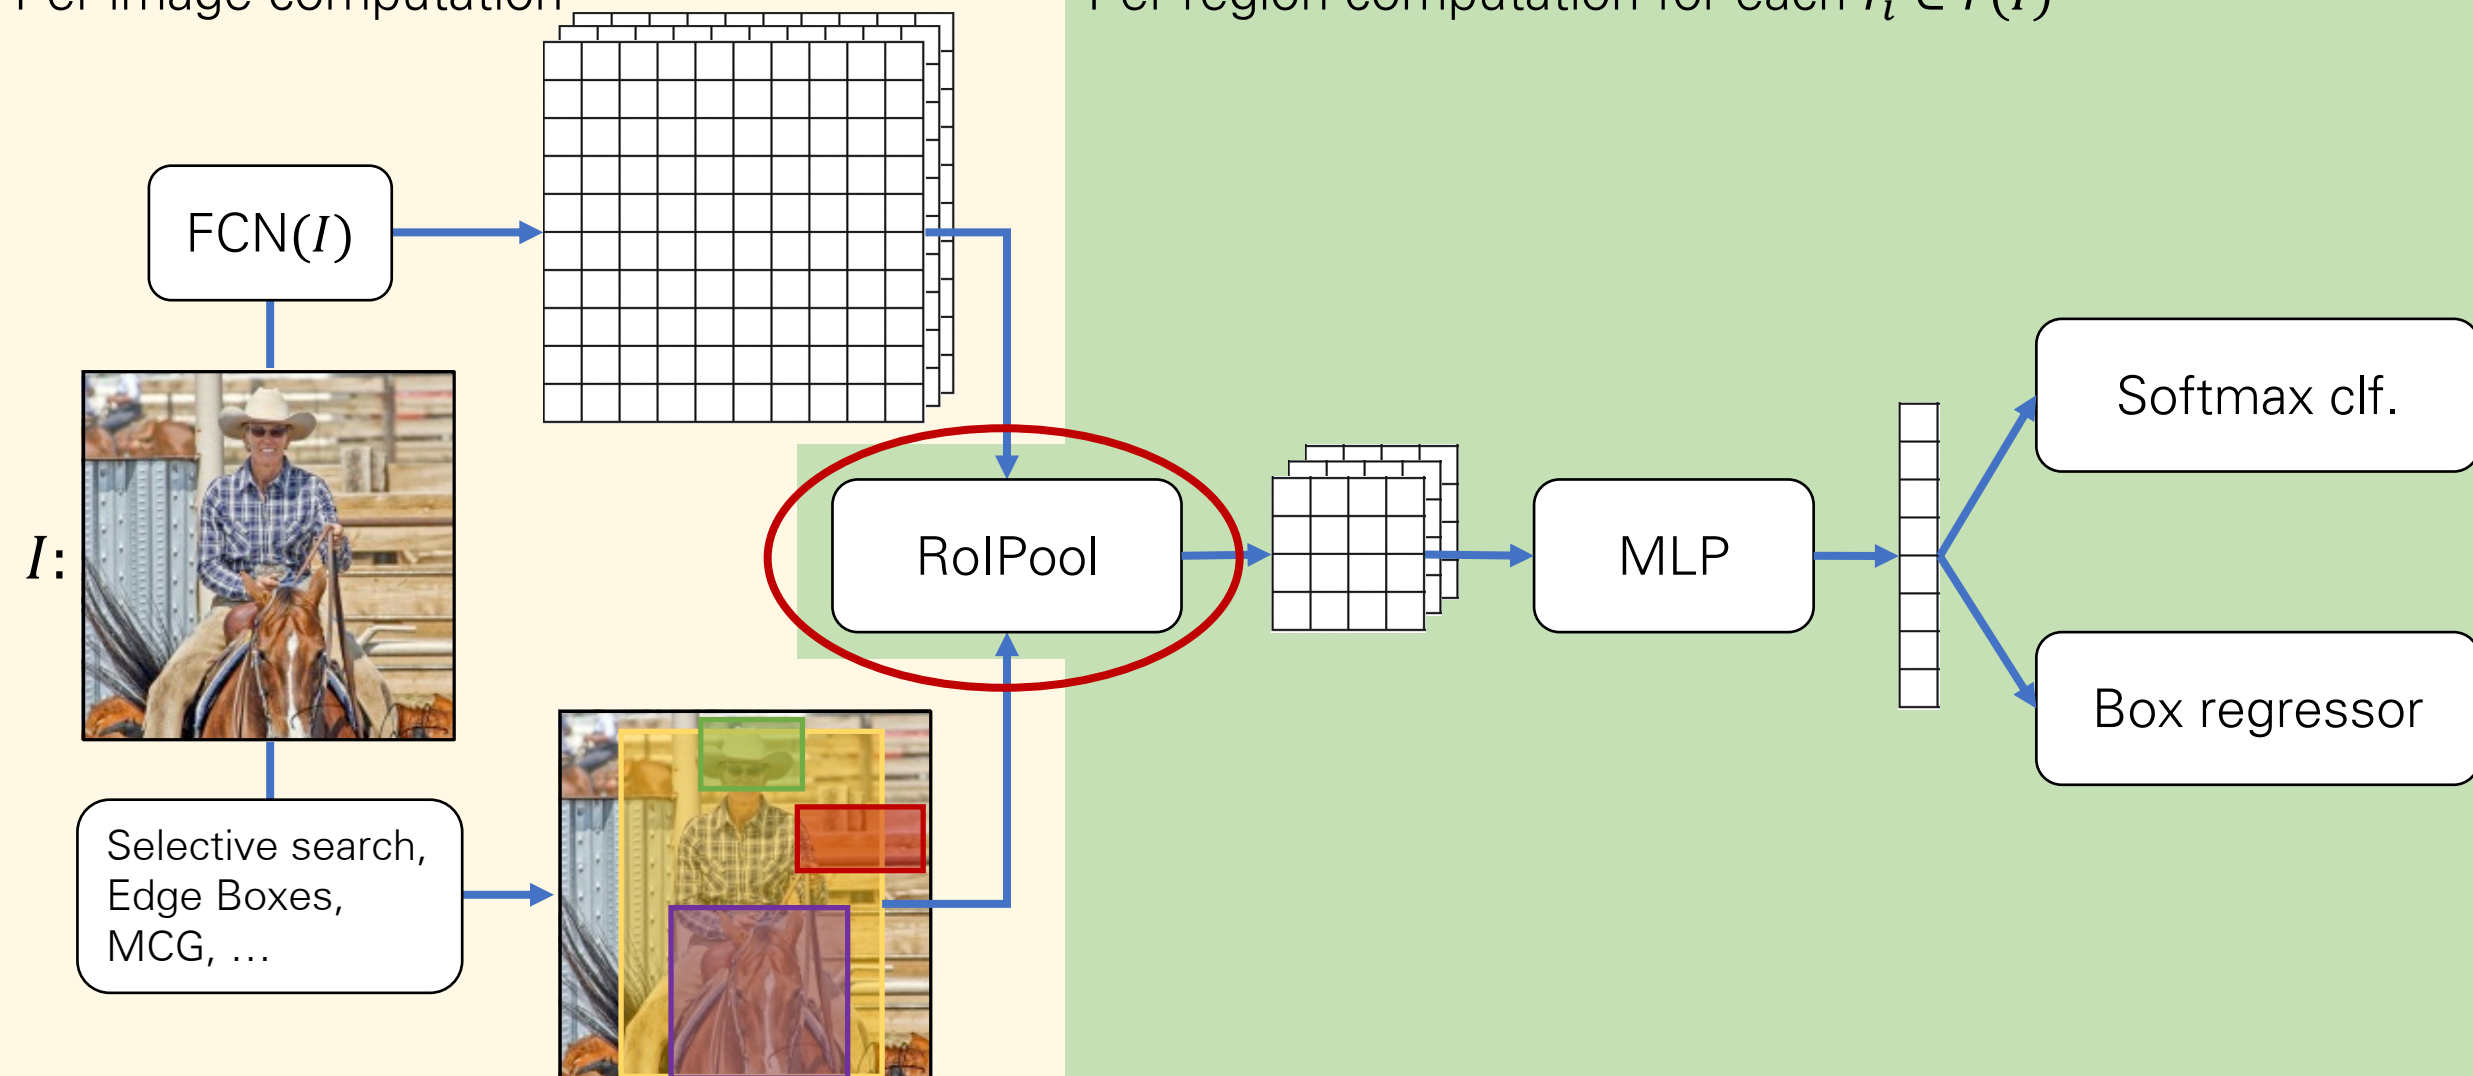


Example feature map dimensions:
(512, H/16, W/16)

Fast R-CNN

Per-image computation

Per-region computation for each $r_i \in r(I)$




RoIPool (on each Proposal)


Transform **arbitrary size proposal** into a **fixed-dimensional** representation (e.g., 2x2)



$$f_I = \text{FCN}(I)$$



Proposal
Region of
Interest
(RoI)




RoIPool (on each Proposal)


Transform **arbitrary size proposal** into a **fixed-dimensional** representation (e.g., 2x2)



$$f_I = \text{FCN}(I)$$



Proposal
Region of
Interest
(RoI)




RoIPool (on each Proposal)

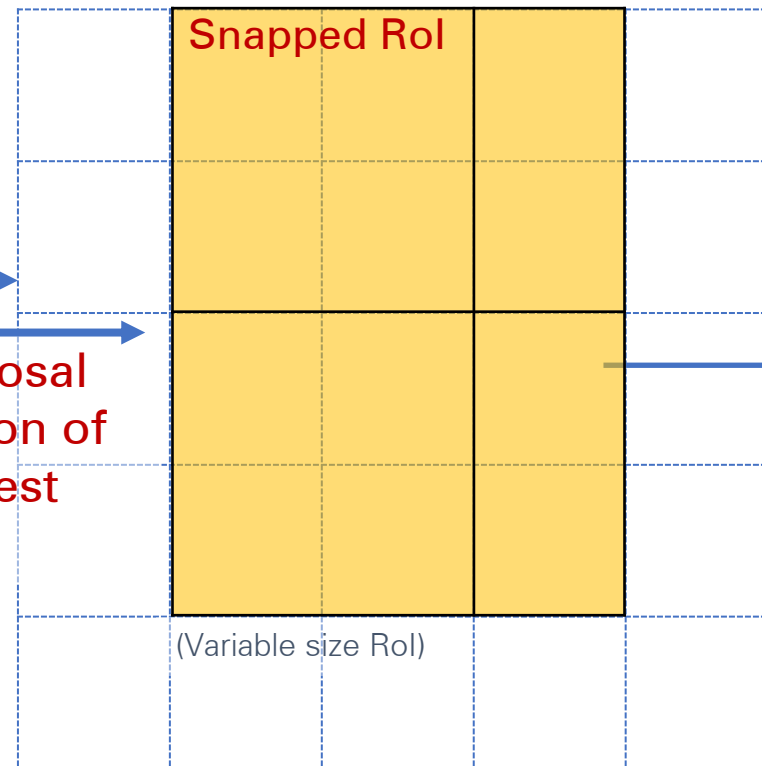
Transform **arbitrary size proposal** into a **fixed-dimensional** representation (e.g., 2x2)



$$f_I = \text{FCN}(I)$$




Proposal
Region of
Interest
(RoI)



RoIPool
transform

(Fixed dimensional
representation)

