

convolutional neural networks with swift for tensorflow: overview

brettkoonce.com/talks

august 7th, 2020

summary

- **goal: explore image recognition w/ s4tf**
- **neural networks + convolutions**
- **mnist, cifar, vgg, resnet**
- **mobilenet, efficientnet**
- **sota, next steps**

why?

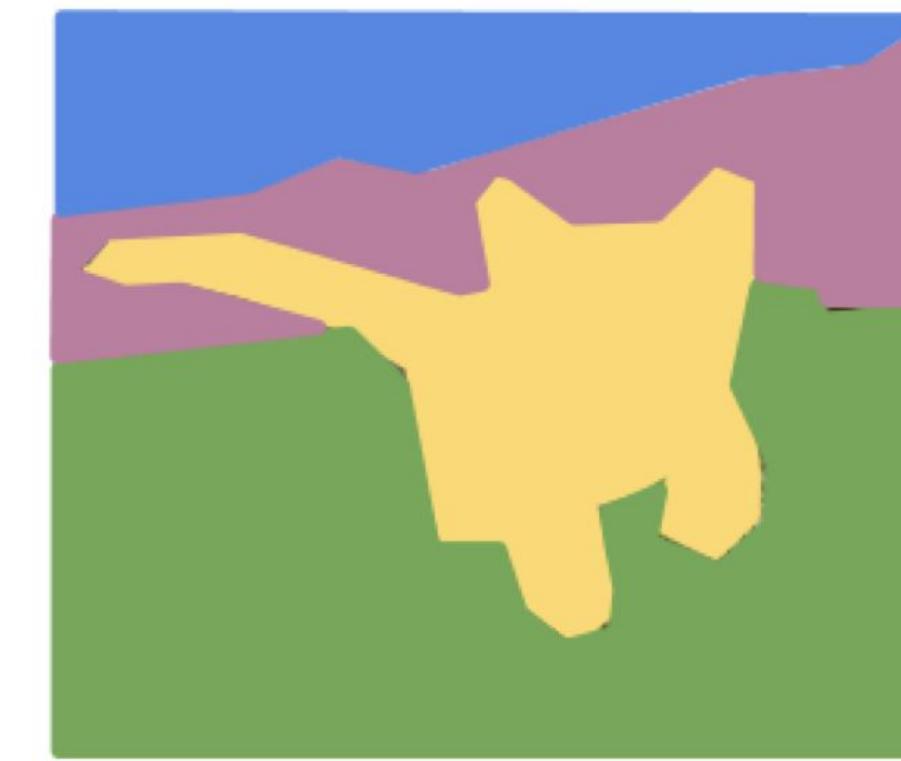
- **me: swift, unix, deep learning**
- **beginner mindset**
- **depth vs breadth**
- **cheap experimentation --> mastery**
- **cnn's are foundational**

computer vision problems

- image recognition
- object detection
- image segmentation
- instance segmentation



CAT



GRASS, CAT,
TREE, SKY

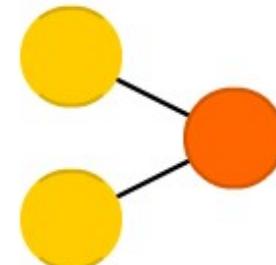


DOG, DOG, CAT

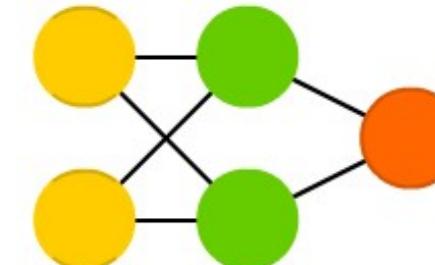
convolutional neural networks

- Input Cell
- Hidden Cell
- Output Cell
- Kernel
- Convolution or Pool

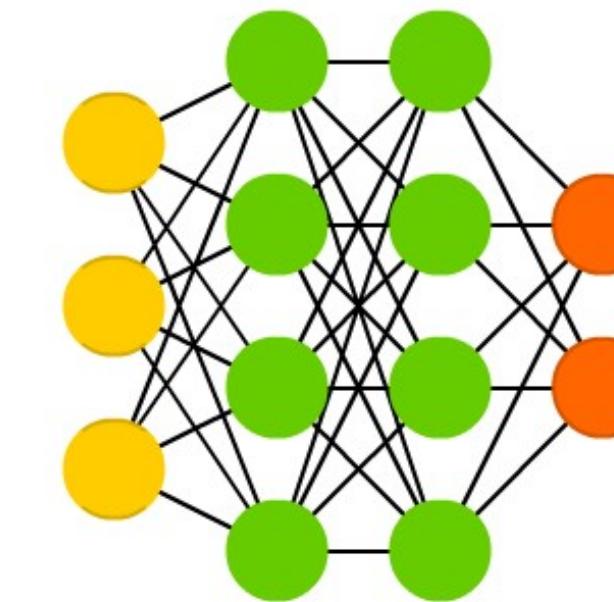
Perceptron (P)



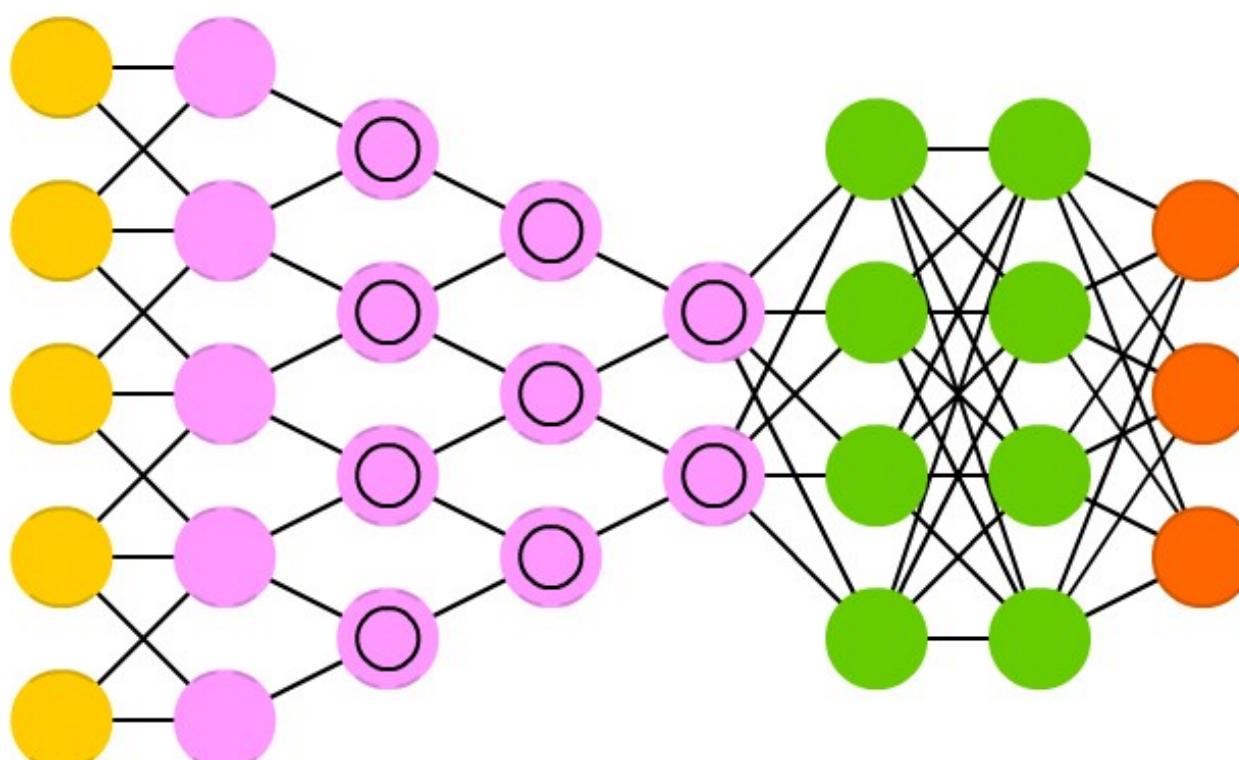
Feed Forward (FF)



Deep Feed Forward (DFF)

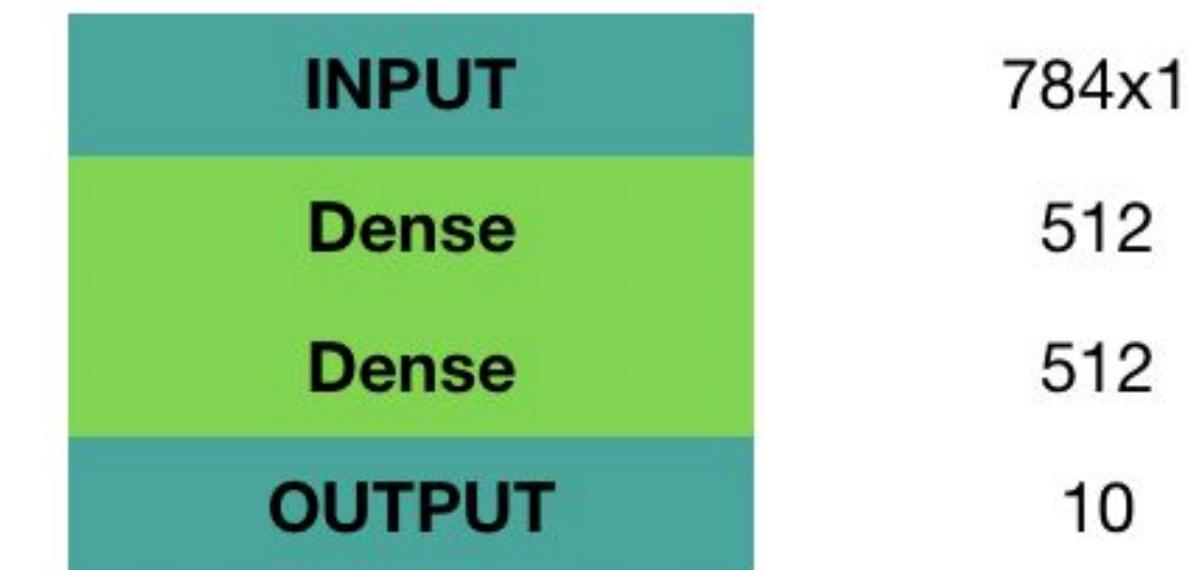


Deep Convolutional Network (DCN)

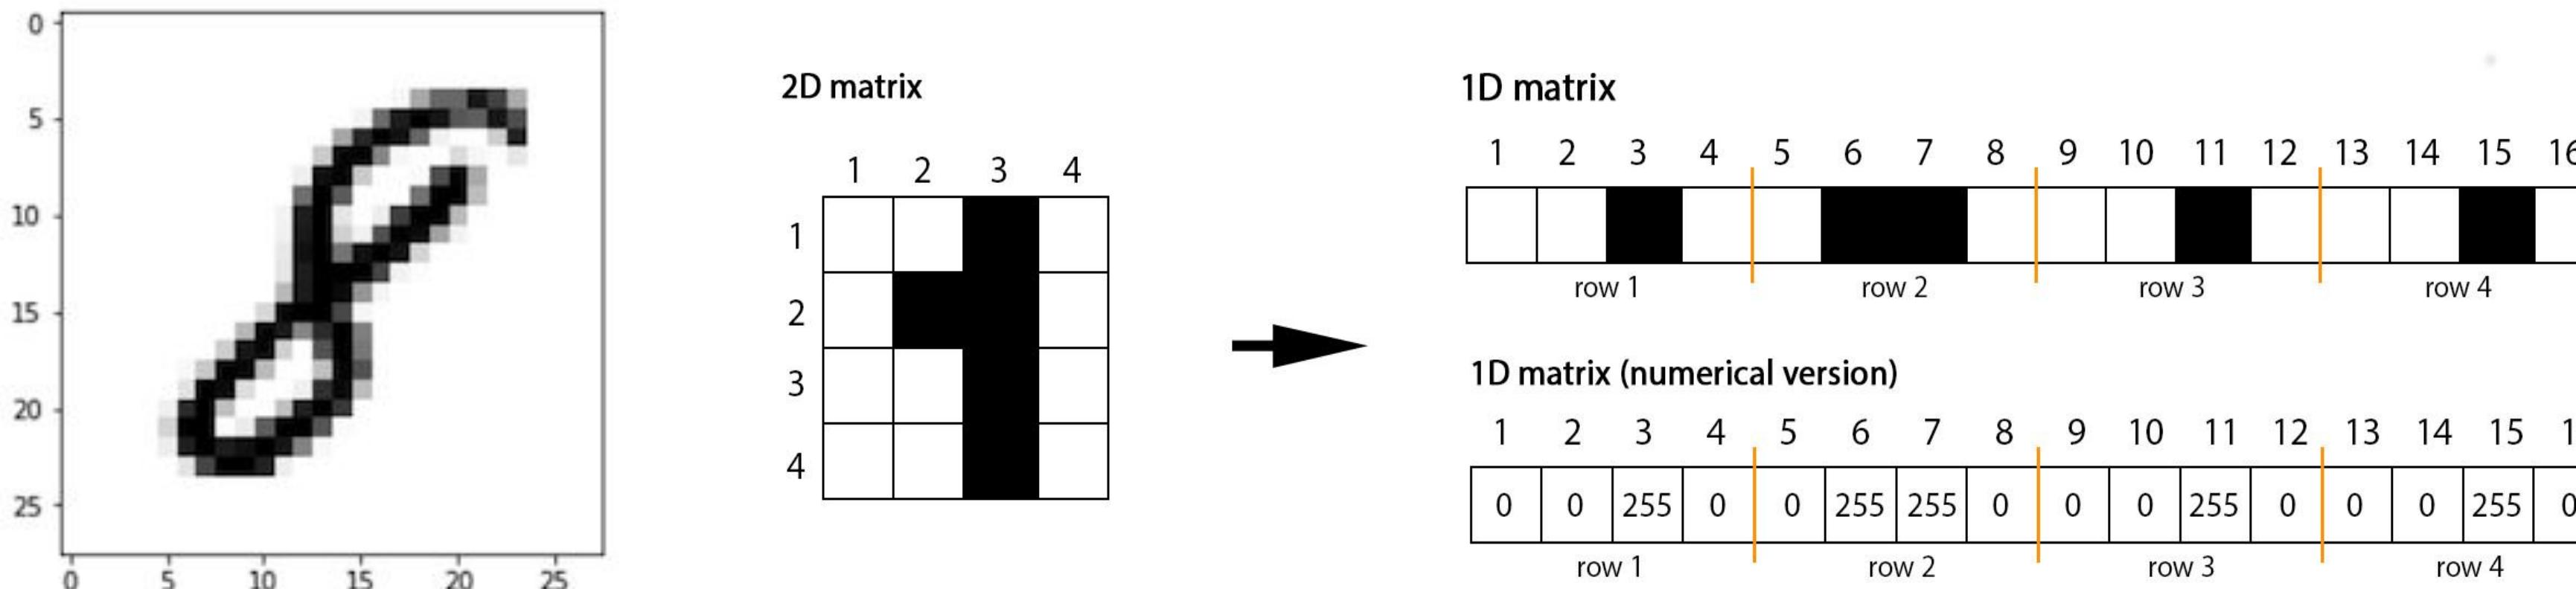


1d mnist

- mnist: $28 \times 28 \times 1$ (h, w, 0..255 greyscale)
- Convert to 1d vector: [row1, row2, row3...row28] == (784×1)
- $2 * 512$ fully connected layers