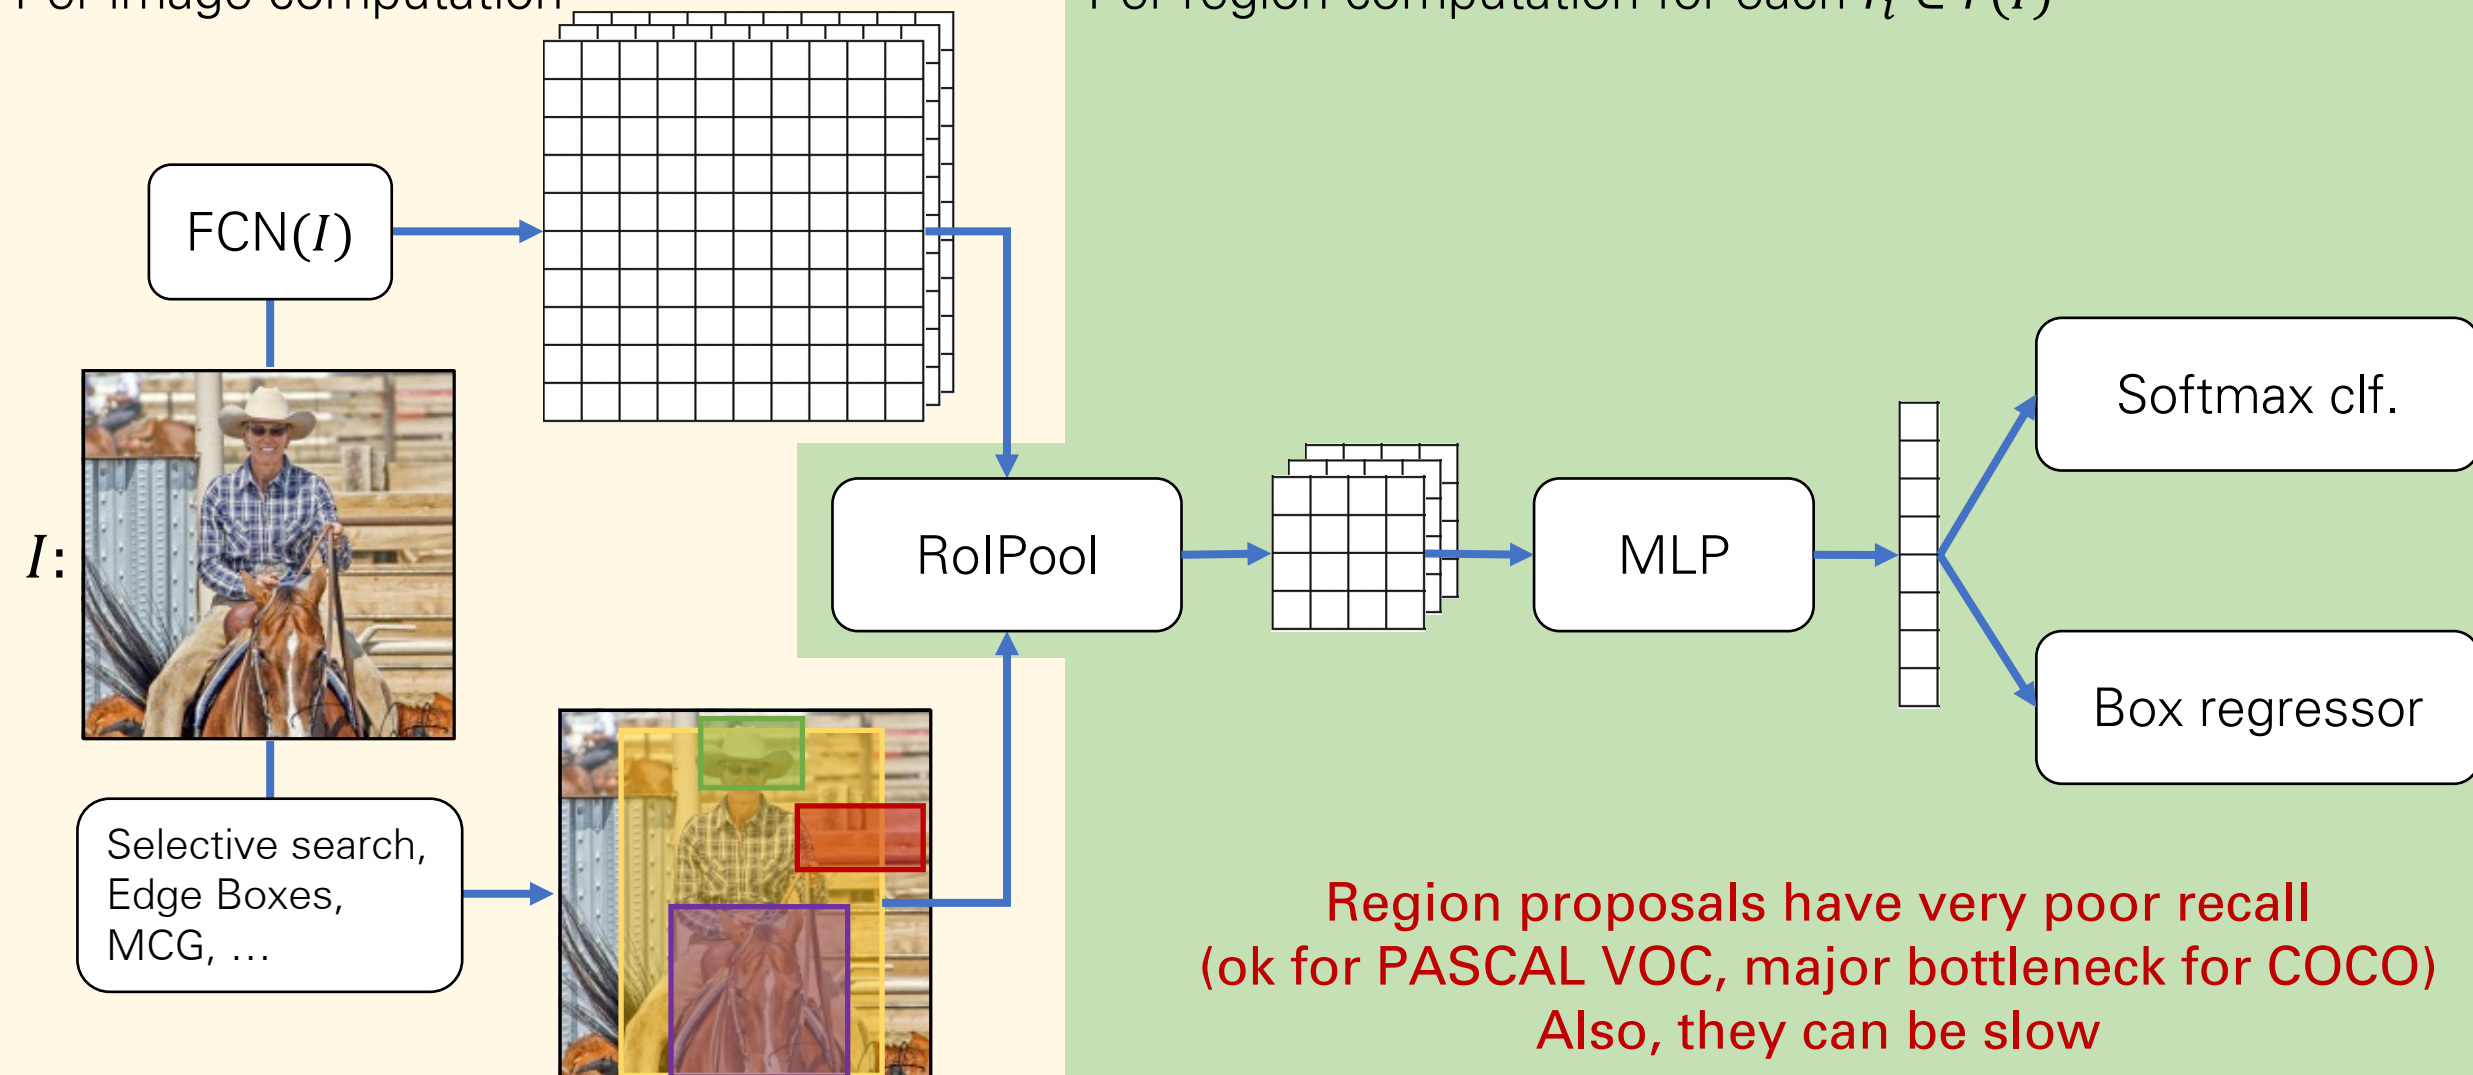


Feature value
is **max** over input
cells

Fast R-CNN

Per-image computation

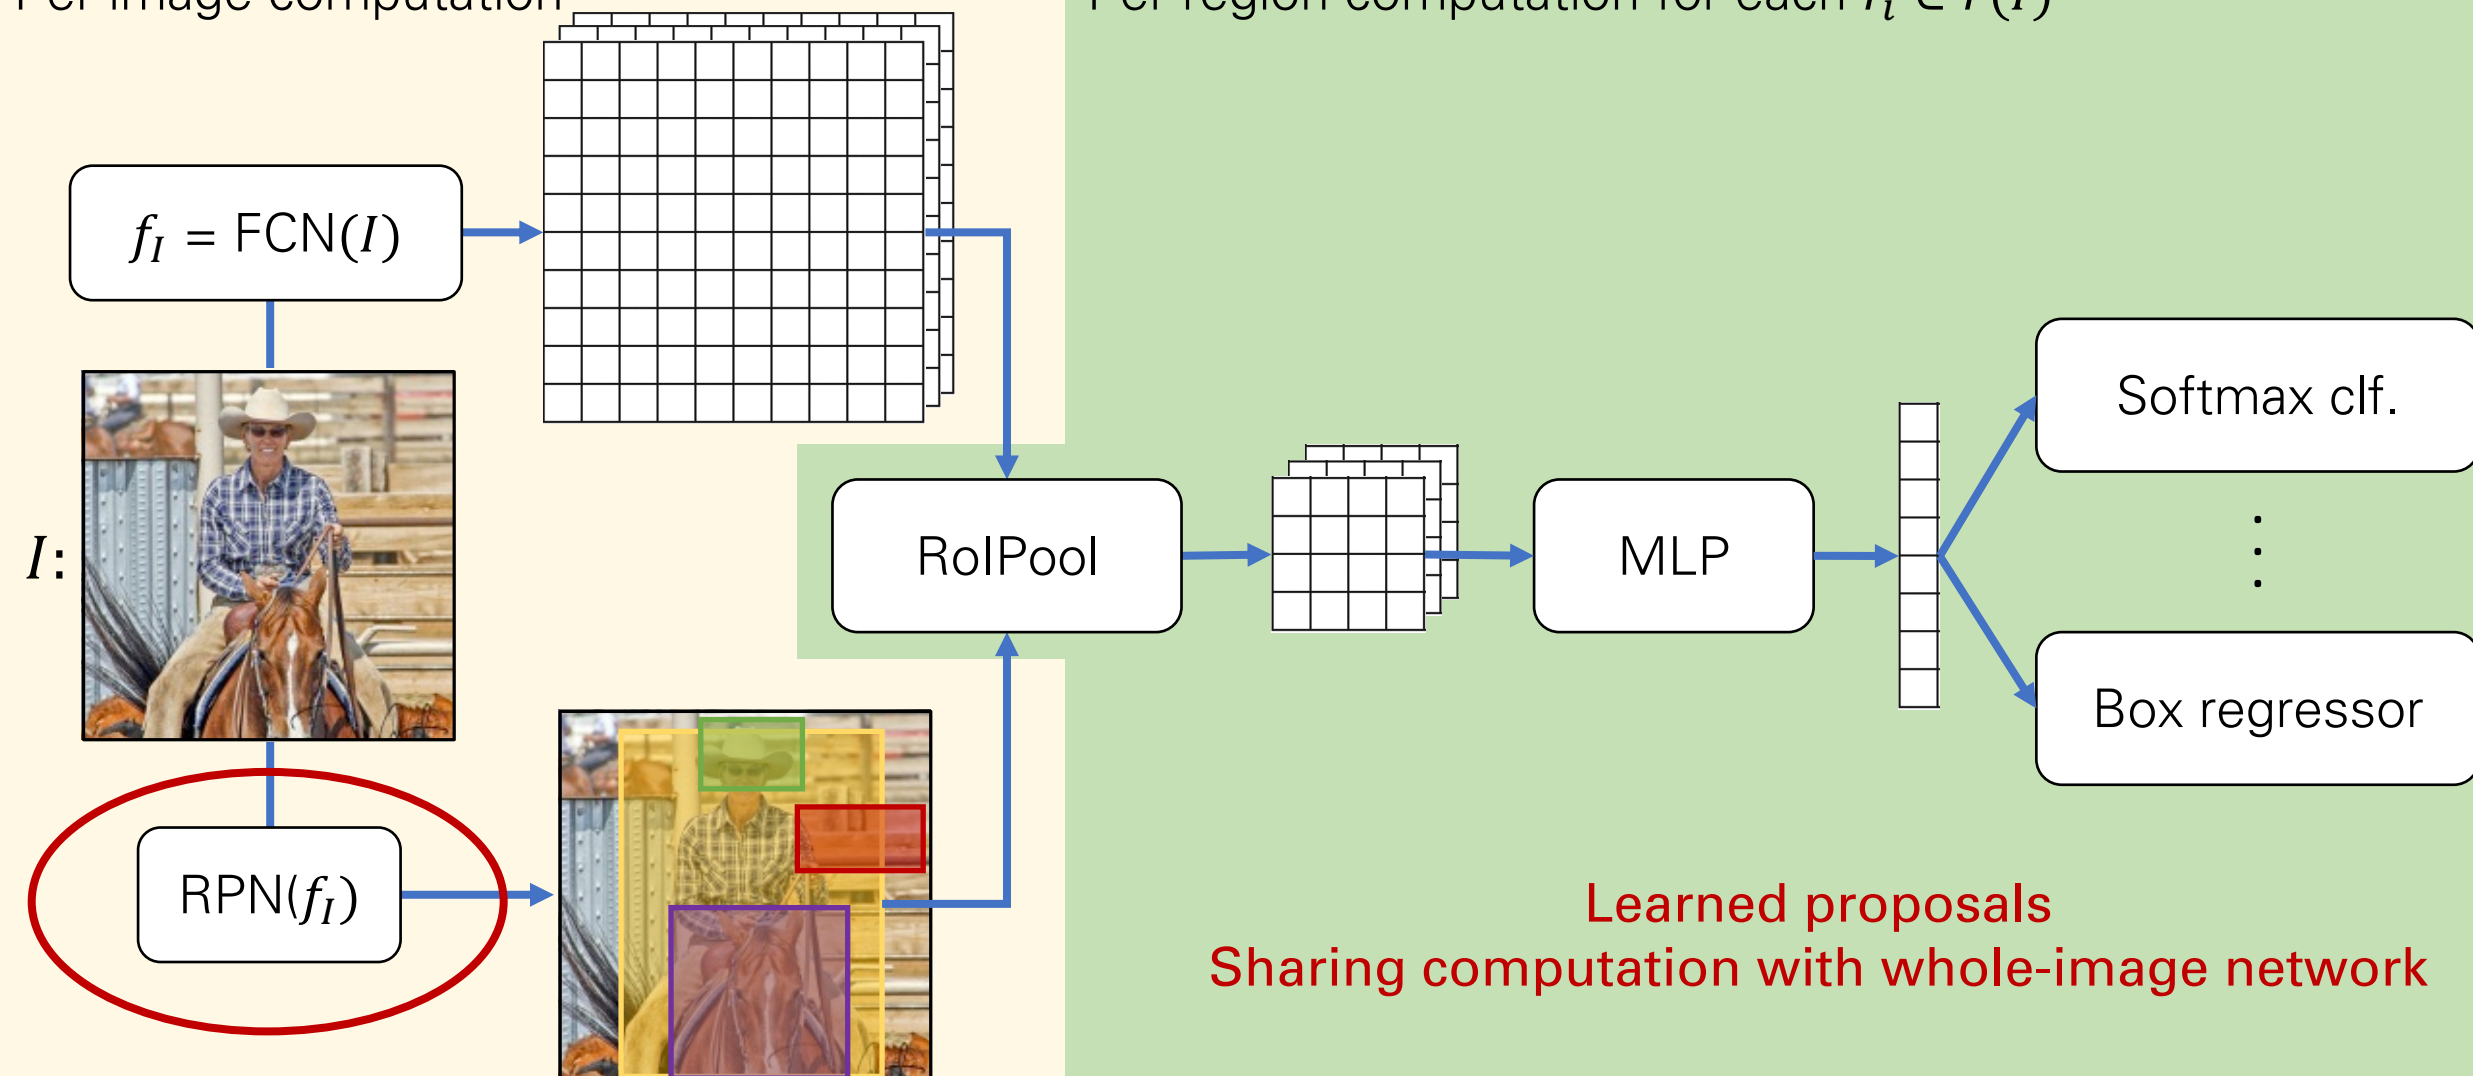
Per-region computation for each $r_i \in r(I)$



Faster R-CNN

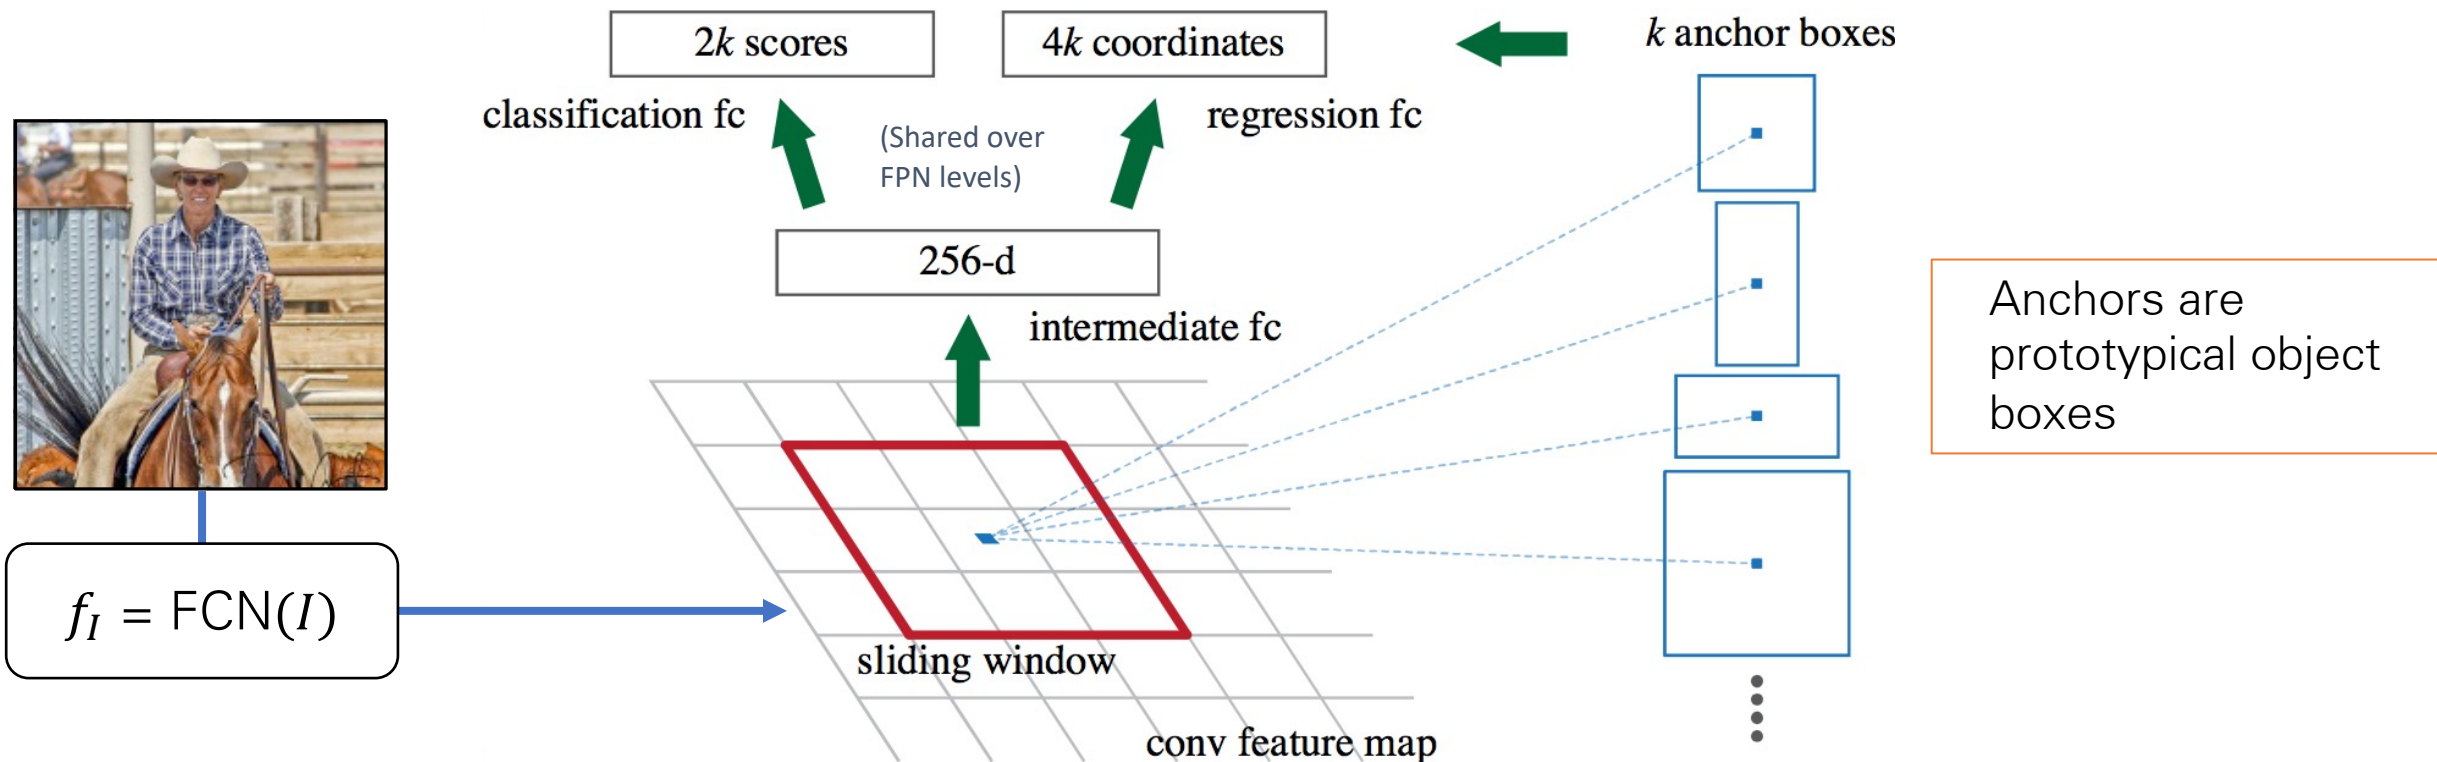
Per-image computation

Per-region computation for each $r_i \in r(I)$



Region Proposal Network (RPN)