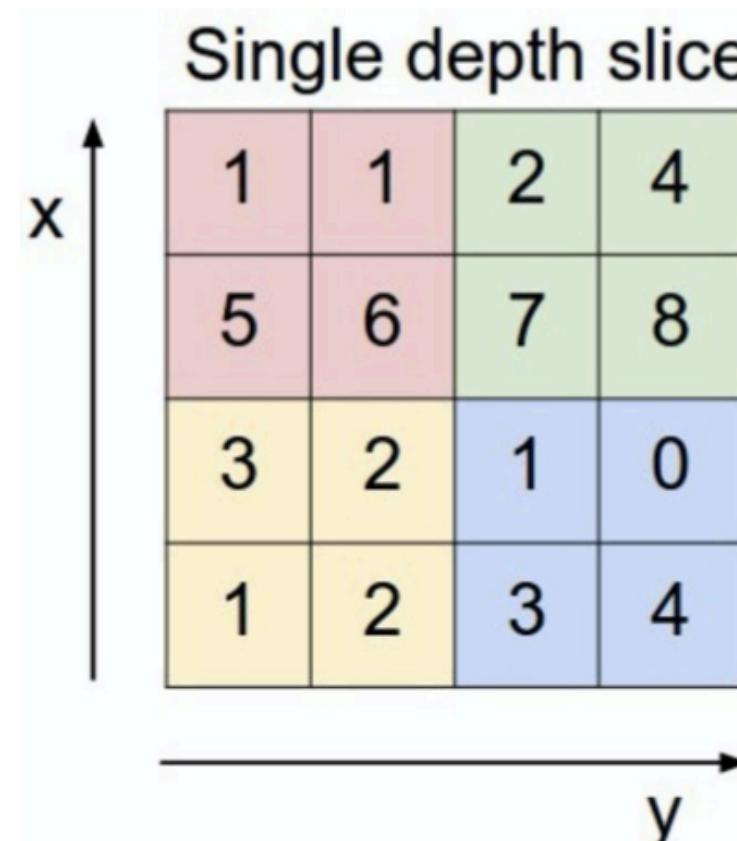


2D matrix to 1D matrix conversion



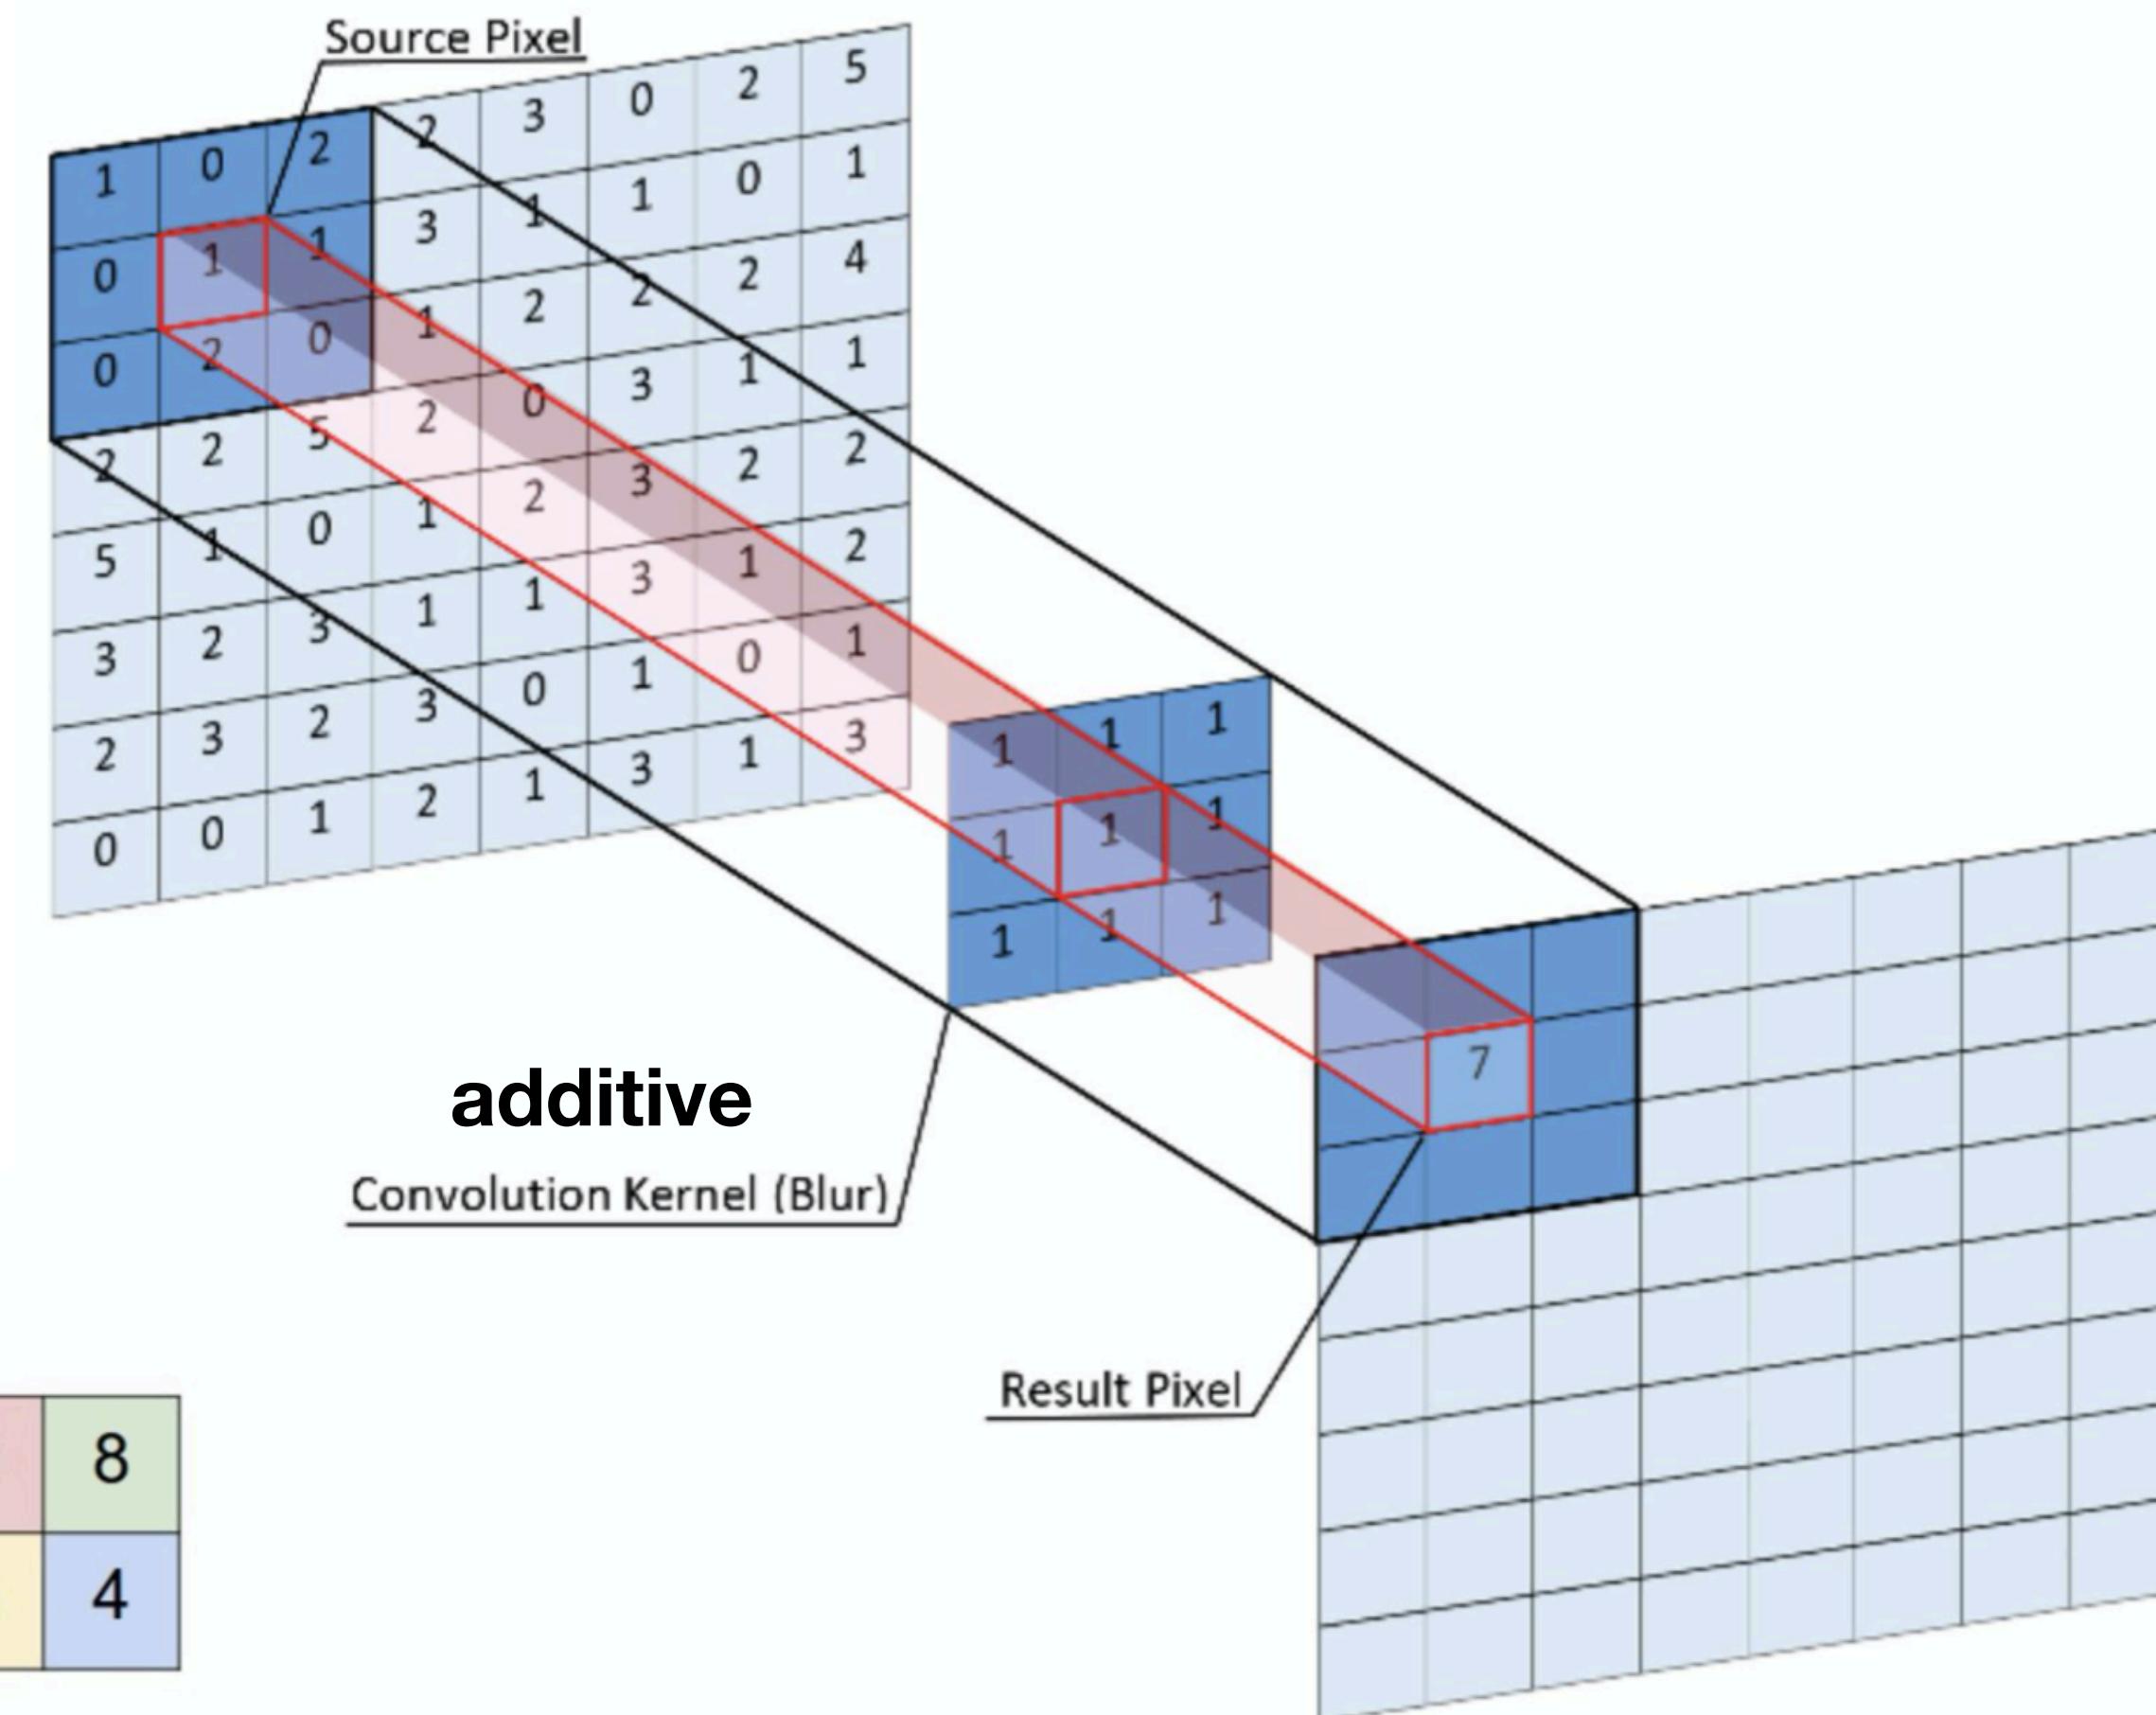
convolutions

- 3x3 blur example
- 3x3 striding
- 2x2 maxpool



max pool with 2x2 filters
and stride 2

6	8
3	4



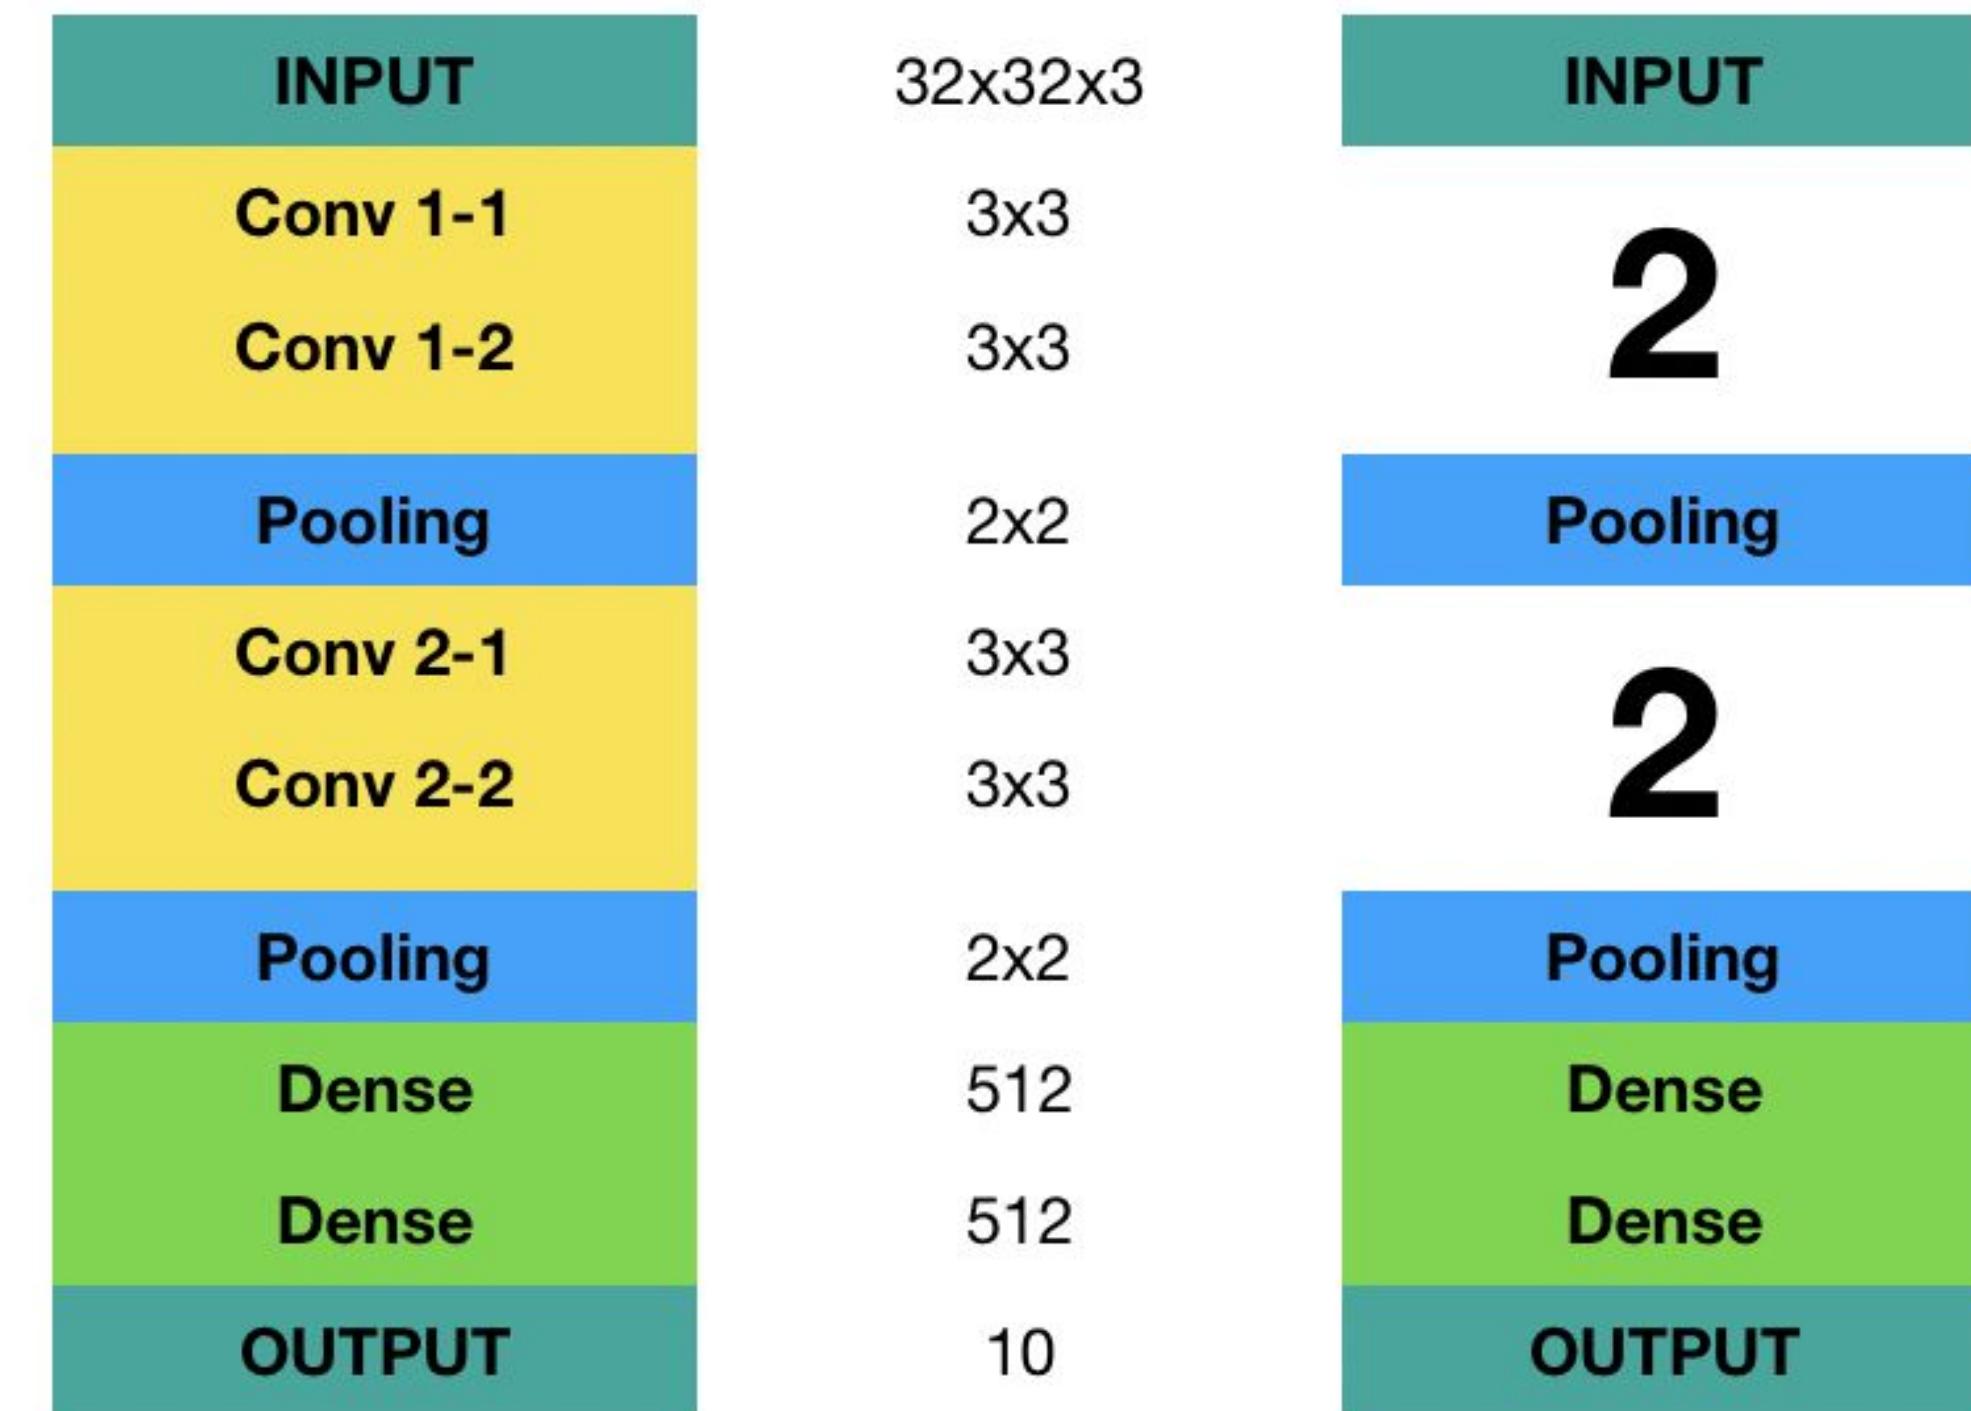
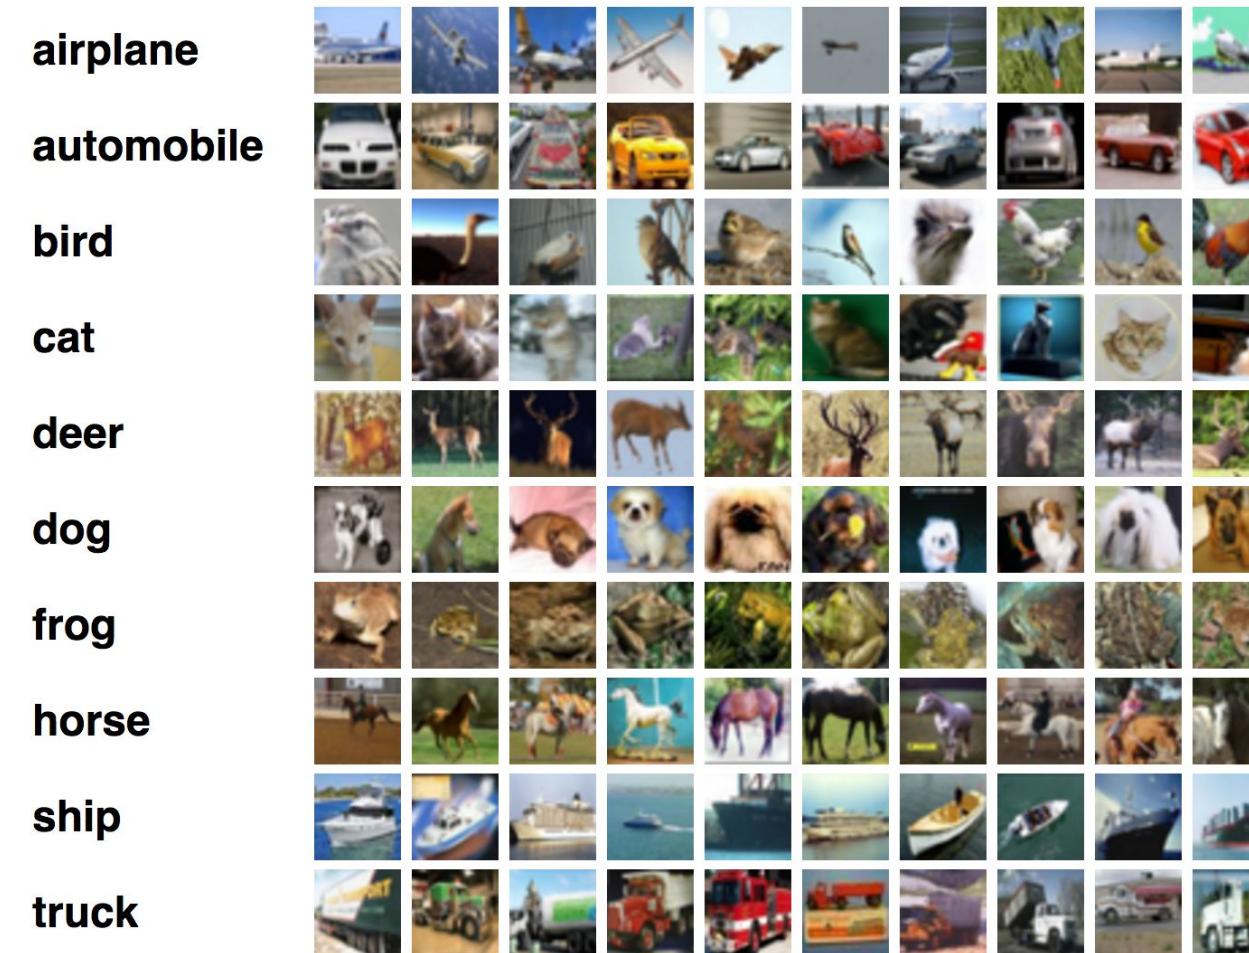
2d mnist

- 3x3 convolution
- 3x3 convolution
- 2x2 maxpool
- 2 * 512 fully connected layers

INPUT	28x28x1
Conv 1-1	3x3
Conv 1-2	3x3
Pooling	2x2
Dense	512
Dense	512
OUTPUT	10

2d stacking, color, cifar

- [3x3 striding, 3x3 striding, 2x2 maxpool]
→ [block1] + [block2]
- 2 * 512 fully connected layers
- Cifar: (32x32x3, 10 categories)



vgg (2014)

- arxiv:1409.1556
- vgg16: [2, 2, 3, 3, 3]
- vgg19: [2, 2, 4, 4, 4]
- 2 * 4096 fully connected layers
- imagenet: (224x224x3, 1000 categories)



residual networks (2015)

- arxiv: 1512.03385
- cnn backbone + skip connections → resnet 34
- resnet 34 + bottleneck blocks ($1 \times 3 \times 1$) → resnet 50