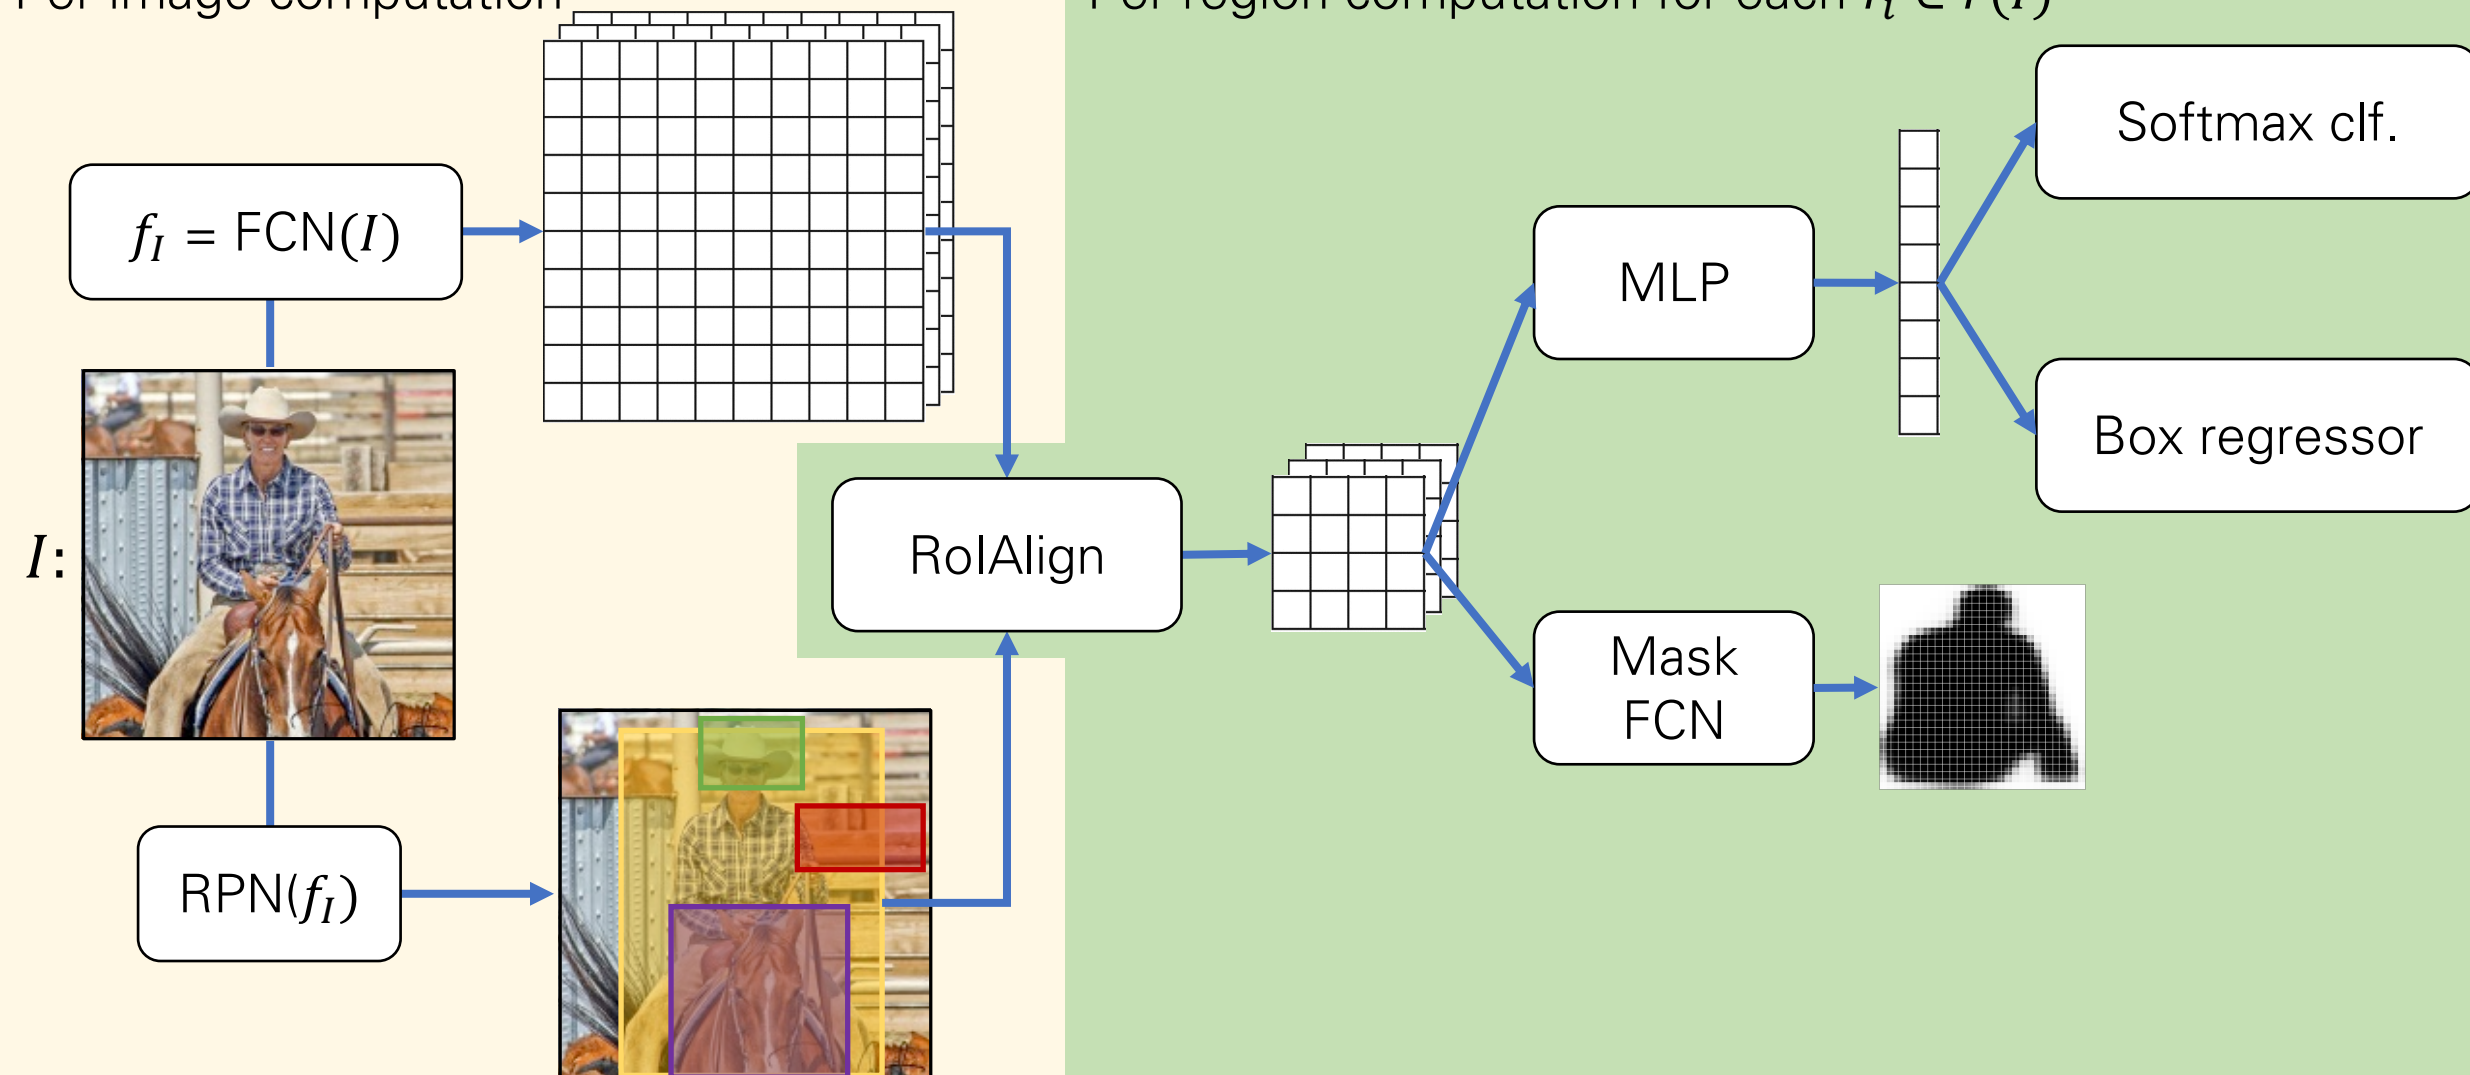
Proposals = sliding window object/not-object classifier + box regression
inside the same network



Mask R-CNN

Per-image computation

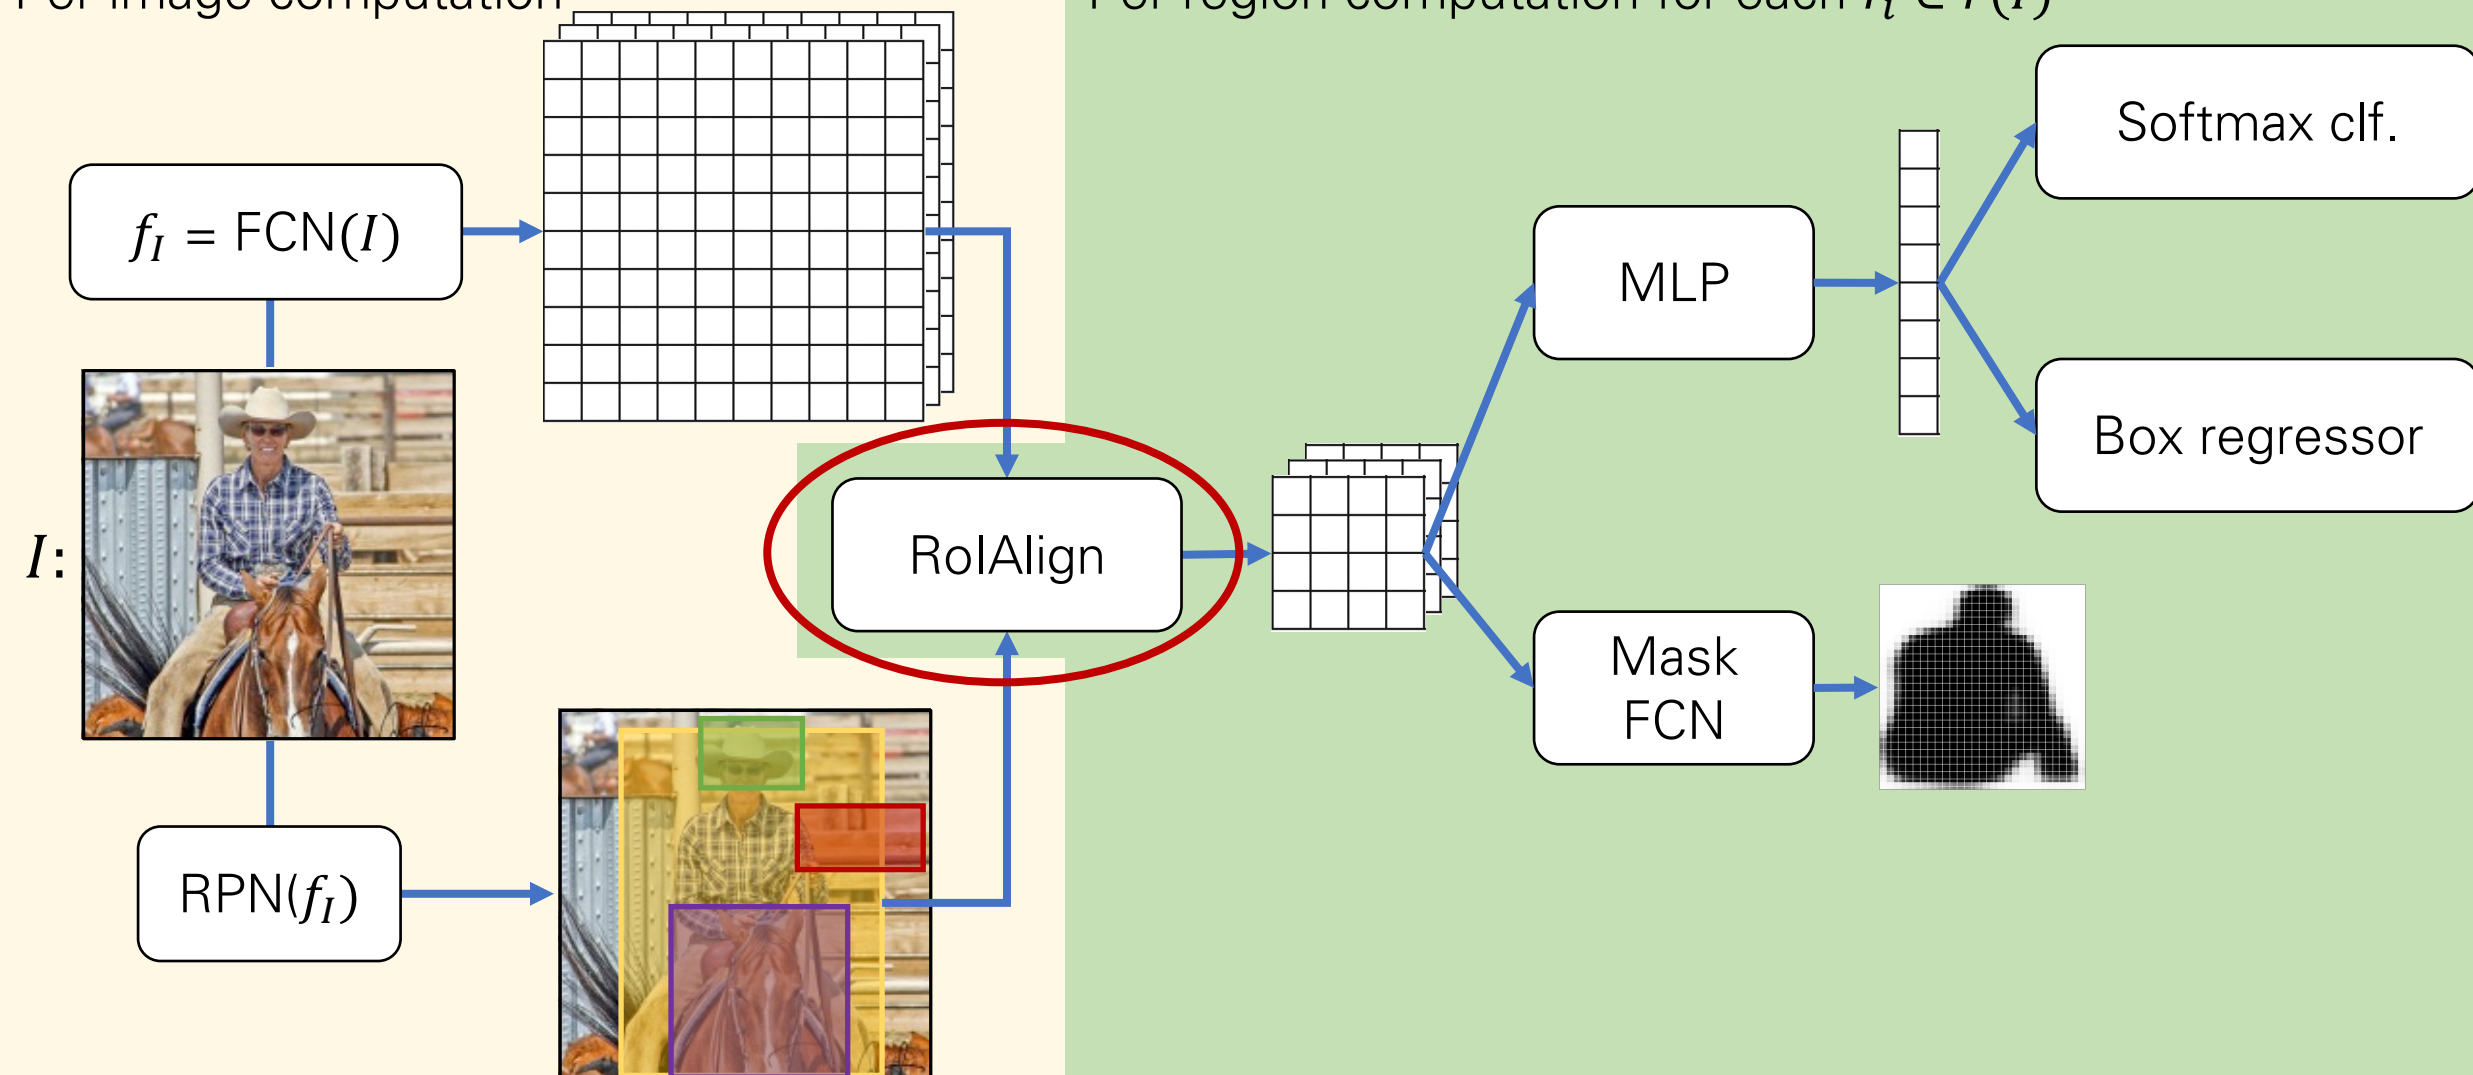
Per-region computation for each $r_i \in r(I)$



Mask R-CNN

Per-image computation

Per-region computation for each $r_i \in r(I)$




RoIAlign (on each Proposal)


Smoothly transform RoI features into a fixed-dimensional representation (e.g., 2x2)



$$f_I = \text{FCN}(I)$$



Proposal
RoI from
RPN




Grid of bilinear
interpolation points

RoIAlign (on each Proposal)


Smoothly transform RoI features into a fixed-dimensional representation (e.g., 2x2)



$$f_I = \text{FCN}(I)$$




Proposal
RoI from
RPN



Grid of bilinear
interpolation points

RoIAlign
transform


(Fixed dimensional
representation)



Feature value is average of
interpolated values on grid

Compare to RoIPool

Preserve alignment or not?


	 align?	bilinear?	agg.	AP	AP ₅₀	AP ₇₅
<i>RoIPool</i> [12]			max	26.9	48.8	26.4
<i>RoIWarp</i> [10]		✓	max	27.2	49.2	27.1
		✓	ave	27.1	48.9	27.1
<i>RoIAlign</i>	✓	✓	max	30.2	51.0	31.8
	✓	✓	ave	30.3	51.2	31.5

+20% relative
at high IoU

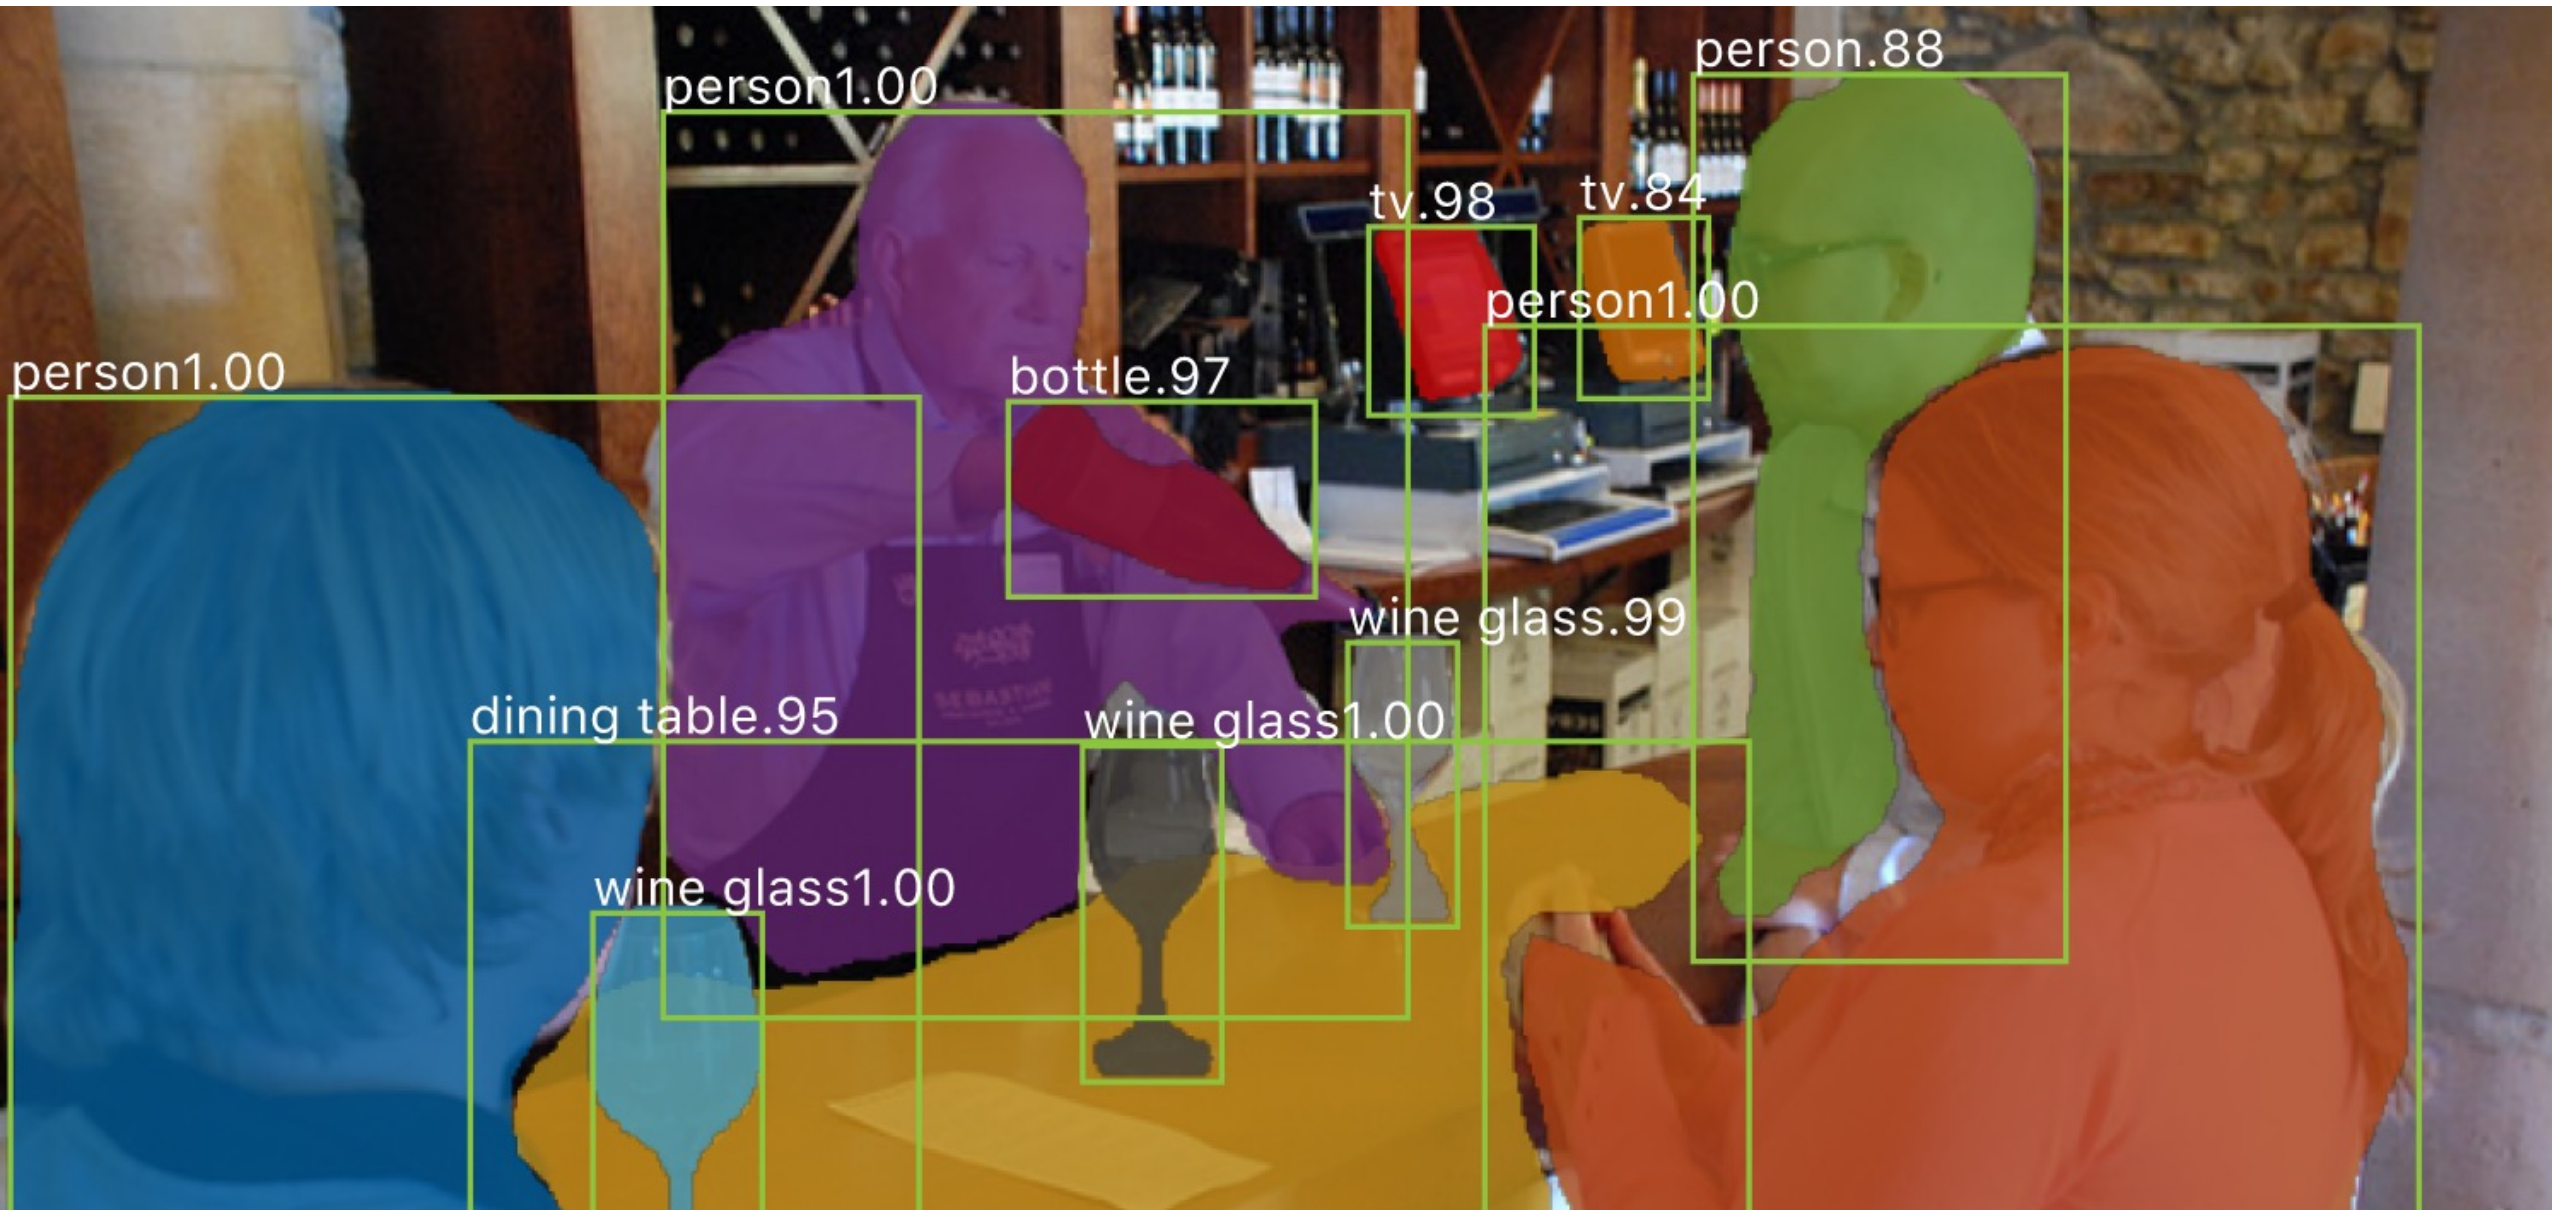
(c) **RoIAlign** (ResNet-50-C4): Mask results with various RoI layers. Our RoIAlign layer improves AP by ~ 3 points and AP₇₅ by ~ 5 points. Using proper alignment is the only factor that contributes to the large gap between RoI layers.