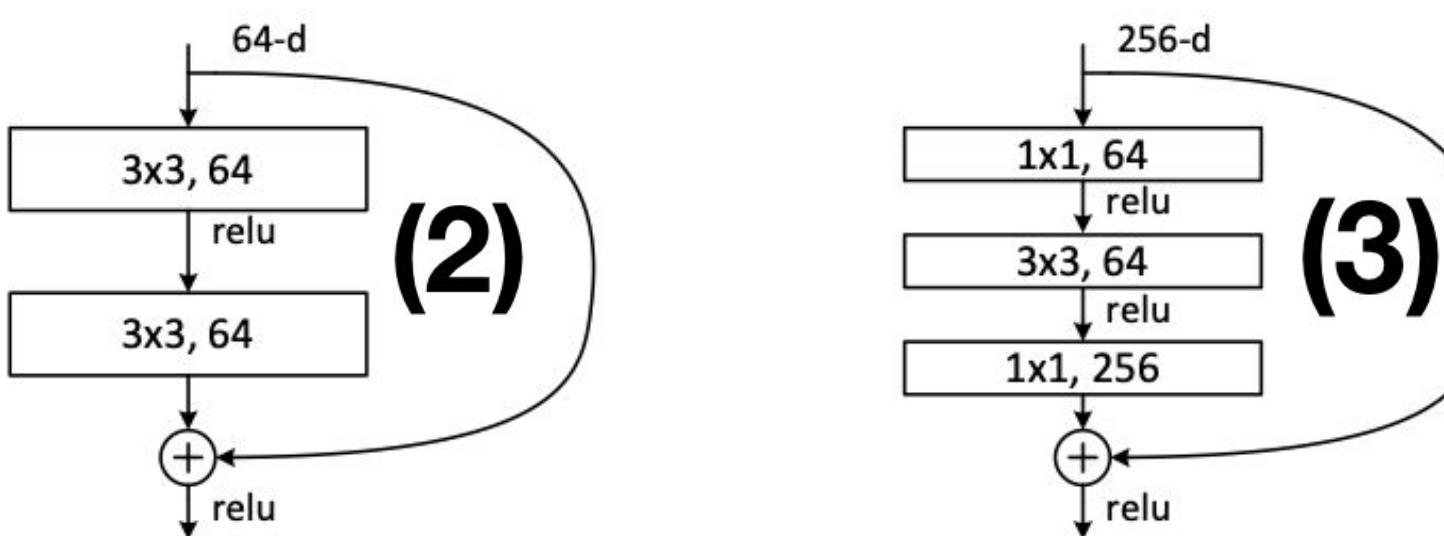
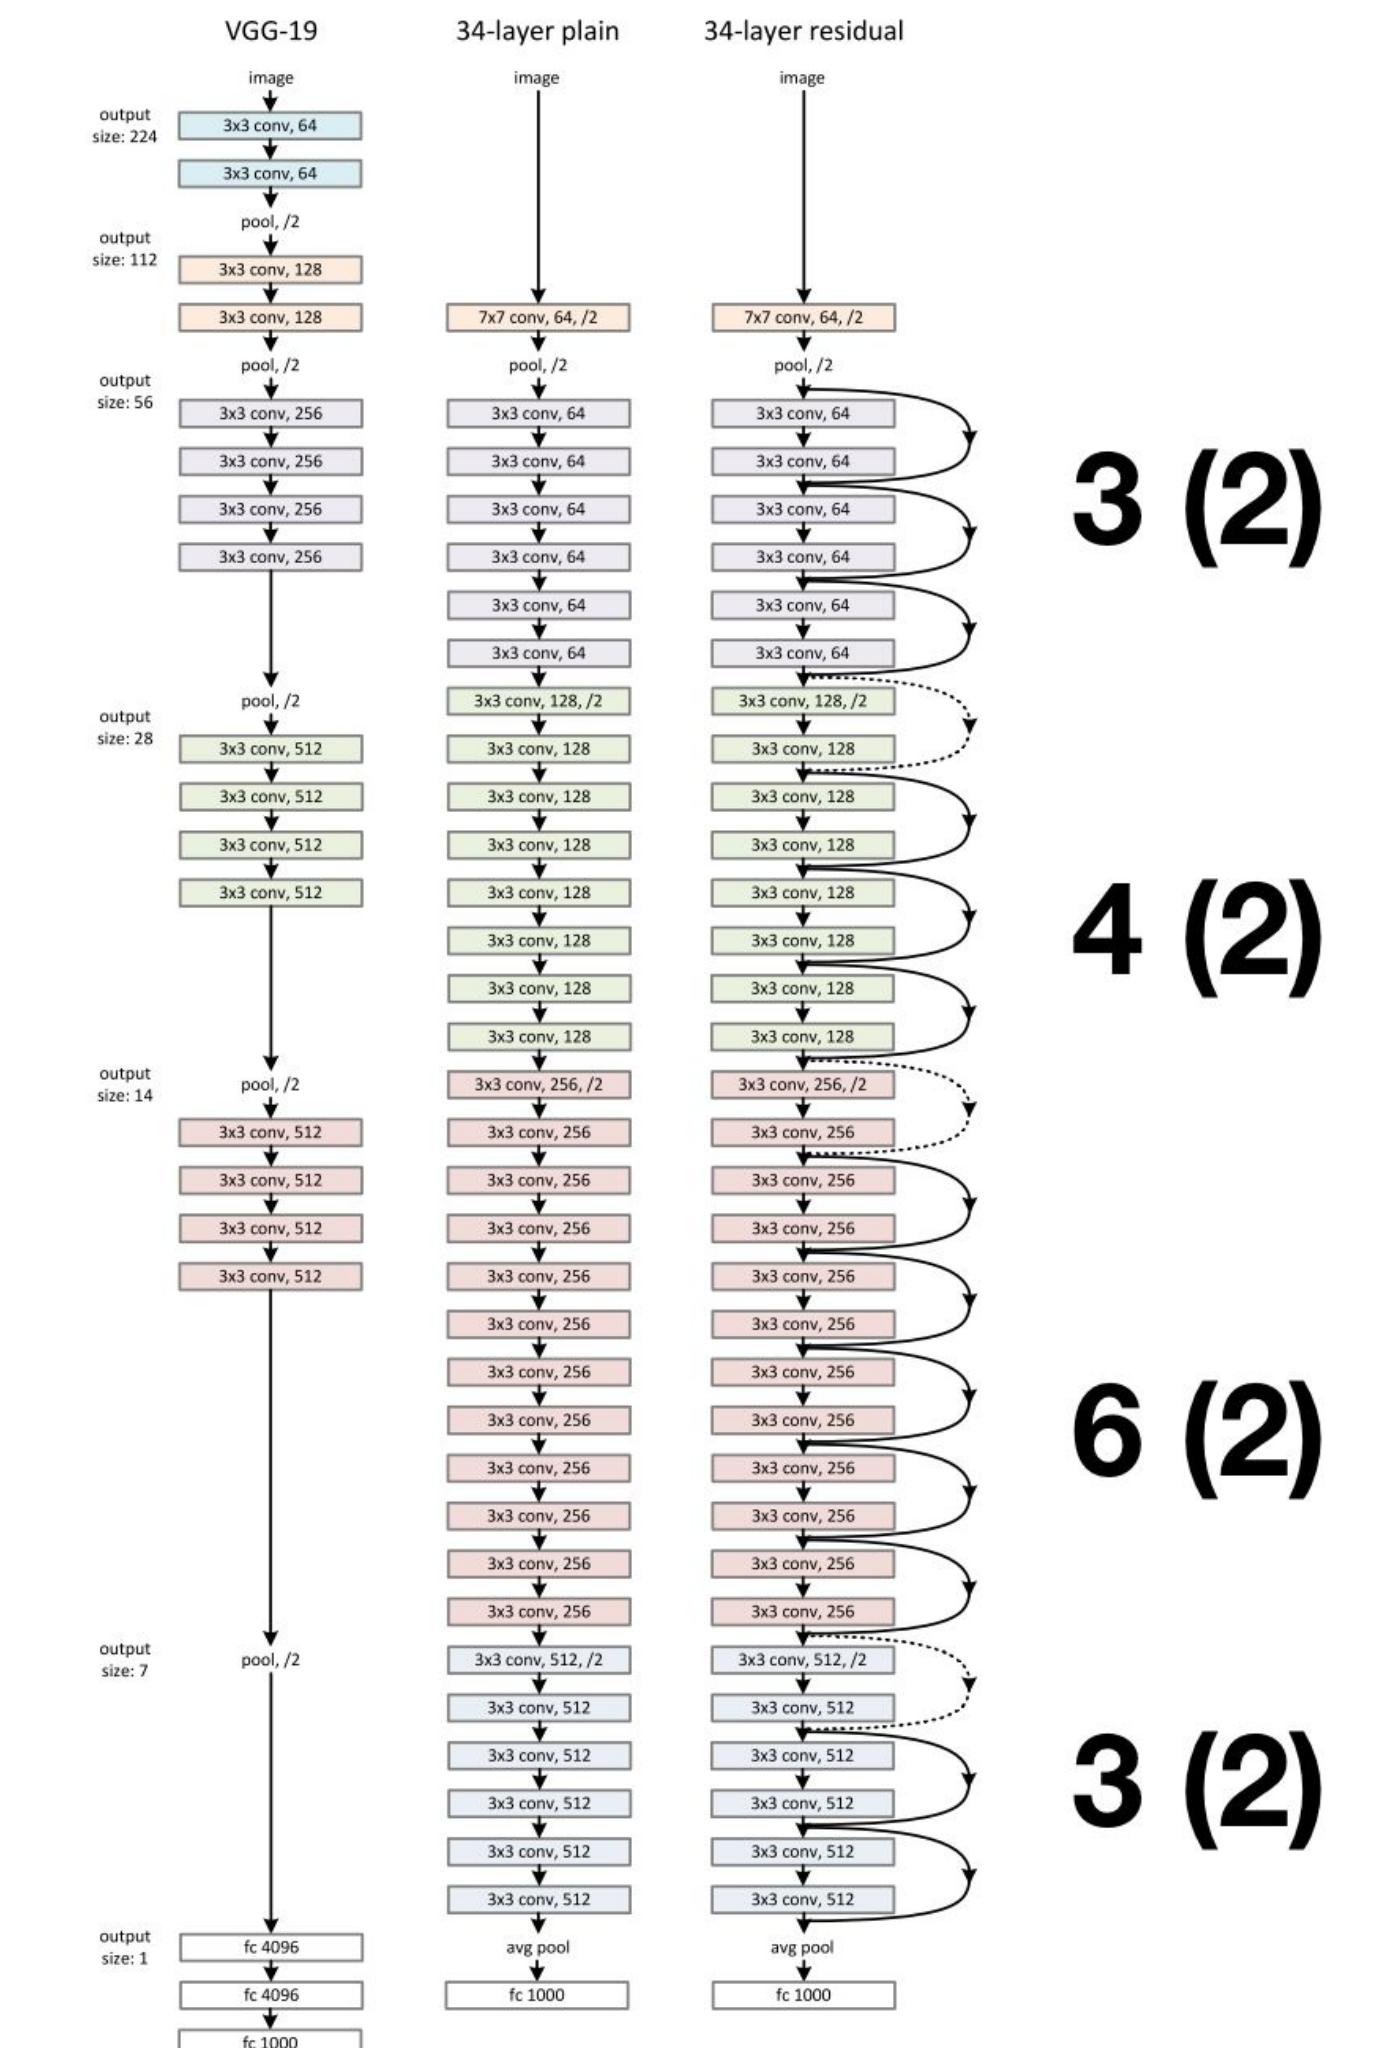
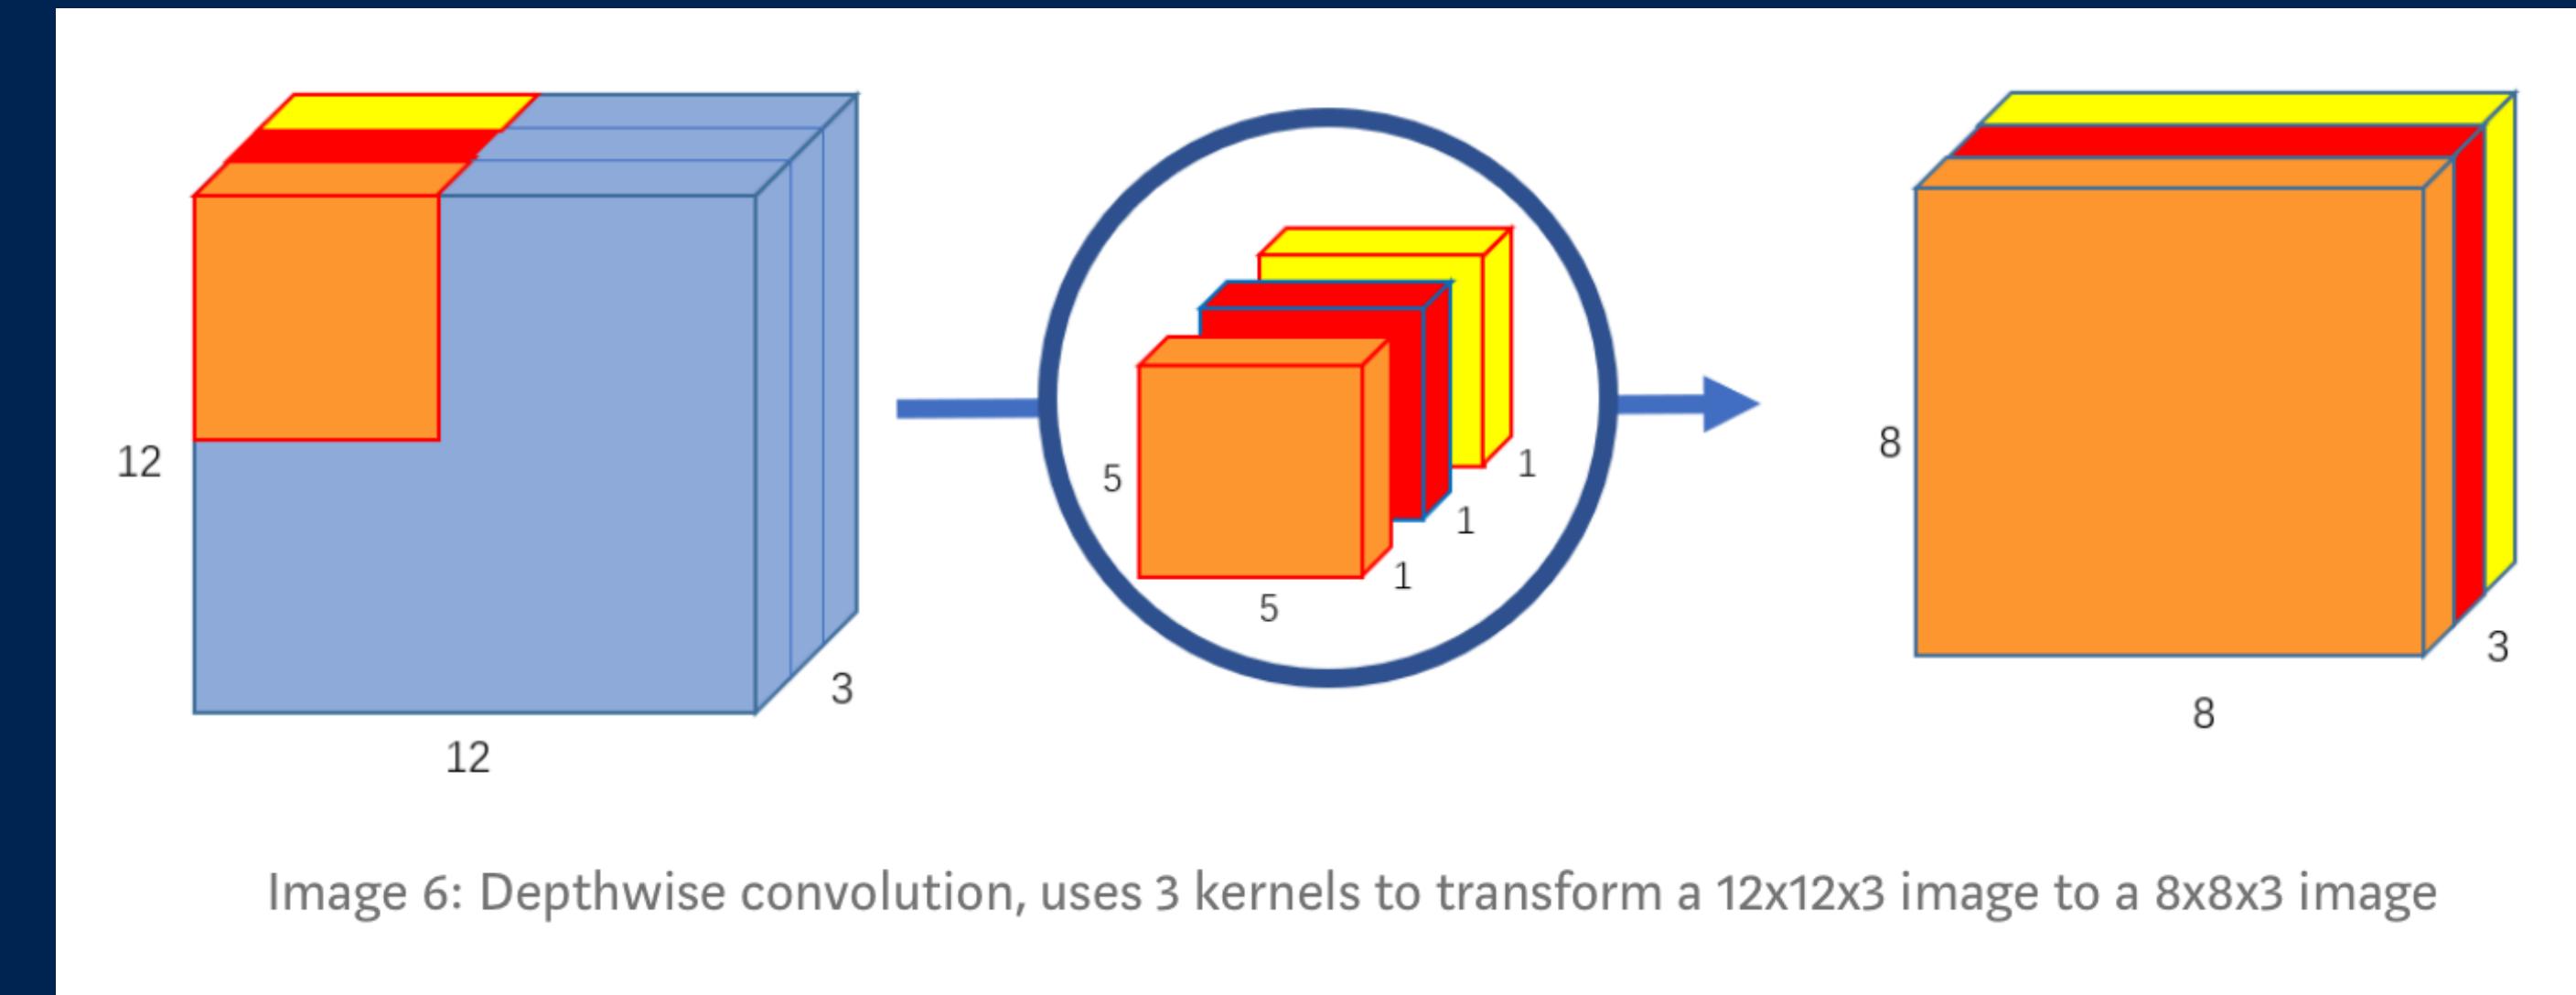


Figure 5. A deeper residual function \mathcal{F} for ImageNet. Left: a building block (on 56×56 feature maps) as in Fig. 3 for ResNet-34. Right: a “bottleneck” building block for ResNet-50/101/152.