Compare to RoIPool

Quantization breaks pixel-to-pixel alignment




Instance Segmentation



Mask R-CNN

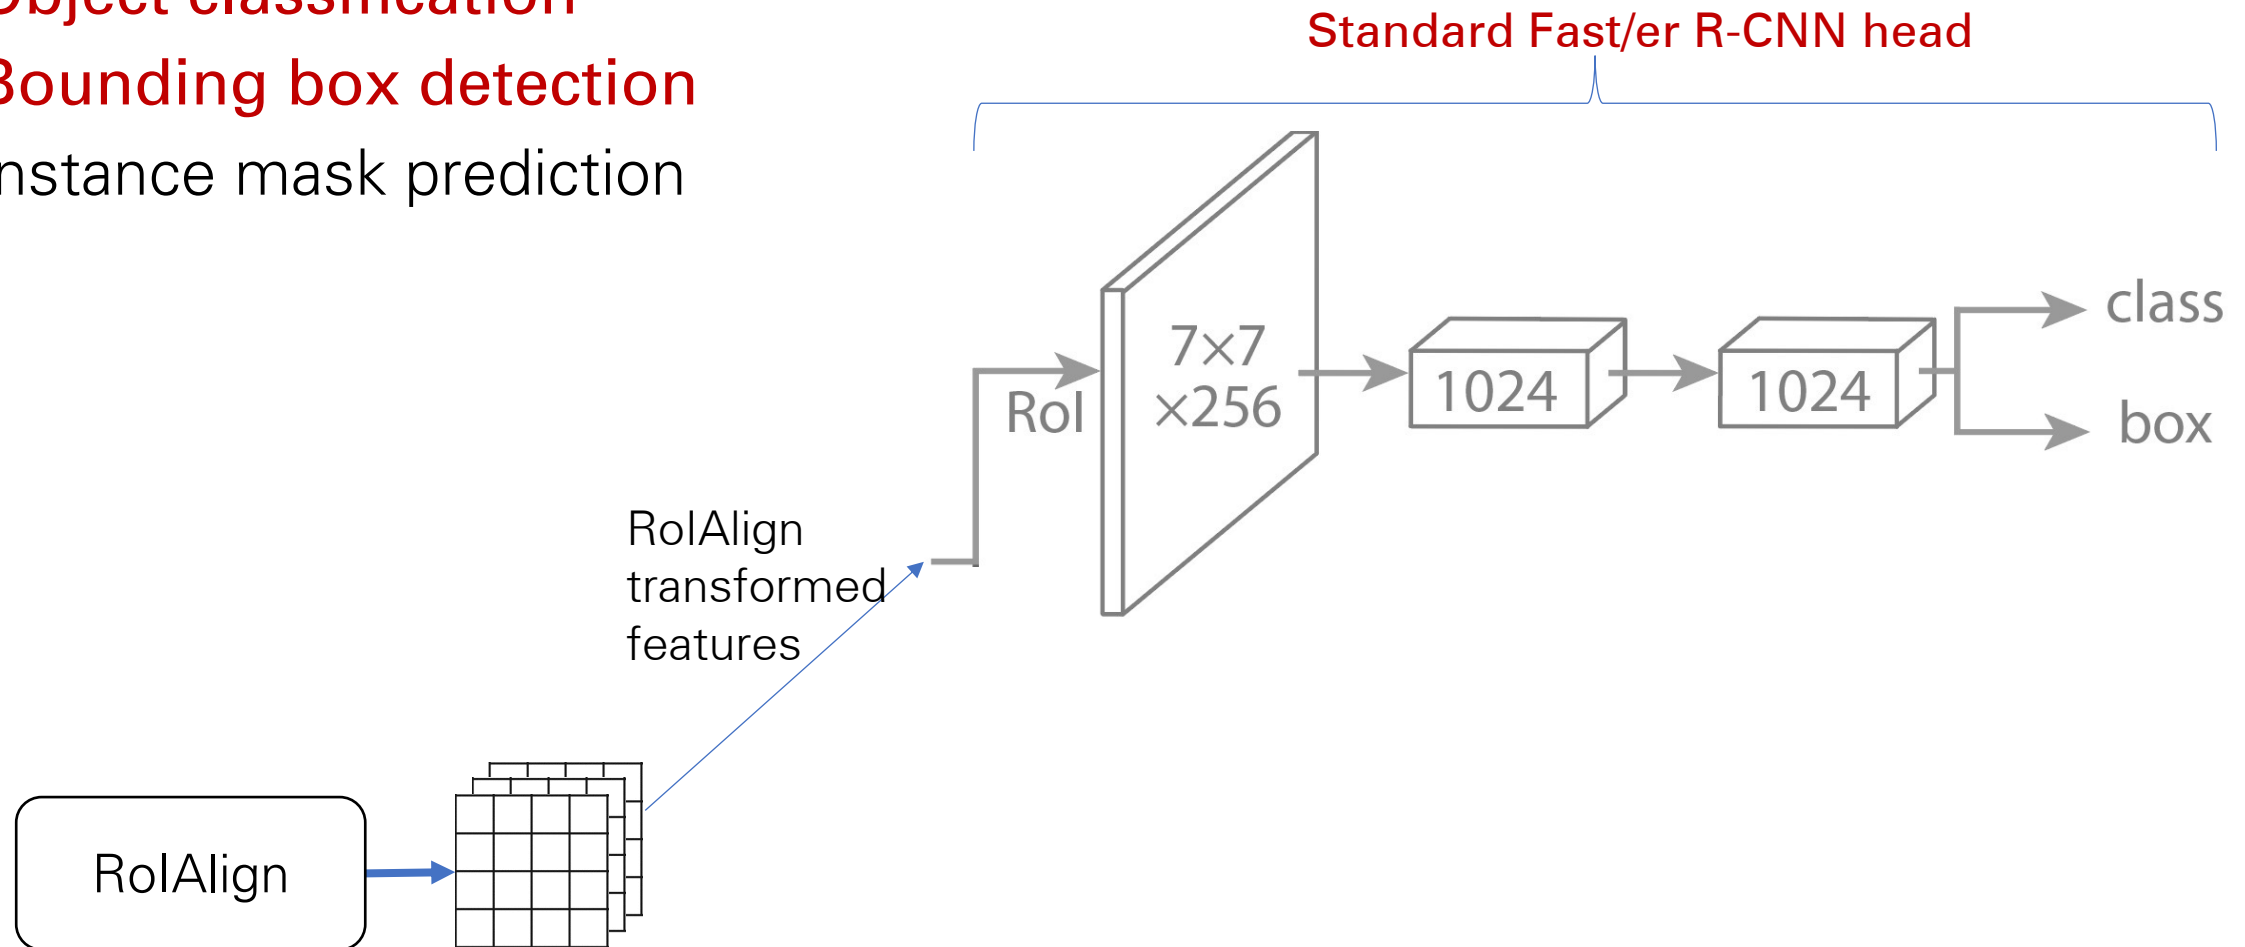
Per-image computation

Per-region computation for each $r_i \in r(I)$



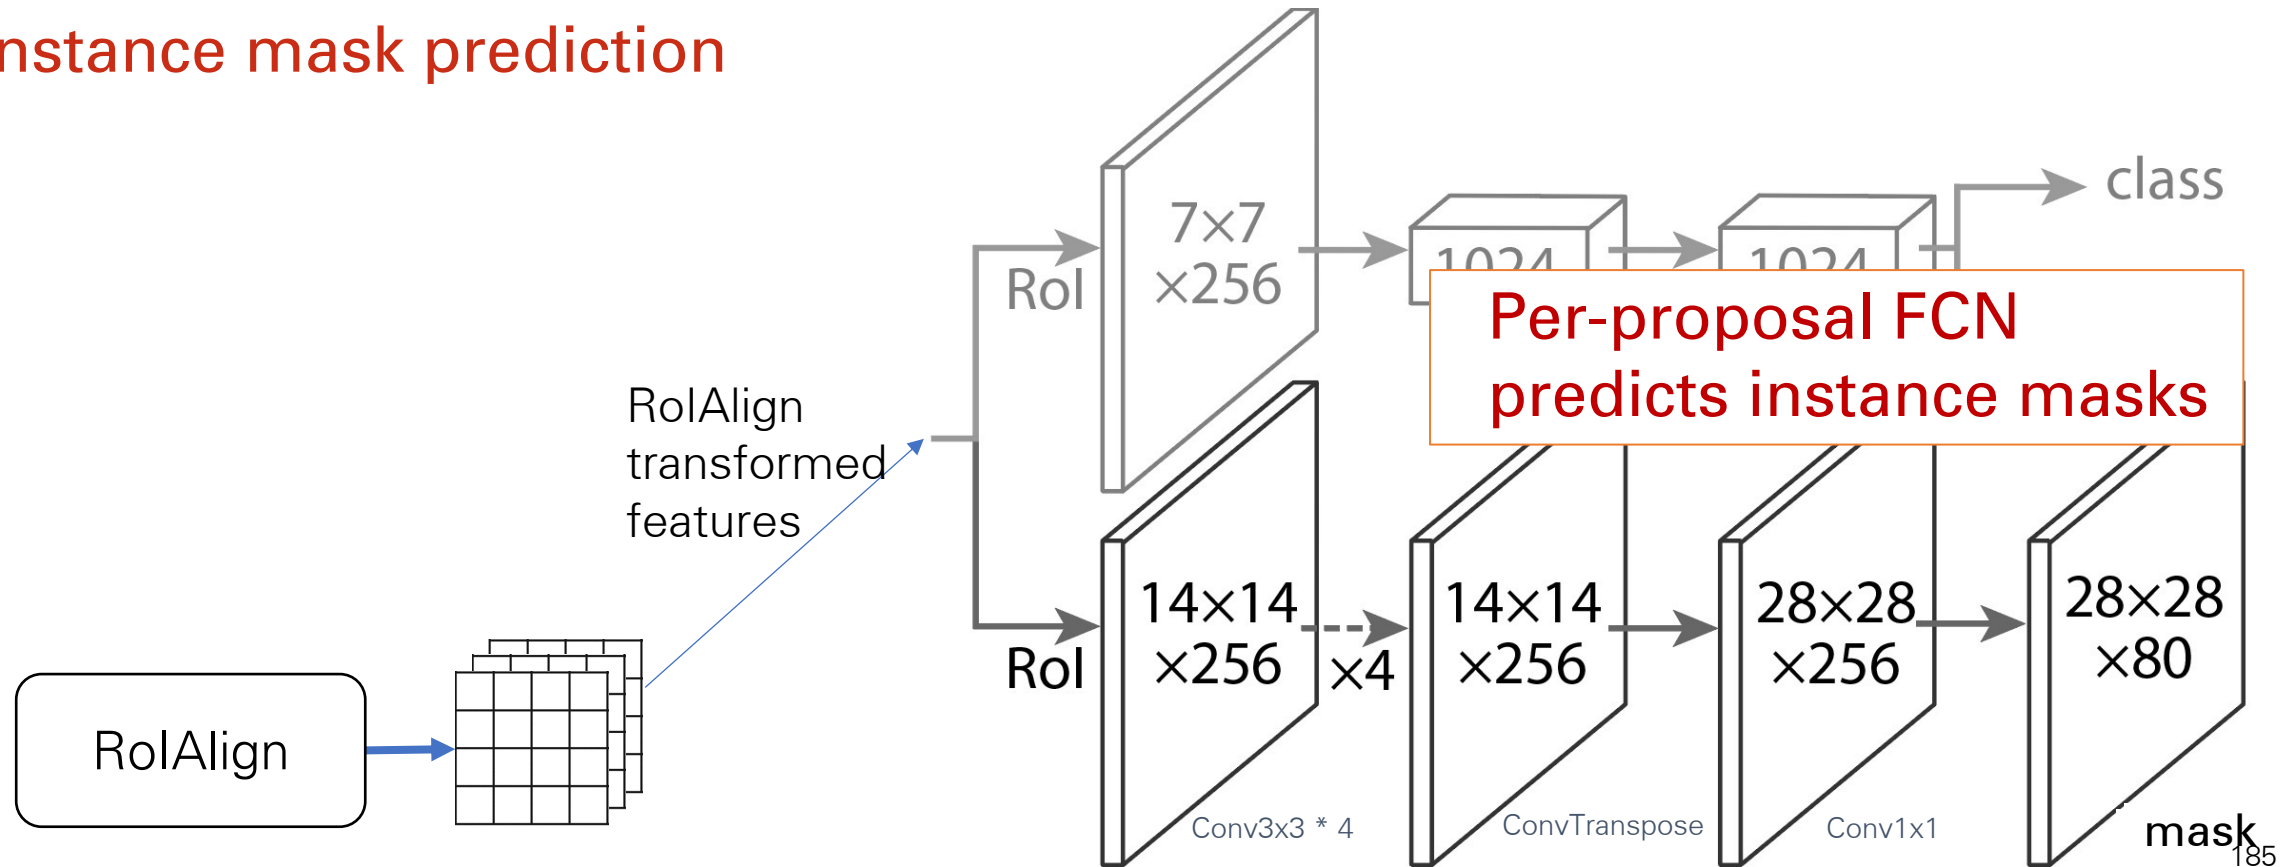
Mask Head (on each Proposal)