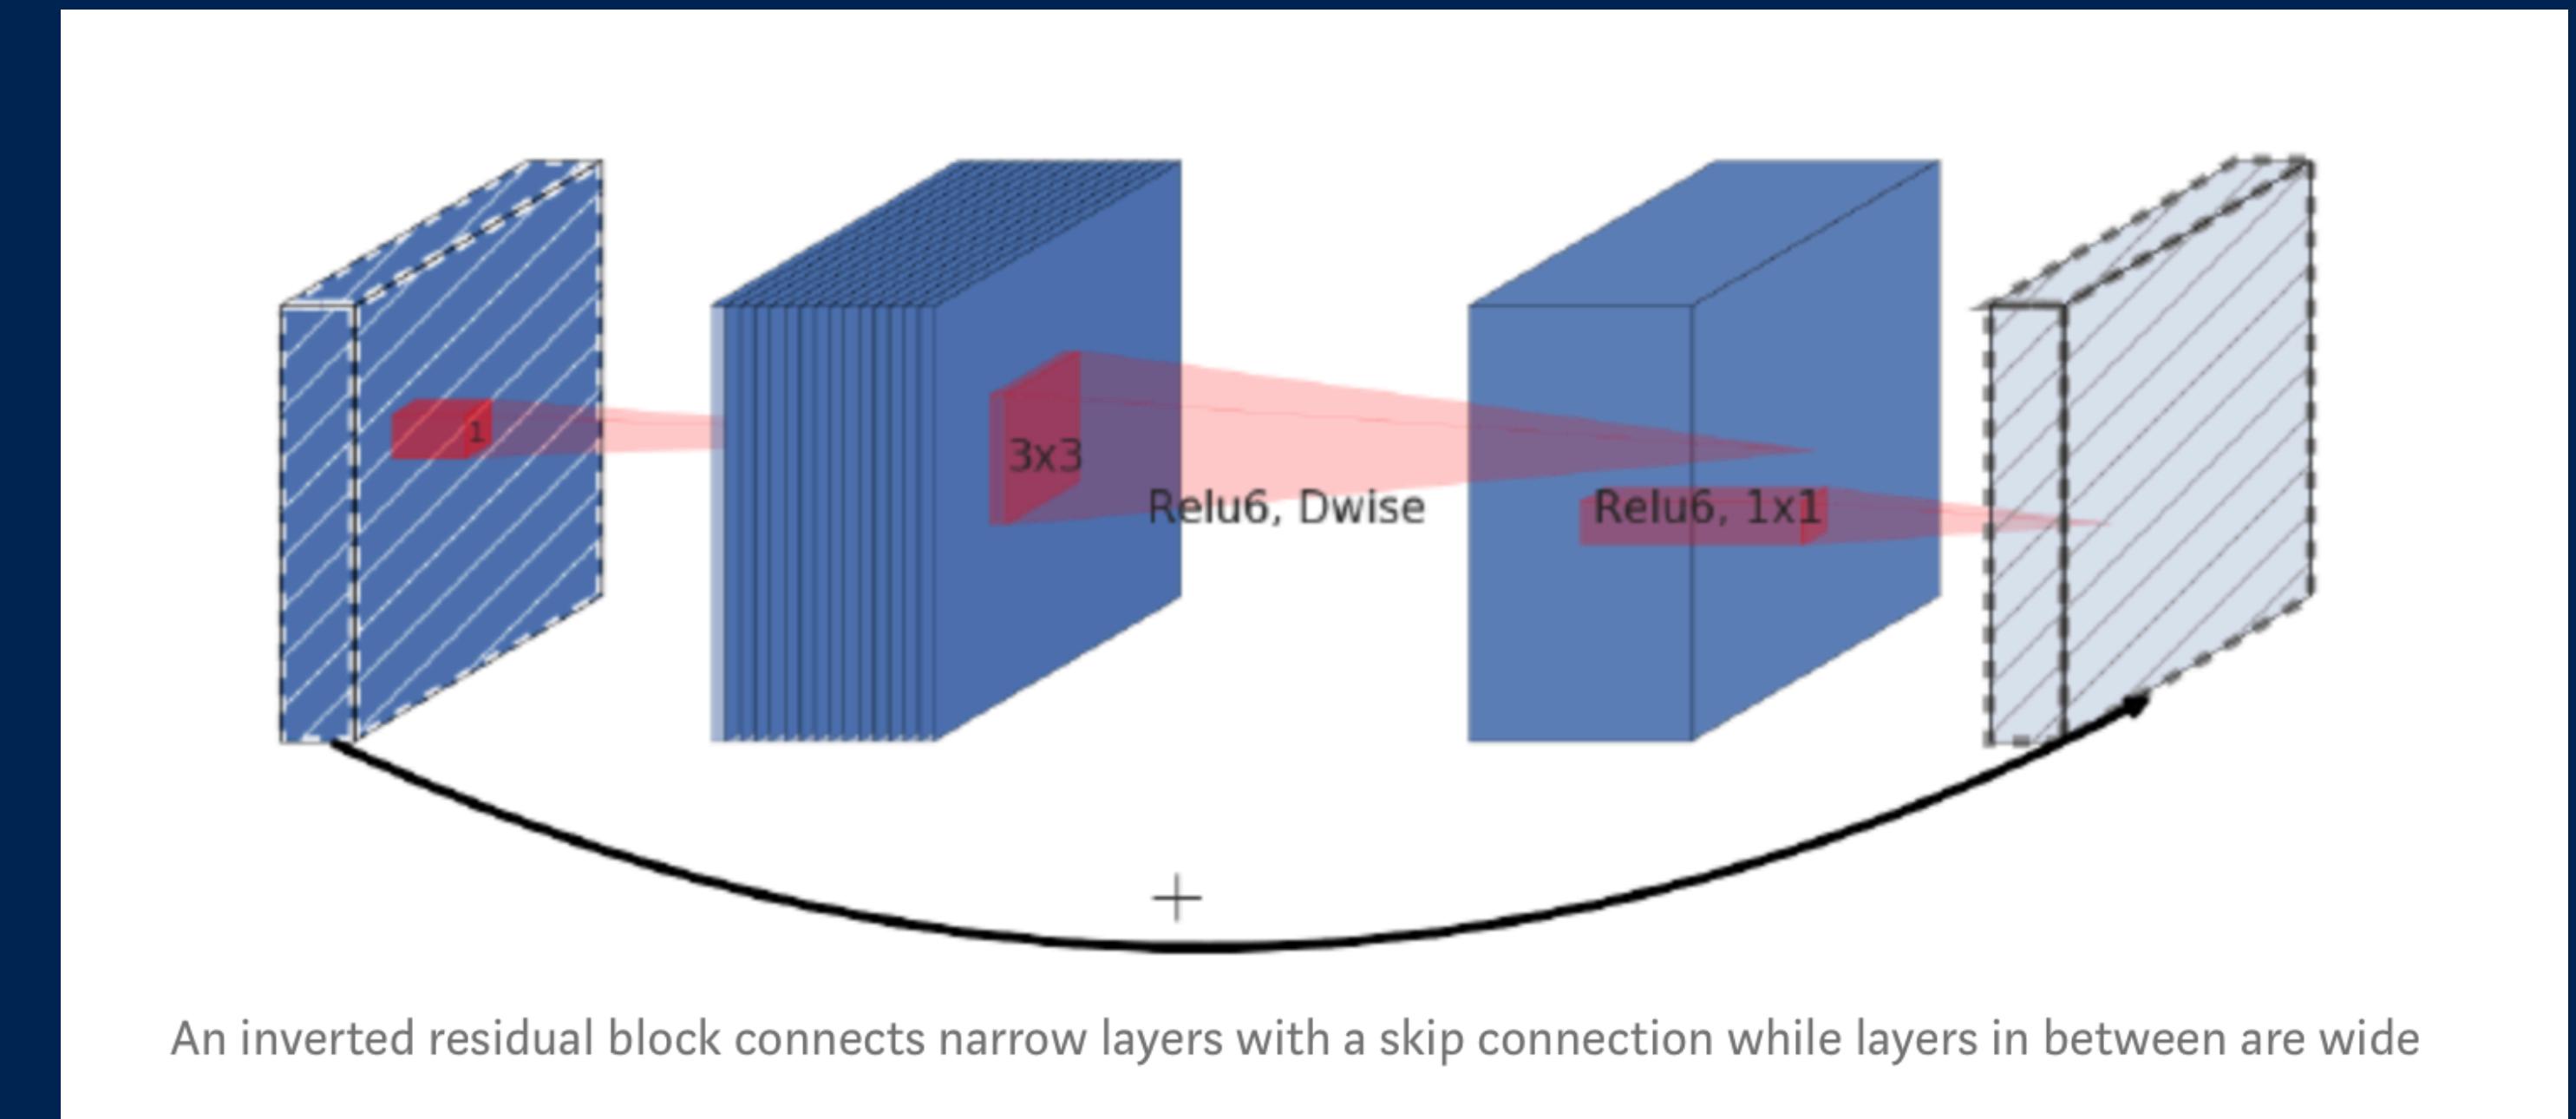
mobilenet v1

- 2017 talk:
**depthwise,
pointwise
convolutions**
- [towardsdatascience.com/a-basic-introduction-to-separable-convolutions-](#)
[b99ec3102728](#)