- Task specific heads for ...
 - Object classification
 - Bounding box detection
 - Instance mask prediction



Mask Head (on each Proposal)


- Task specific heads for ...
 - Object classification
 - Bounding box detection
 - **Instance mask prediction**



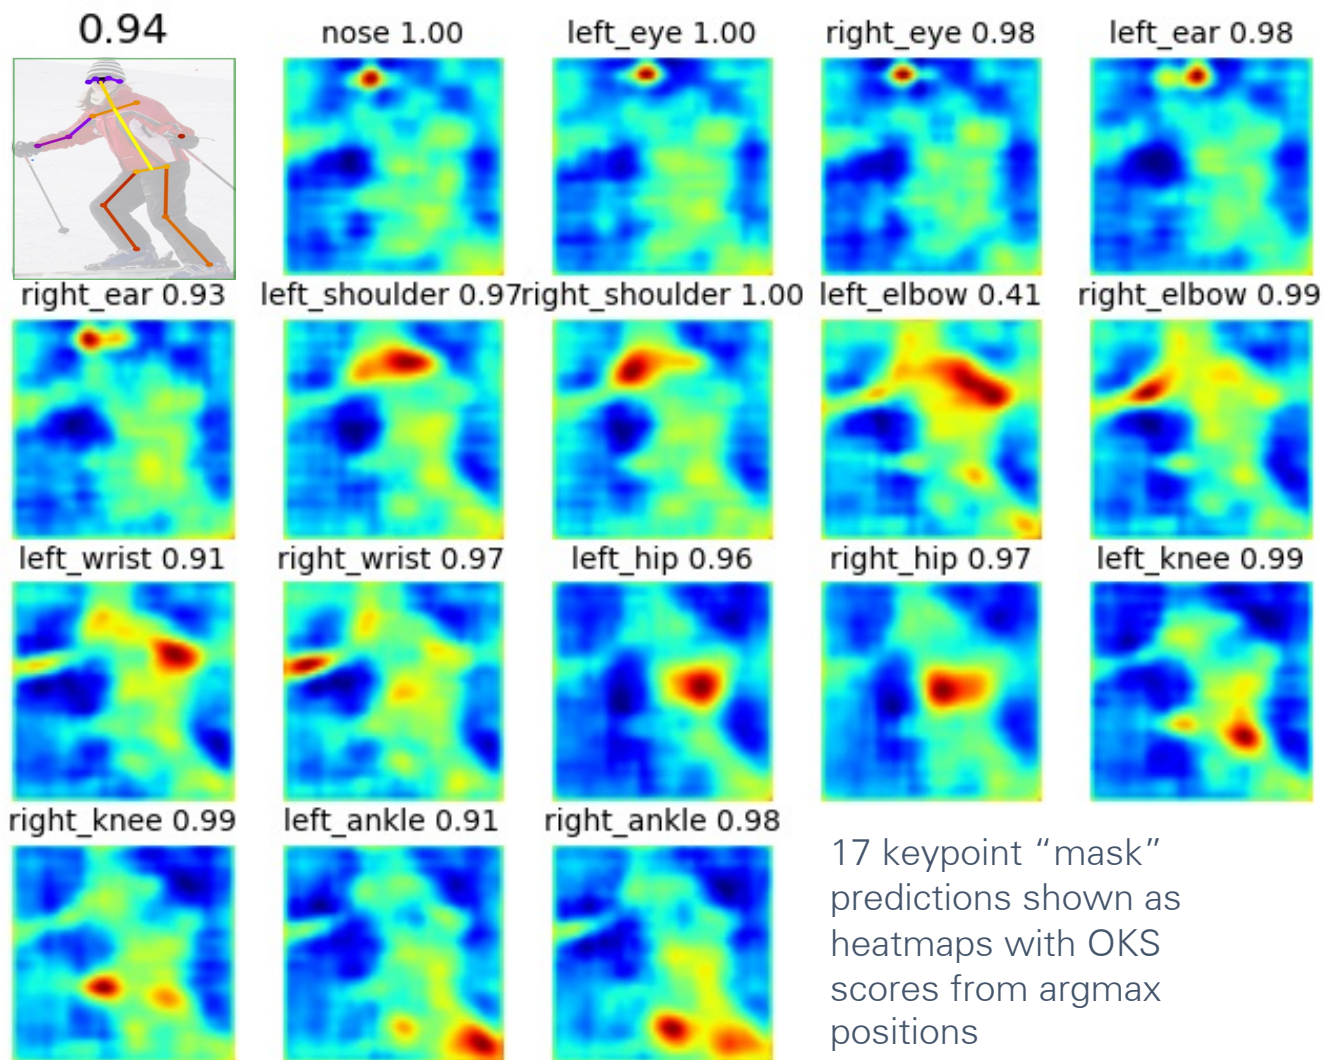

Mask R-CNN: Extension to 2D Human Pose

Per-image computation

Per-region computation for each $r_i \in r(I)$



Pose Head




- Add keypoint head (28x28x17)
- Predict one "mask" for each keypoint
- Softmax over **spatial locations** (encodes one keypoint per mask "prior")

Mask R-CNN: Training

- Same as “image centric” Fast/er R-CNN training
- But with **training targets for masks**

Example Mask Training Targets

Image with training proposal



28x28 mask target









Image with training proposal



28x28 mask target



Mask R-CNN: Inference

1. Perform Faster R-CNN inference


- Run backbone FCN
- Generate proposals with RPN
- Score the proposals with clf. head
- Refine proposals with box regressor
- Apply NMS and take the top K (= 100, e.g.)

2. Run RoIAlign and mask head on top- K refined, post-NMS boxes

- Fast (only compute masks for top- K detections)
- Improves accuracy (uses refined detection boxes, not proposals)

Mask Prediction

28x28 soft prediction from Mask R-CNN
(enlarged)



Soft prediction **resampled to image coordinates**
(bilinear and bicubic interpolation work equally well)




Final prediction (threshold at 0.5)



Validation image with box detection shown in red

Mask Prediction


28x28 soft prediction from Mask R-CNN
(enlarged)




Soft prediction **resampled to image coordinates**
(bilinear and bicubic interpolation work equally well)




Final prediction (threshold at 0.5)




Validation image with box detection shown in red



Quantization breaks pixel-to-pixel alignment




RoiPool coordinate quantization



Mask Prediction

28x28 soft prediction



Resized soft prediction



Final mask




Validation image with box detection shown in red

Mask Prediction



28x28 soft prediction




Resized Soft prediction




Final mask



Validation image with box detection shown in red



person1.00 person1.00 person1.00 person.98
person1.00 person.91 surfboard1.00 surfboard.98 surfboard1.00 surfboard1.00
surfboard1.00 person.74



person1.00

person.88

person1.00

tv.98

tv.84

person1.00


bottle.97

wine glass.99

dining table.95


wine glass1.00

wine glass1.00



Is Object Detection Solved?

- Obviously no; there are **frequently silly errors**
- But it is getting frustratingly good
- The errors are often reasonable
- The bottlenecks are raw recognition and “reasoning”





tv 1.00

person 1.00

bottle 0.92

bottle 0.93

sink 0.99

sink 0.99

sink 0.96

bed 0.98




bed 1.00

bed 0.98

bed 0.90

bed 0.94





person 1.00

tie 0.96






book 0.

person 0.95

vase 0.97


orange 0.96

orange 0.97




Addressing other tasks...

Addressing other tasks...




A block of compute with a few million parameters.

Addressing other tasks...



A block of compute with a few million parameters.


Addressing other tasks...





A block of compute with a few million parameters.

Image Classification



thing = a vector of probabilities for different classes



Segmentation




224x224x3





7x7x512

deconv layers




224x224x20
array of class
probabilities
at each pixel.

Localization




image

224x224x3




CNN



features

7x7x512




Class probabilities
(as before)

4 numbers:


- X coord
- Y coord
- Width
- Height

Image Captioning



image

224x224x3




CNN



features

7x7x512

RNN




A sequence of 10,000-dimensional vectors giving probabilities of different words in the caption.