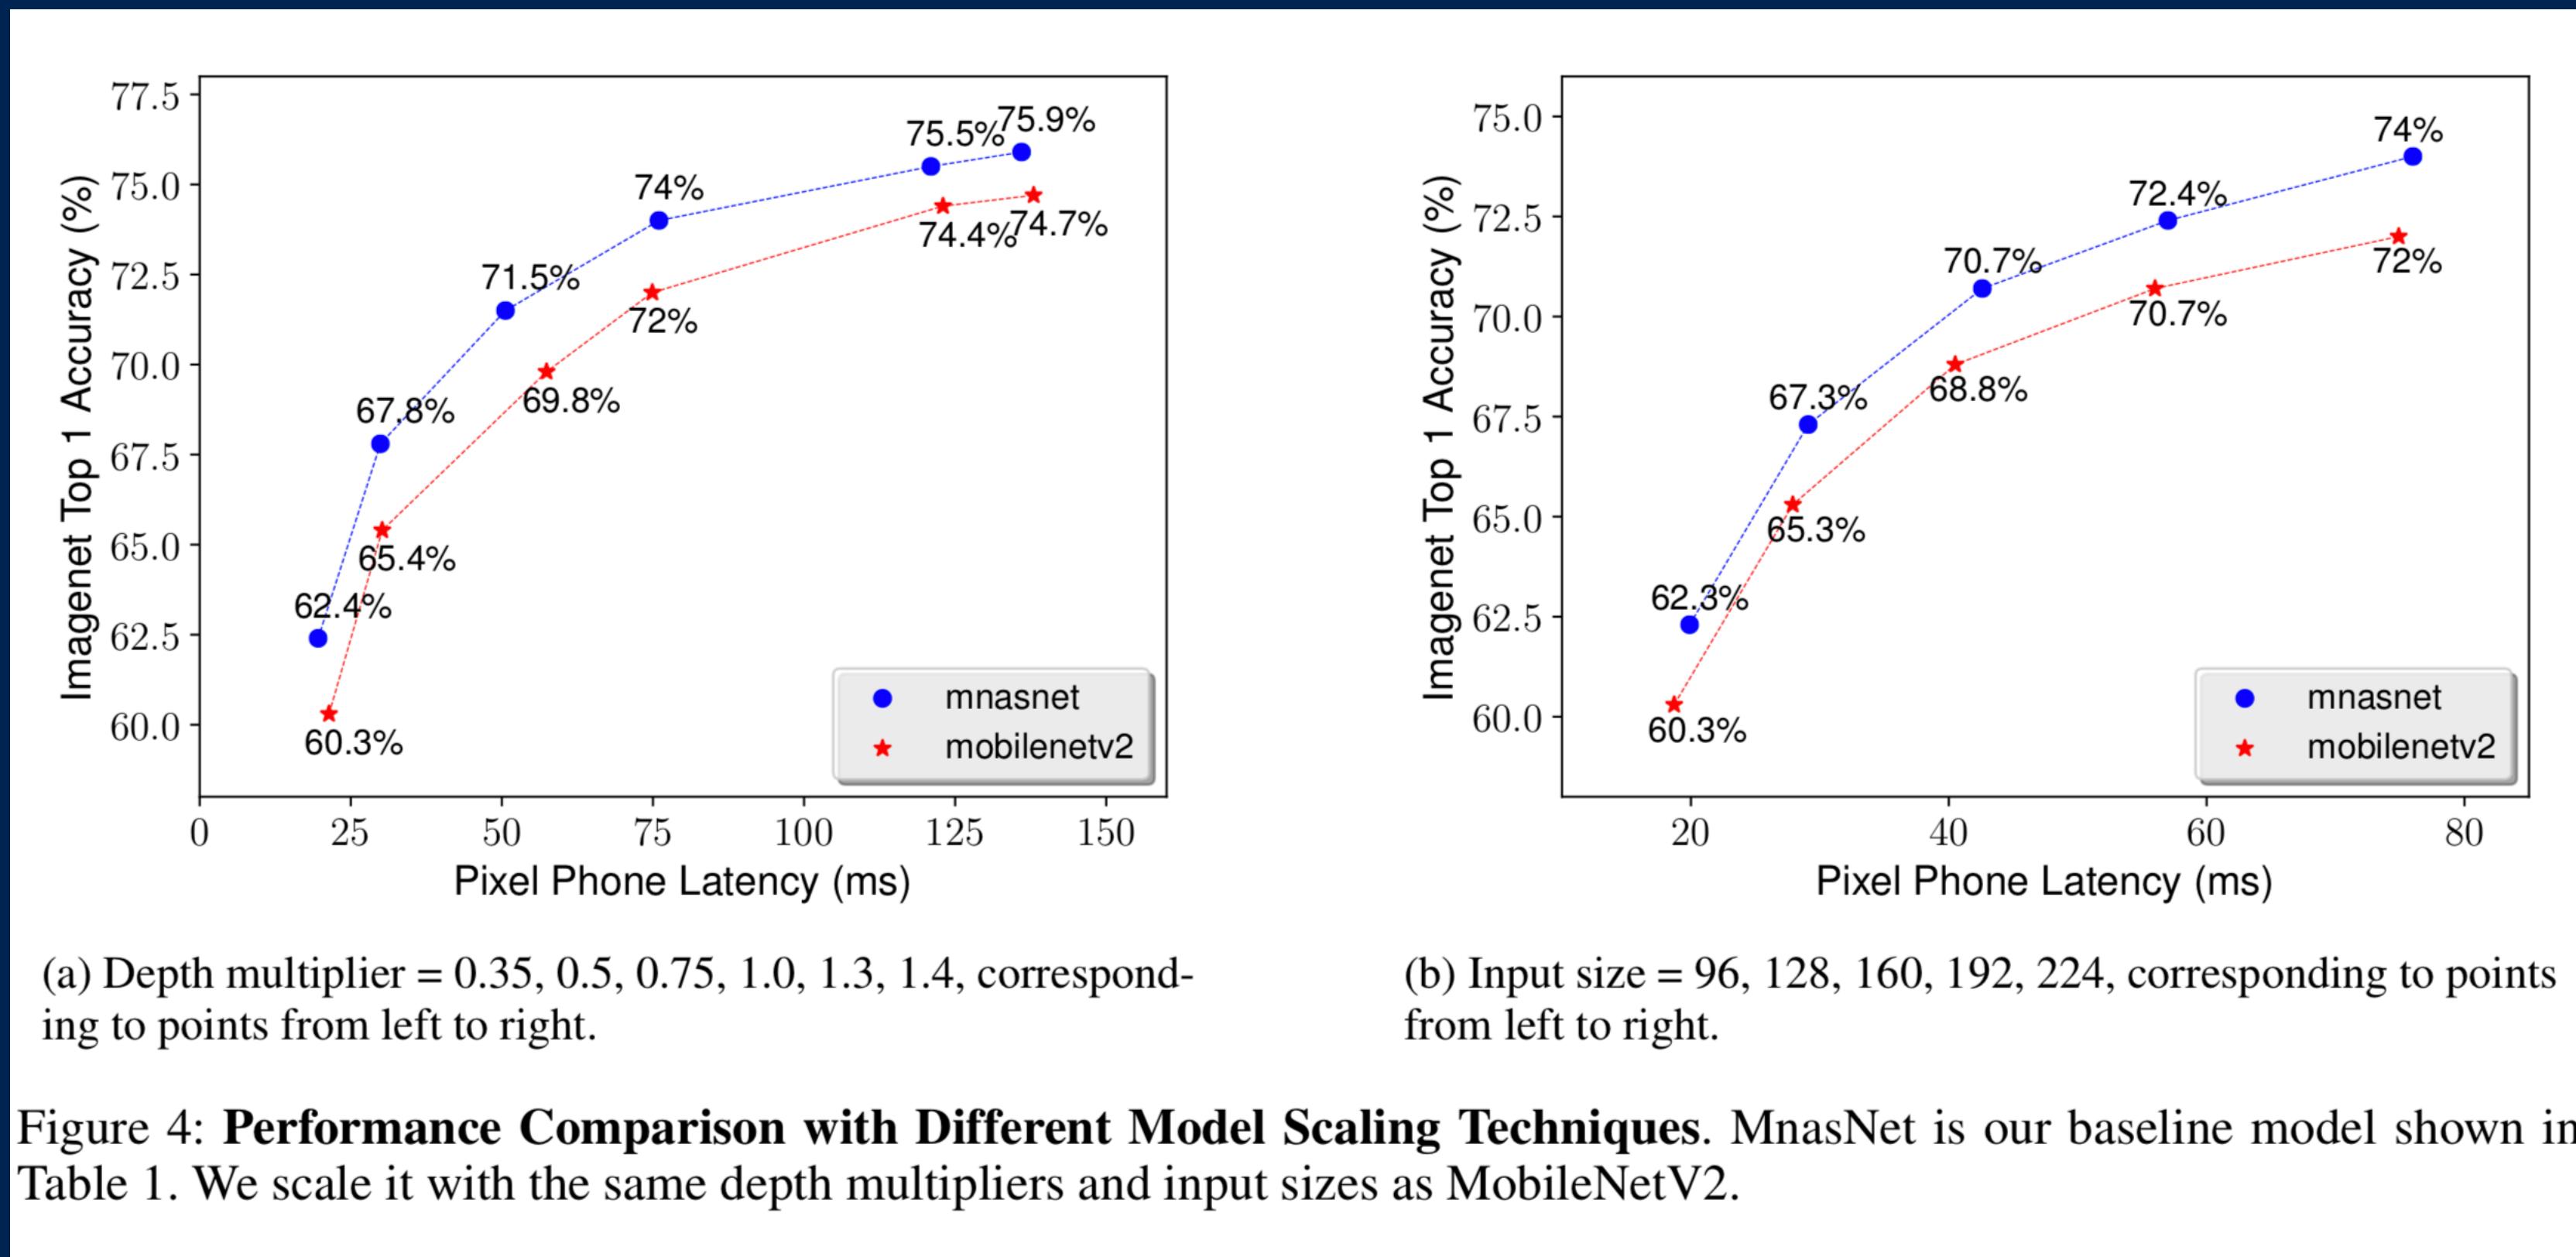


mobilenet v2

- 2018 talk: linear activations, inverted bottleneck layers
- towardsdatascience.com/mobilenetv2-inverted-residuals-and-linear-bottlenecks-8a4362f4ffd5



mnasnet



senet

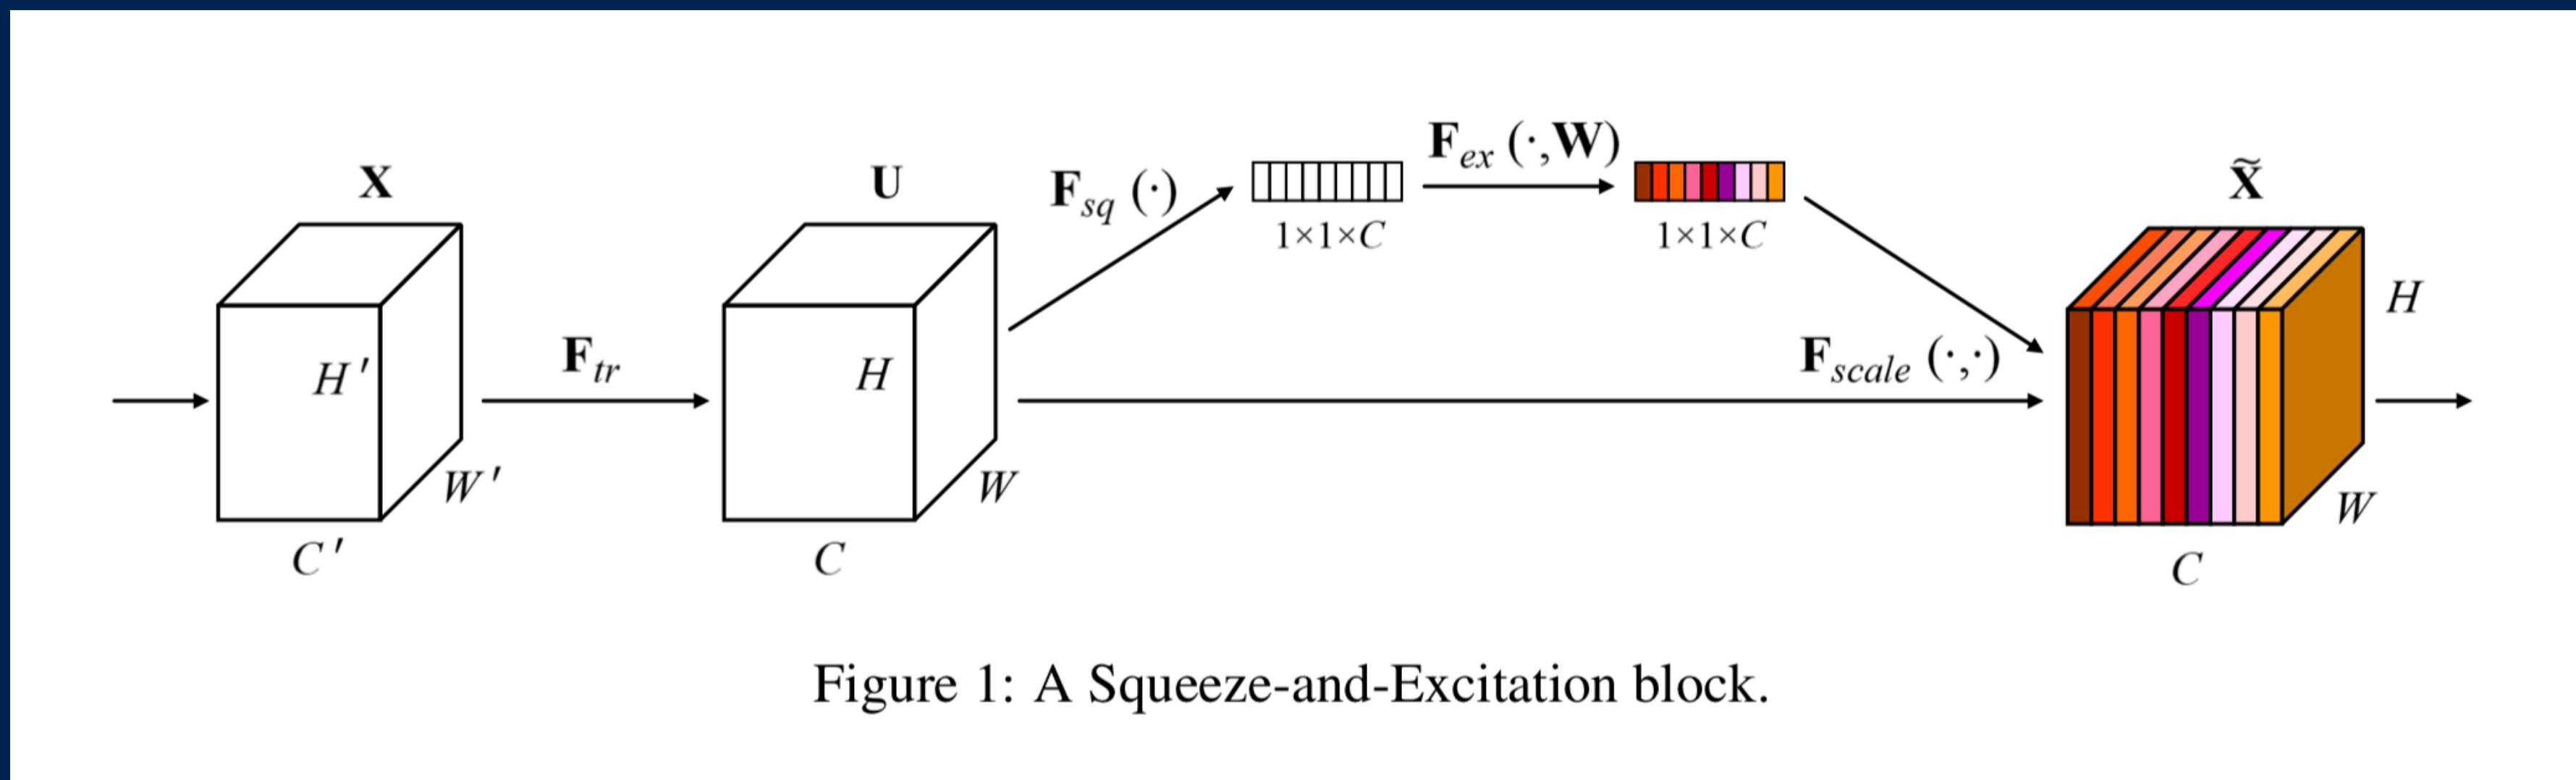


Figure 1: A Squeeze-and-Excitation block.

- **squeeze + excitation networks (2018)**
- **<https://www.robots.ox.ac.uk/~vgg/publications/2018/Hu18/presentation.pdf>**

efficientnet (may 2019)