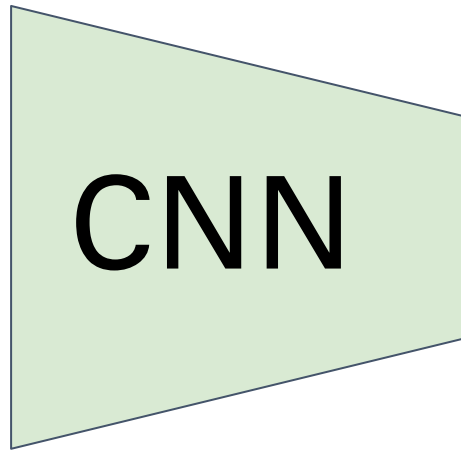
Reinforcement Learning




Mnih et al. 2015



160x210x3

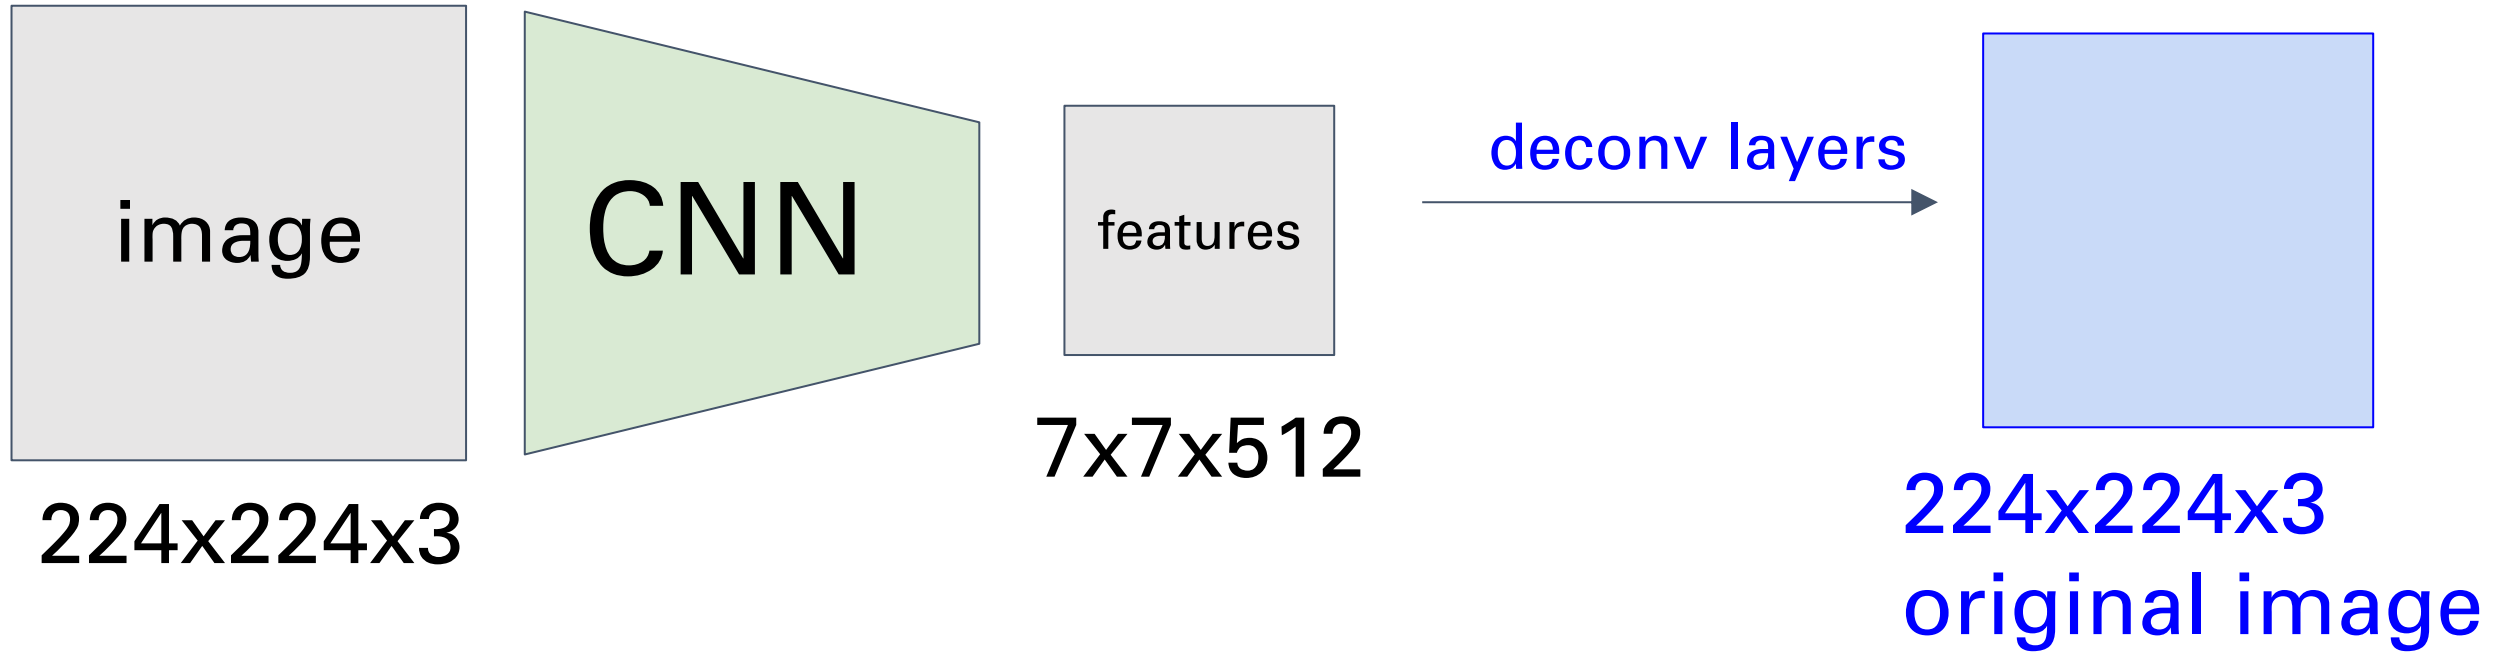


fully connected

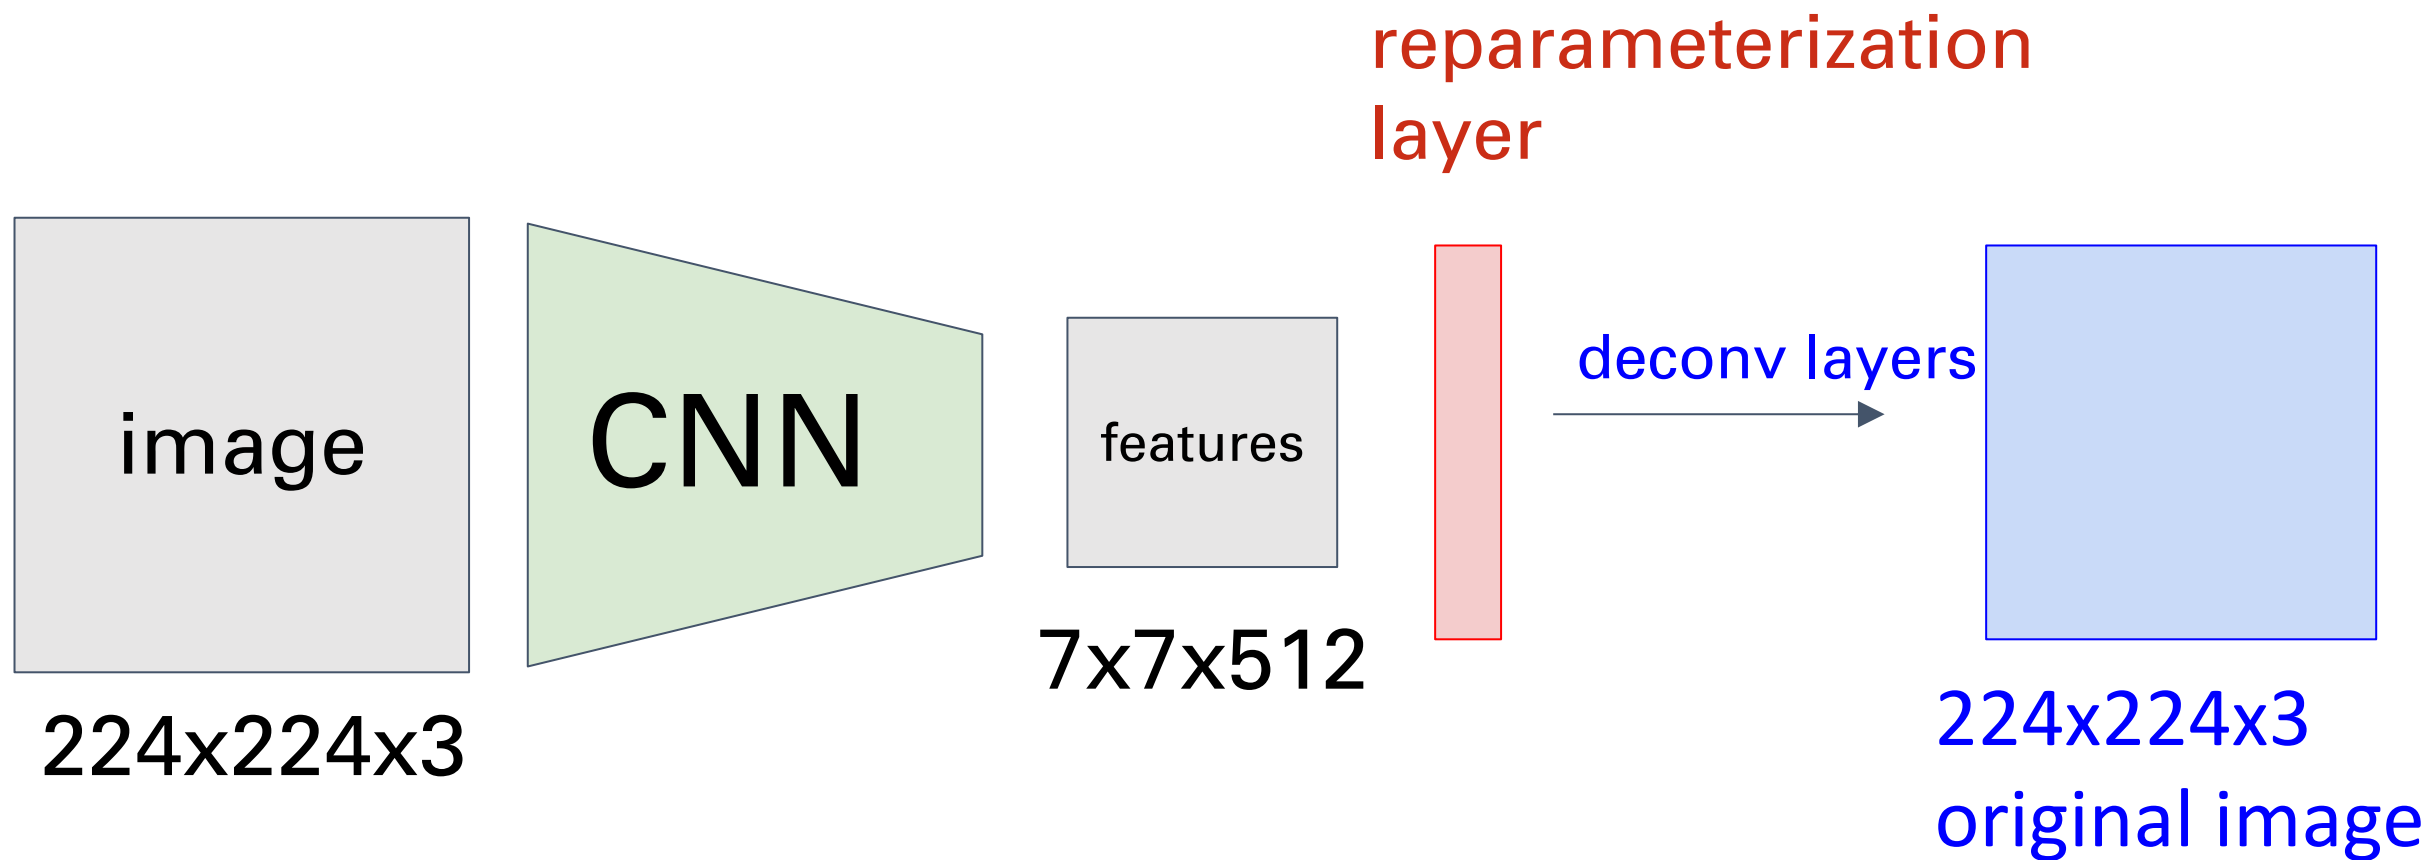


e.g. vector of 8 numbers giving probability of wanting to take any of the 8 possible ATARI actions.

Autoencoders




Variational Autoencoders




[Kingma et al.], [Rezende et al.], [Salimans et al.]

Addressing other tasks...



- 1D convolution \approx Time Delay Neural Networks (Waibel et al. 1989, Collobert and Weston 2011)
- Two main paradigms:
 - **Context window modeling:** For tagging, etc. get the surrounding context before tagging
 - **Sentence modeling:** Do convolution to extract n-grams, pooling to combine over whole sentence

Addressing other tasks...



- CNNs for audio processing: MFCC features + Time Delay Neural Networks

Next lecture:
Understanding and Visualizing
ConvNets