- arxiv: 1905.11946

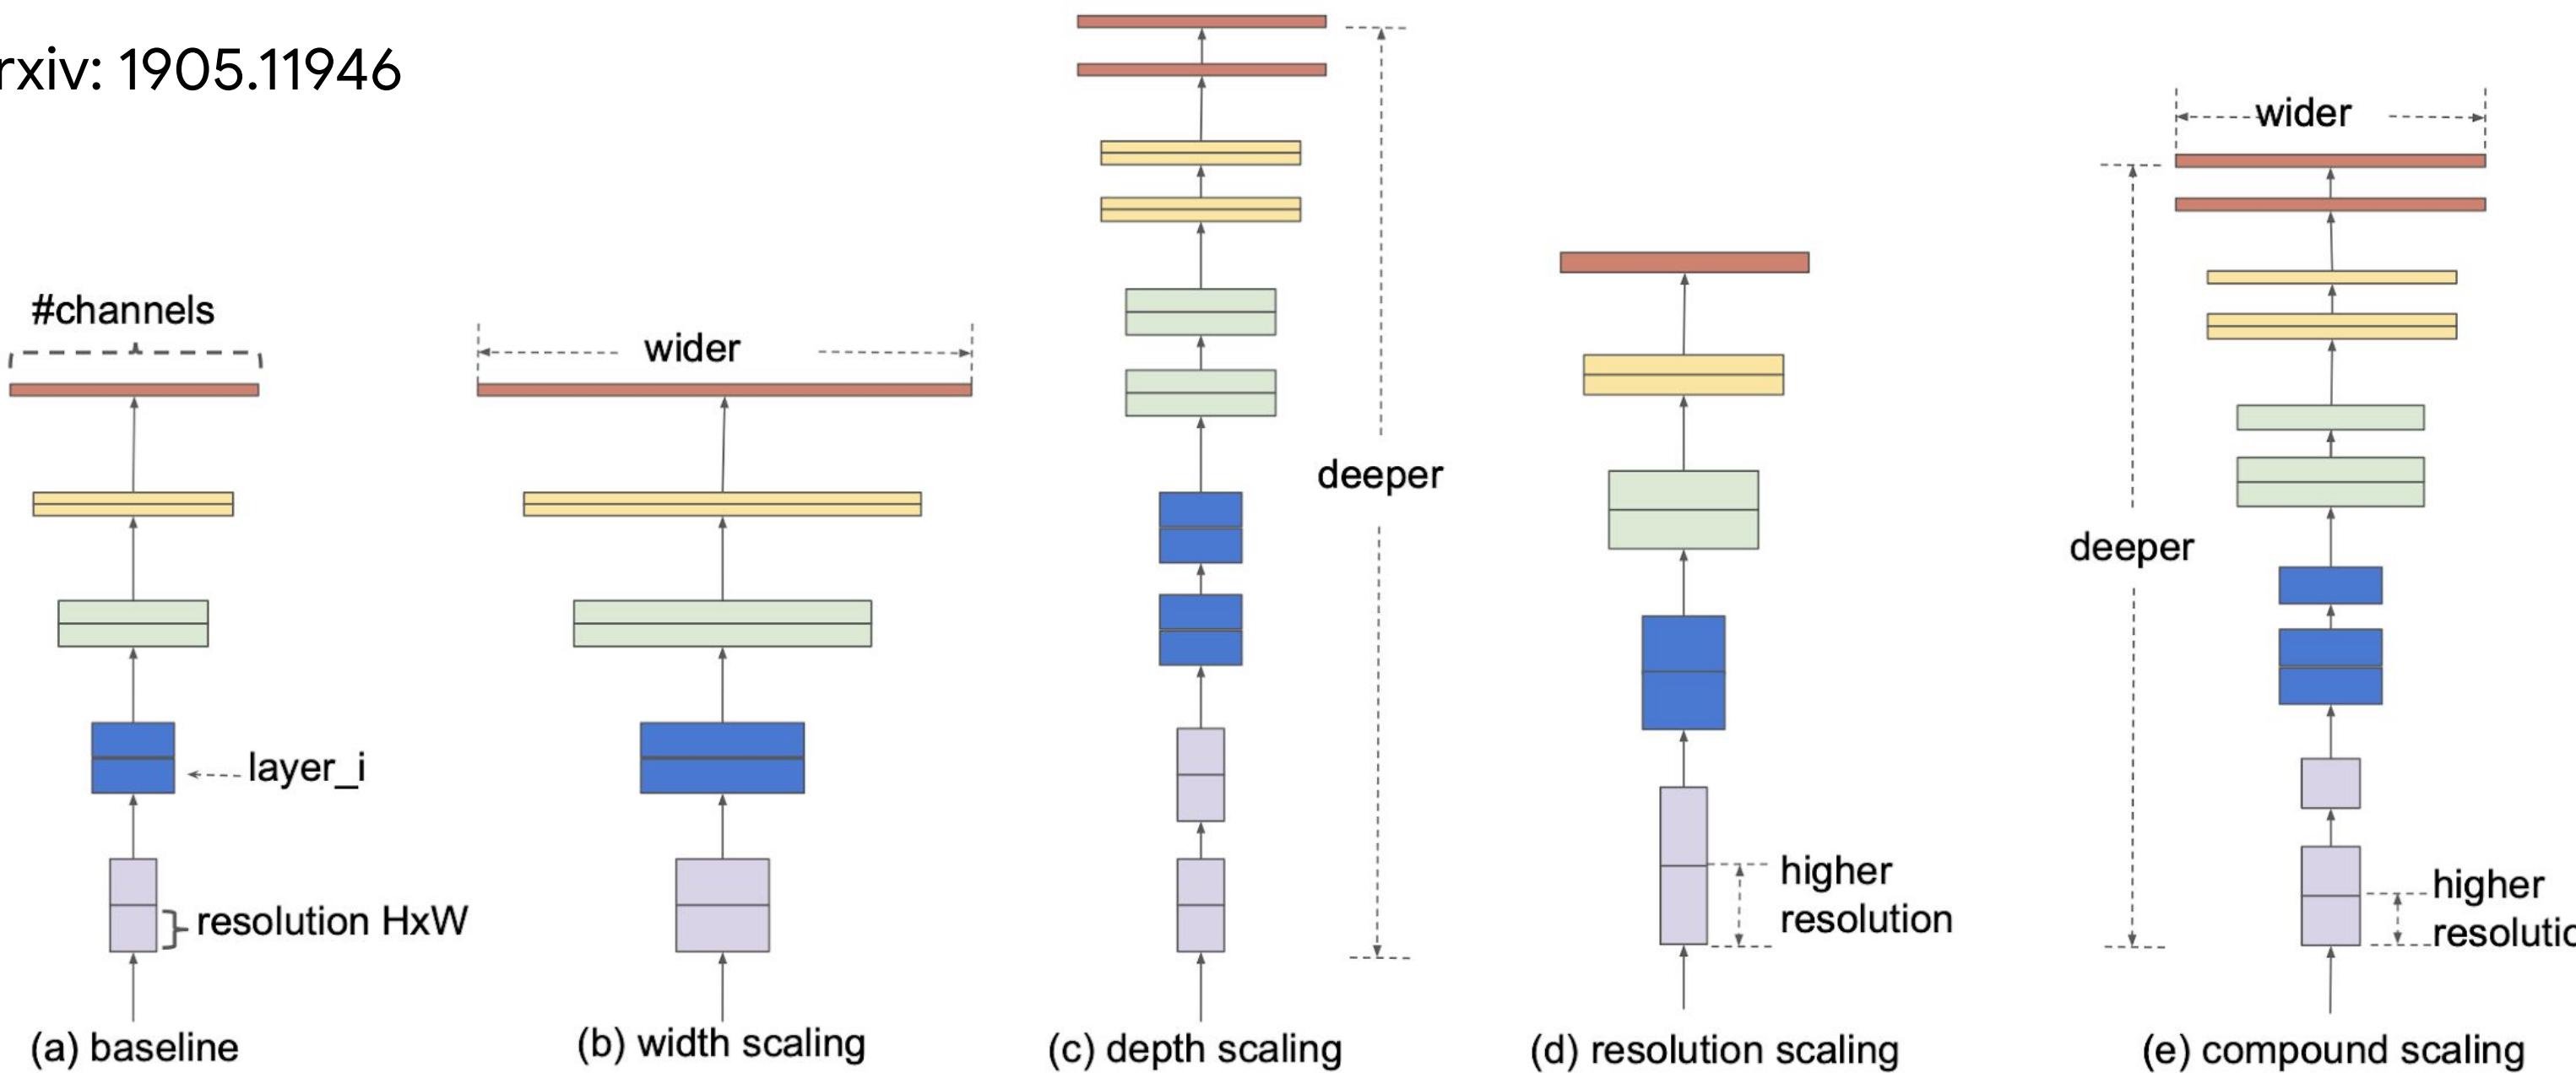
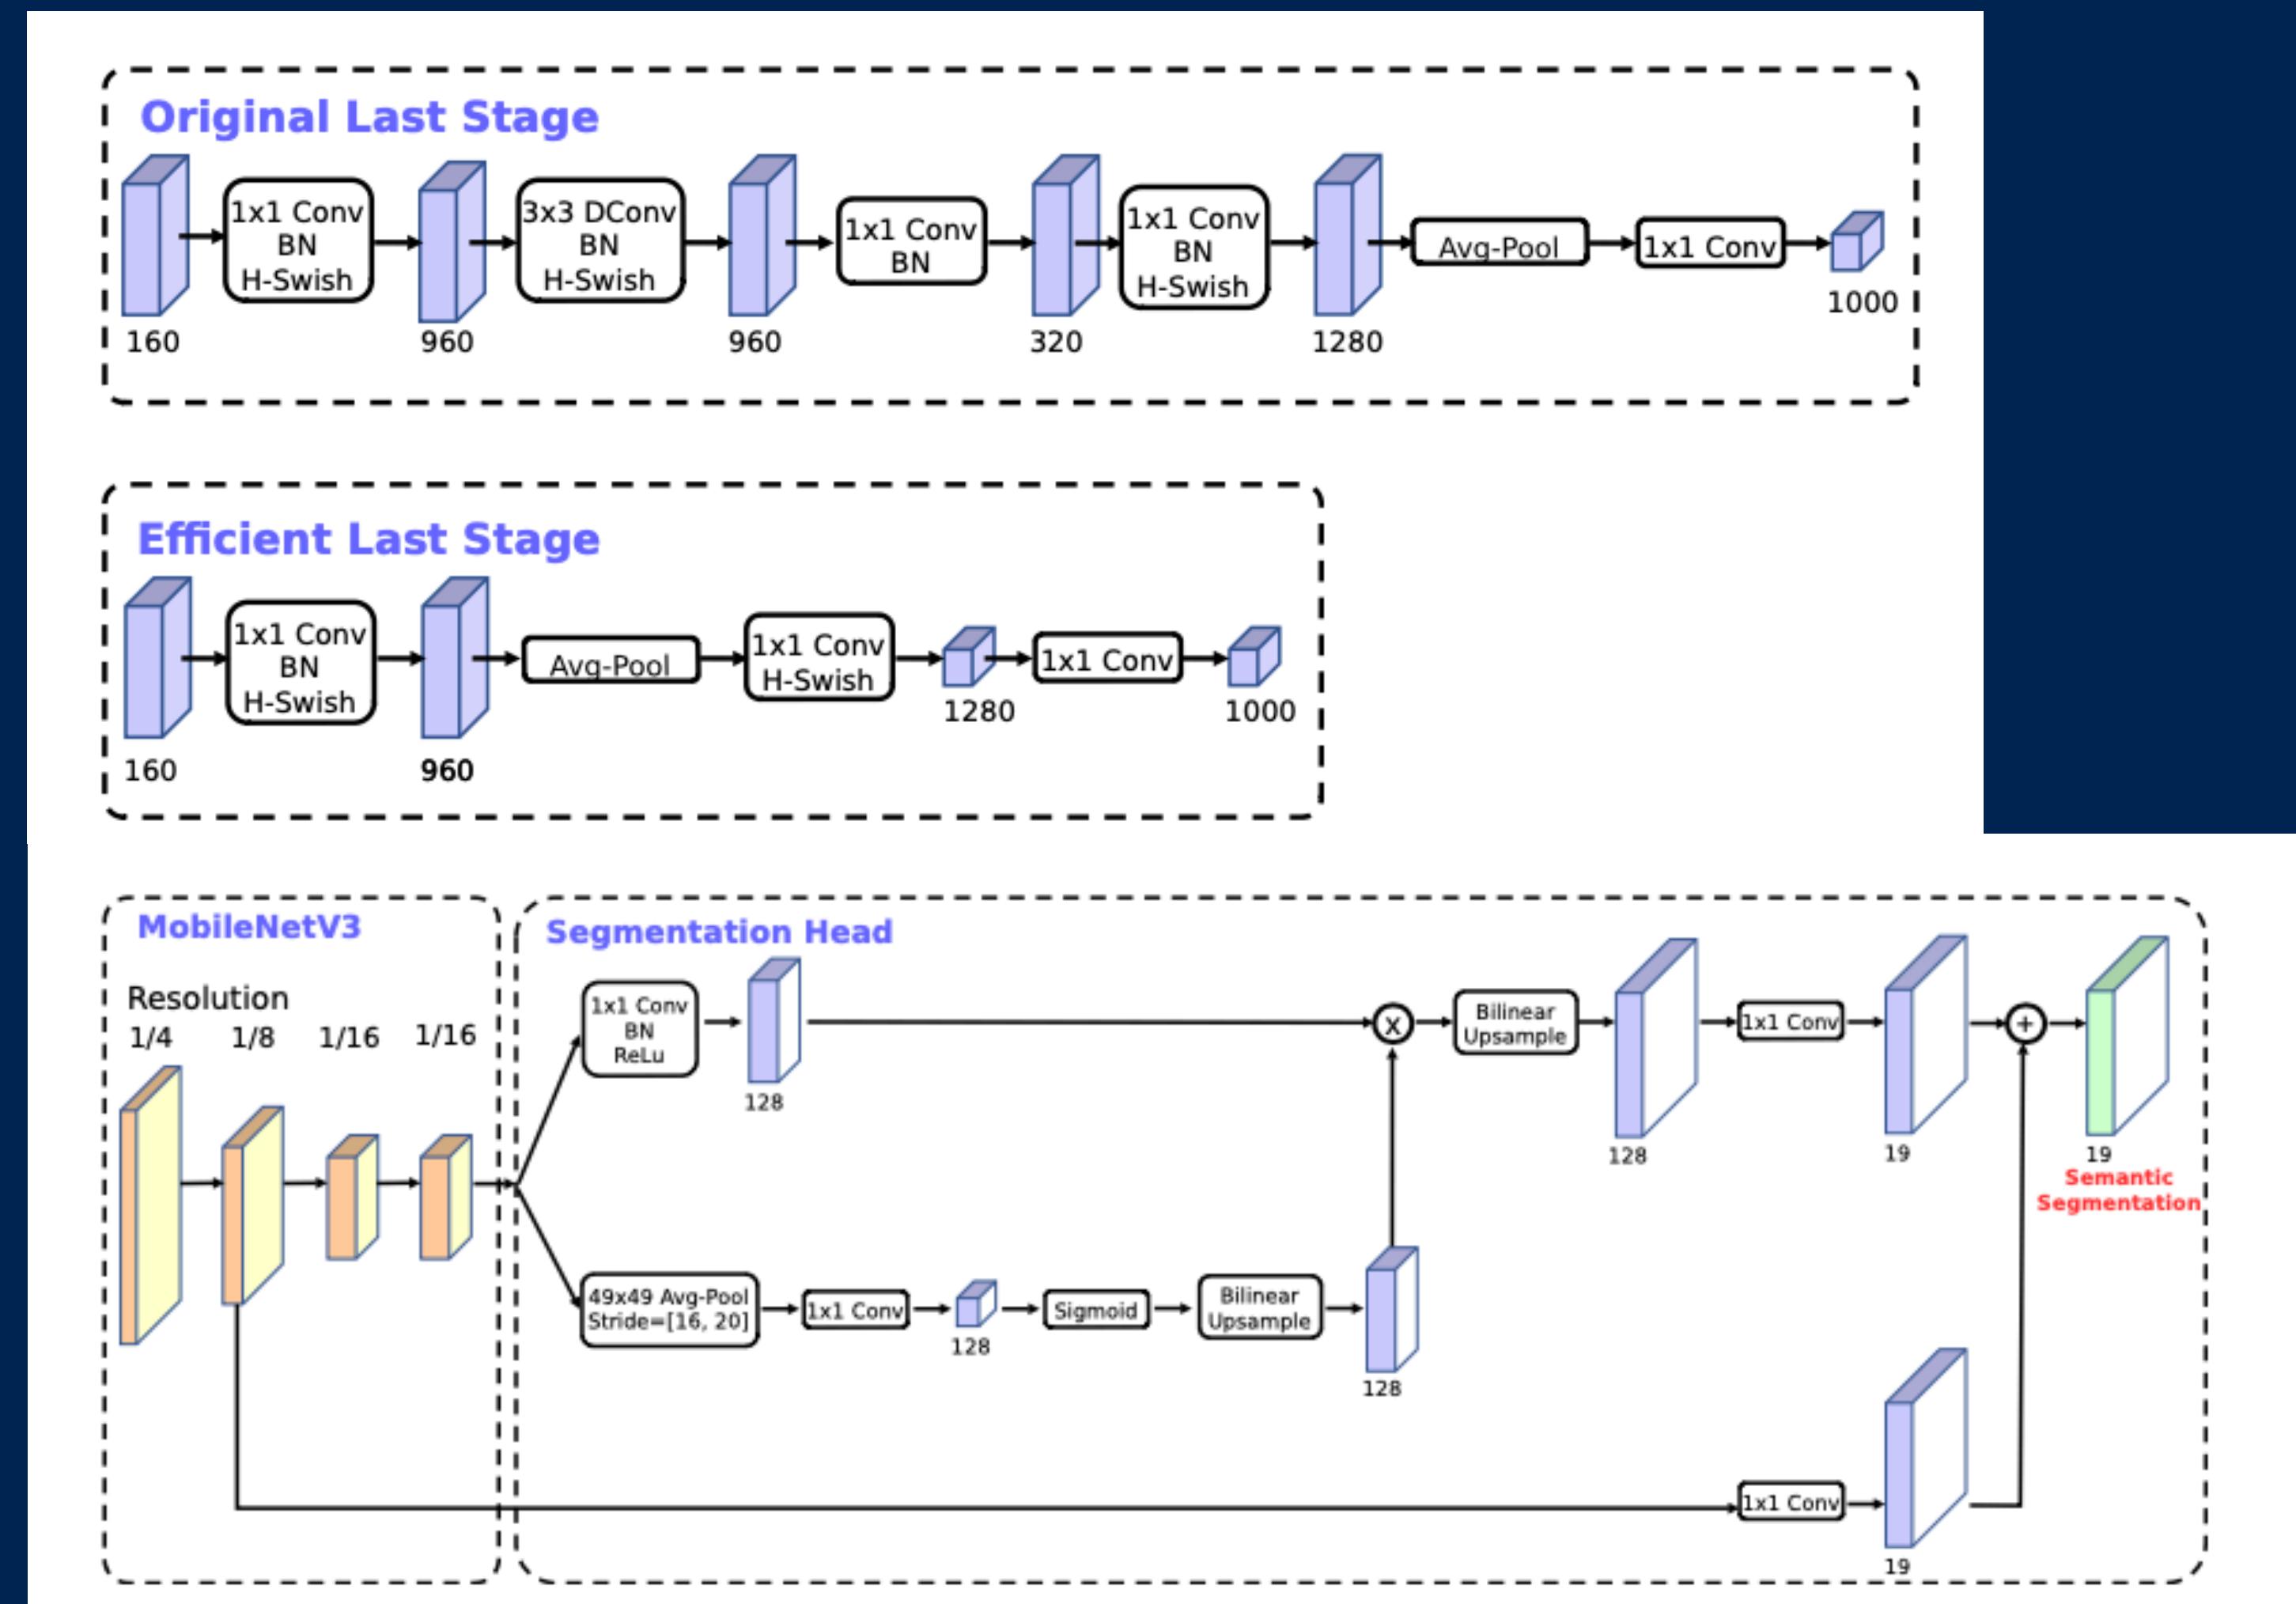


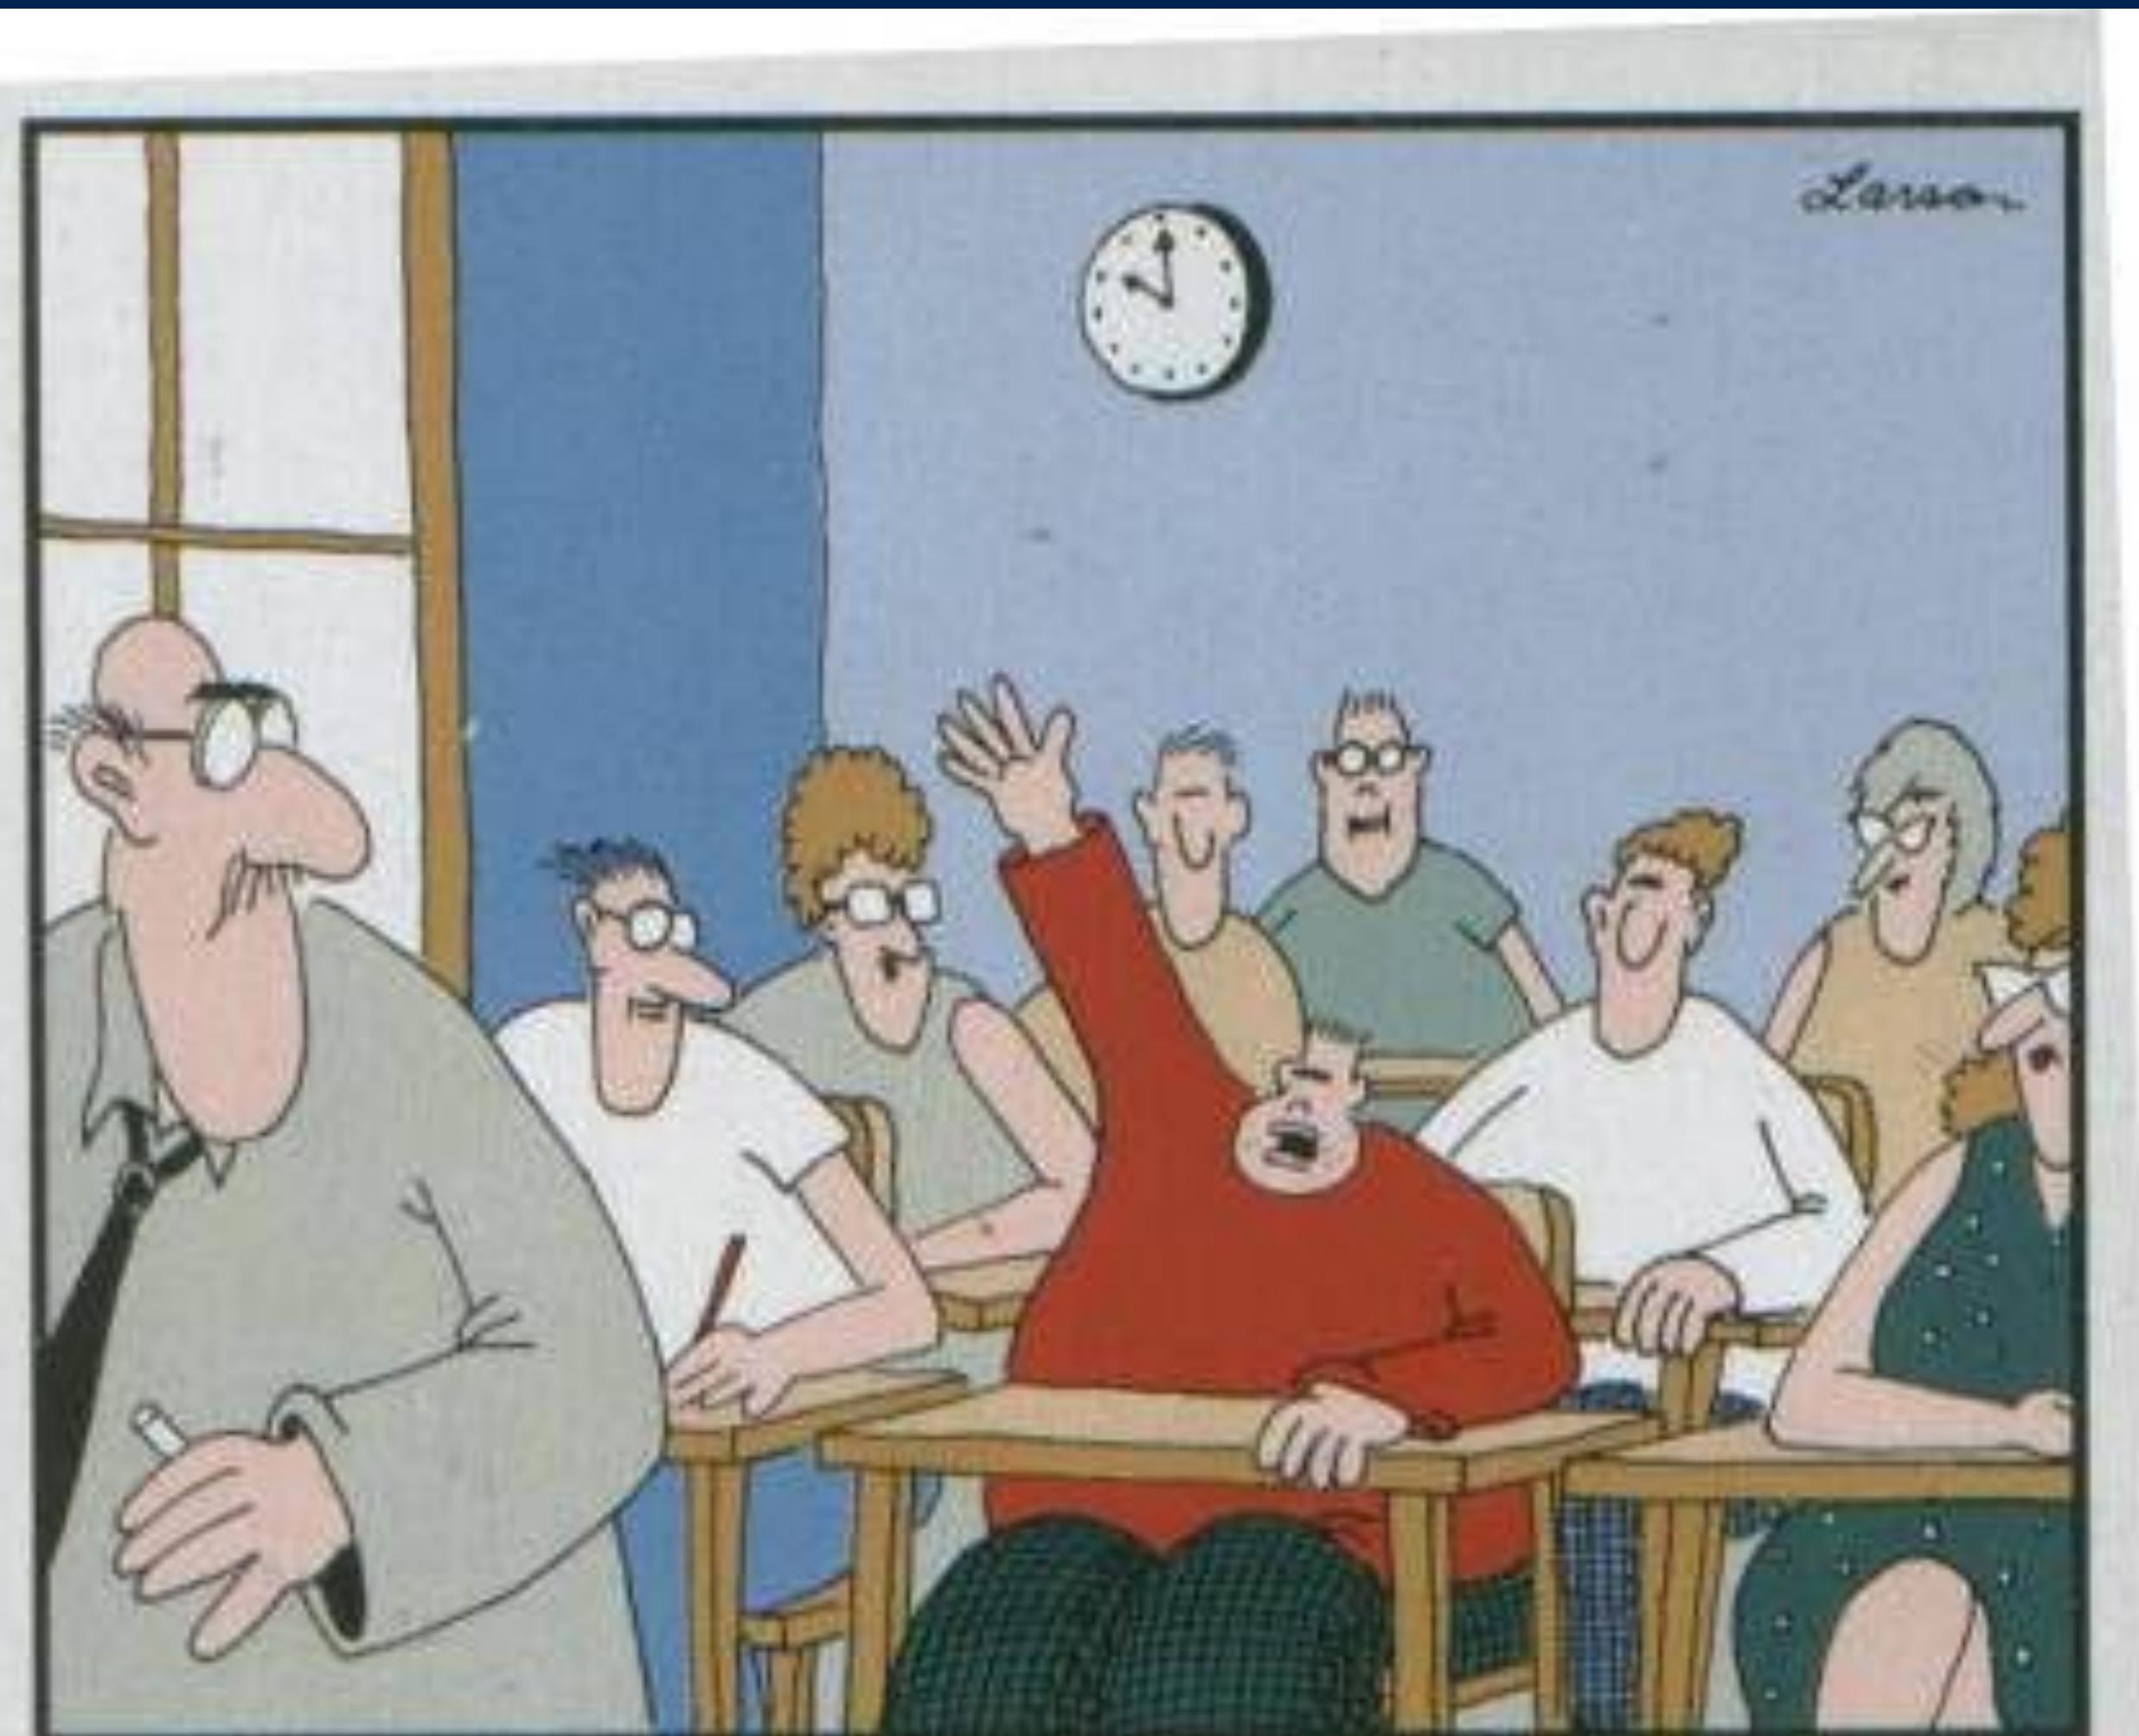
Figure 2. Model Scaling. (a) is a baseline network example; (b)-(d) are conventional scaling that only increases one dimension of network width, depth, or resolution. (e) is our proposed compound scaling method that uniformly scales all three dimensions with a fixed ratio.

mobilenet v3

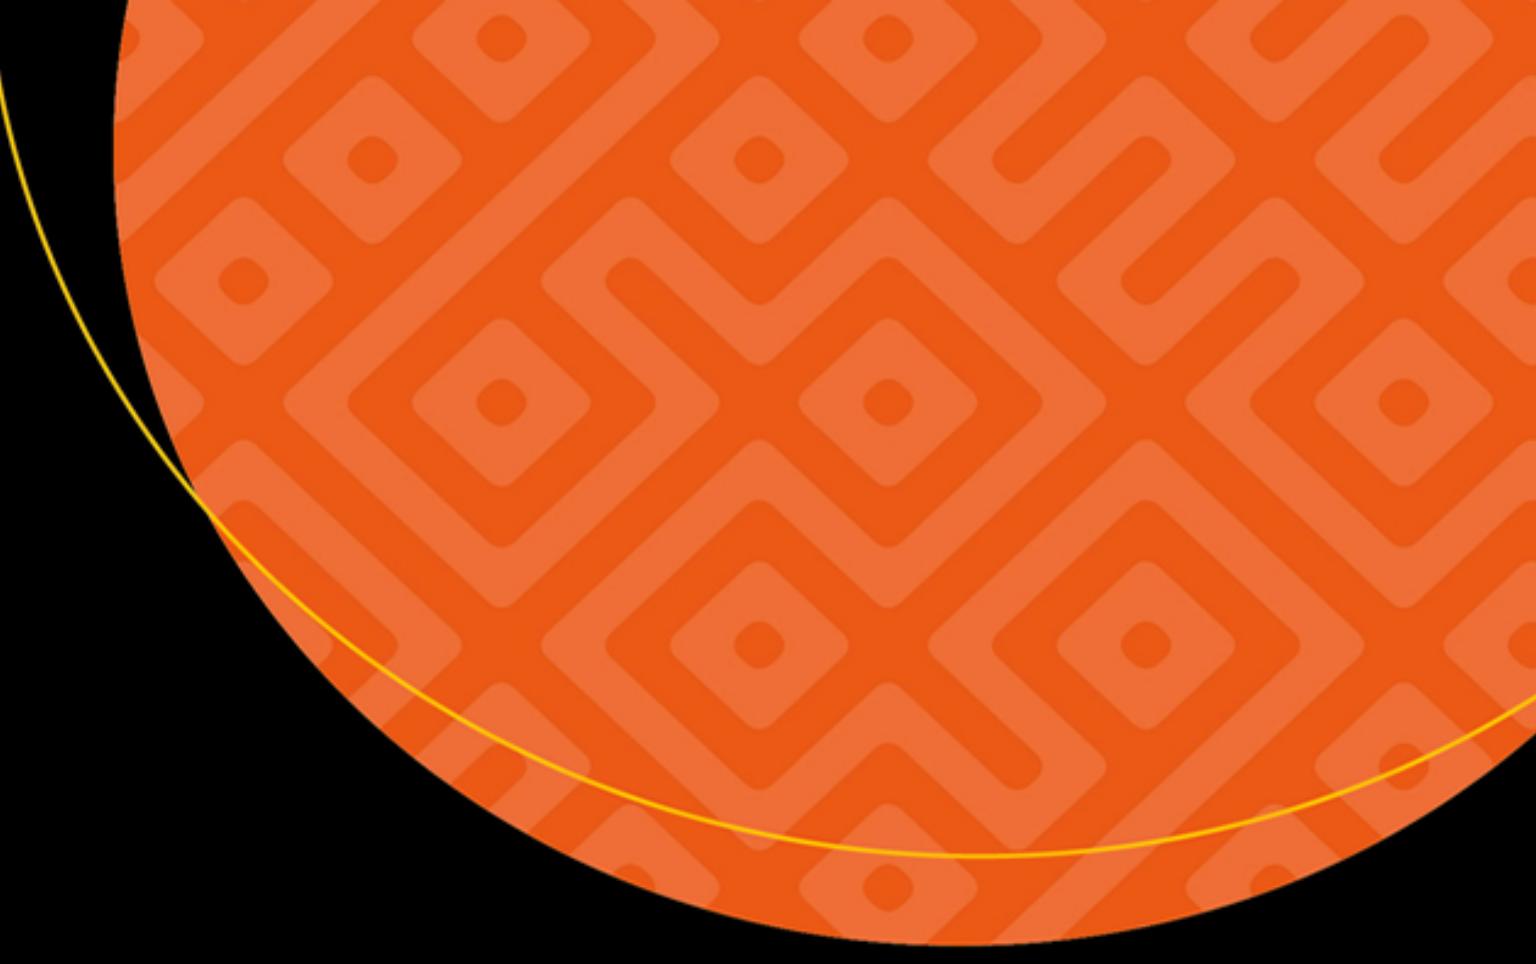


state of the art (arxiv)

- **fb/instagram: 1905.00546**
- **randaugment: 1909.13719**
- **efficientdet: 1911.09070**
- **noisy teacher: 1911.04252**



**"Mr. Osborne, may I be excused?
My brain is full."**



Convolutional Neural Networks with Swift for TensorFlow

Image Recognition and
Dataset Categorization

—
Brett Koonce

Apress®

- **apress: aaron
black, jessica
valenti, james
markham,
vishwesh shrimali**
- **notes: james maki,
brennan saeta**
- **parents,
quarkworks**

wishlist



- **checkpointing --> done!**
- **tpu training/scaling**
- **mixed precision training (bfloating16)**
- **data pipelines/augmentation**

thanks for coming!