Convolutional neural networks
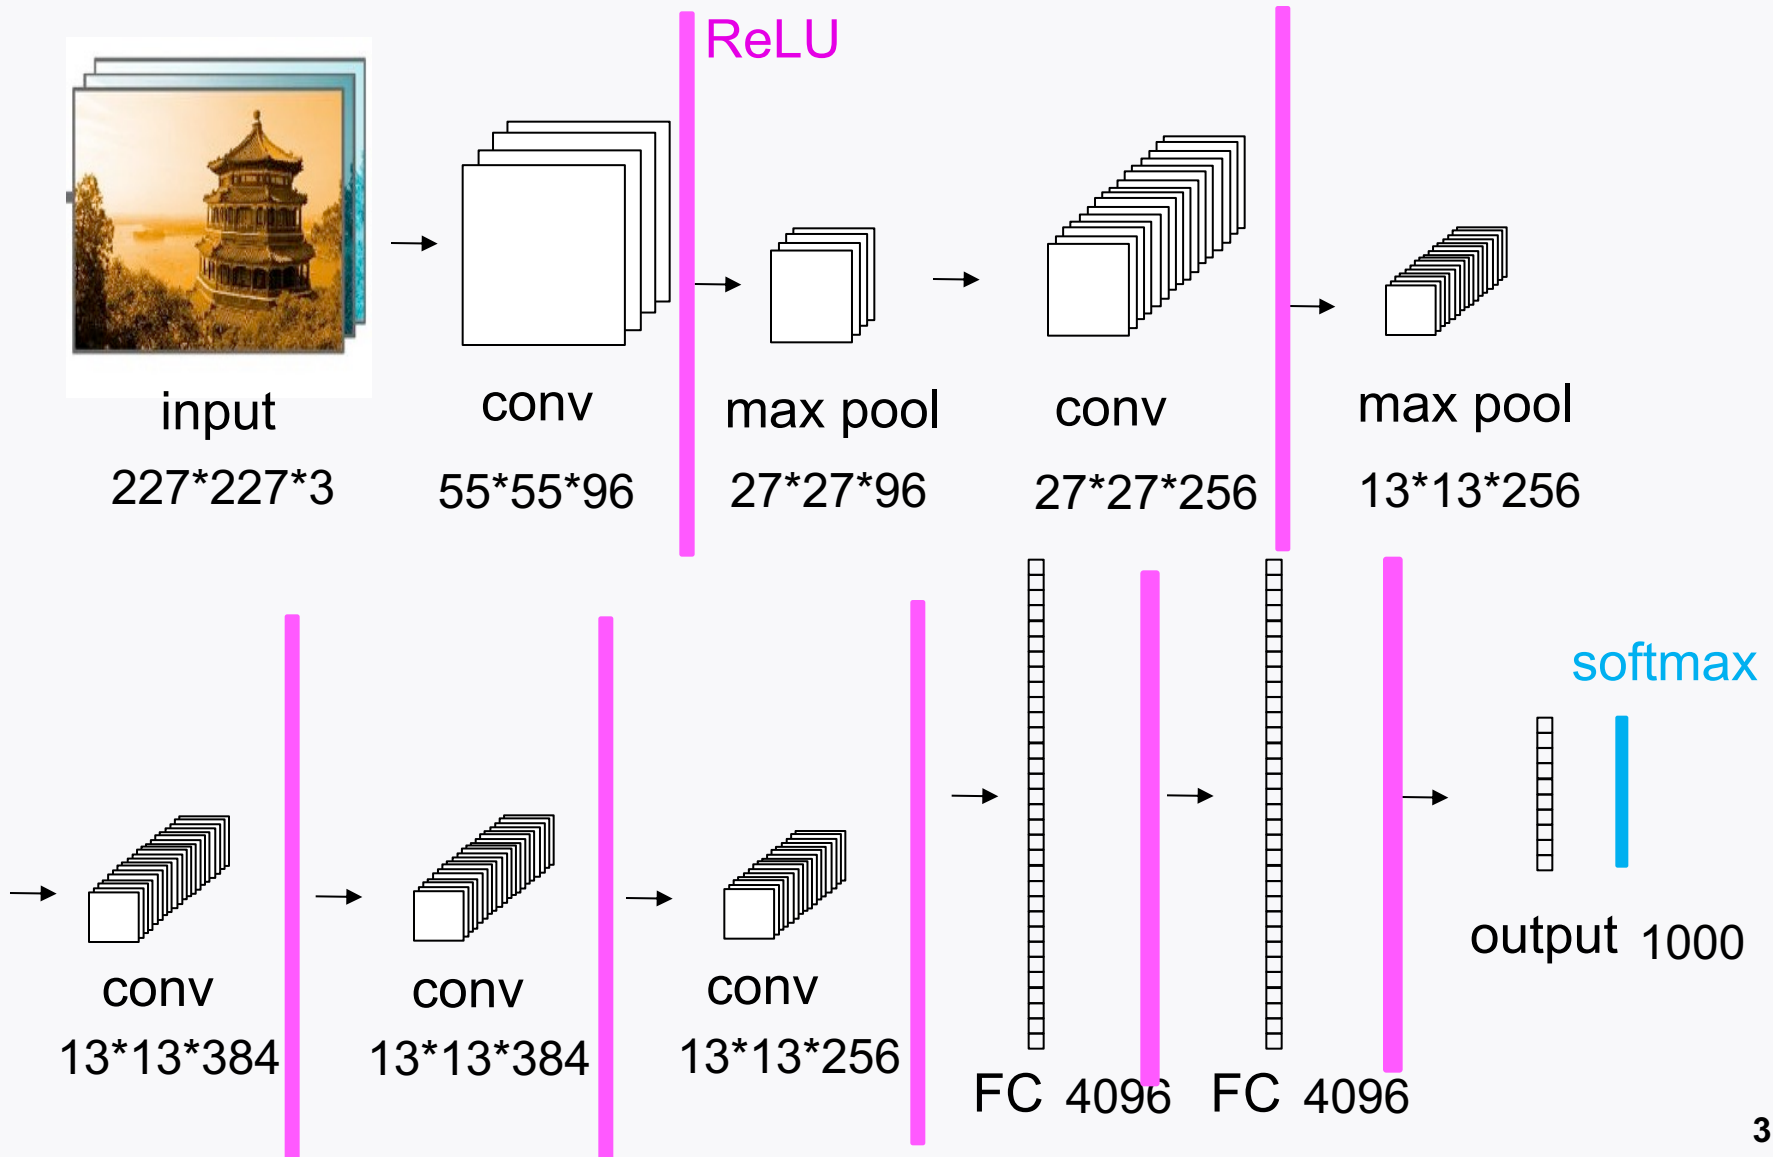
Lecture 9

Changho Suh
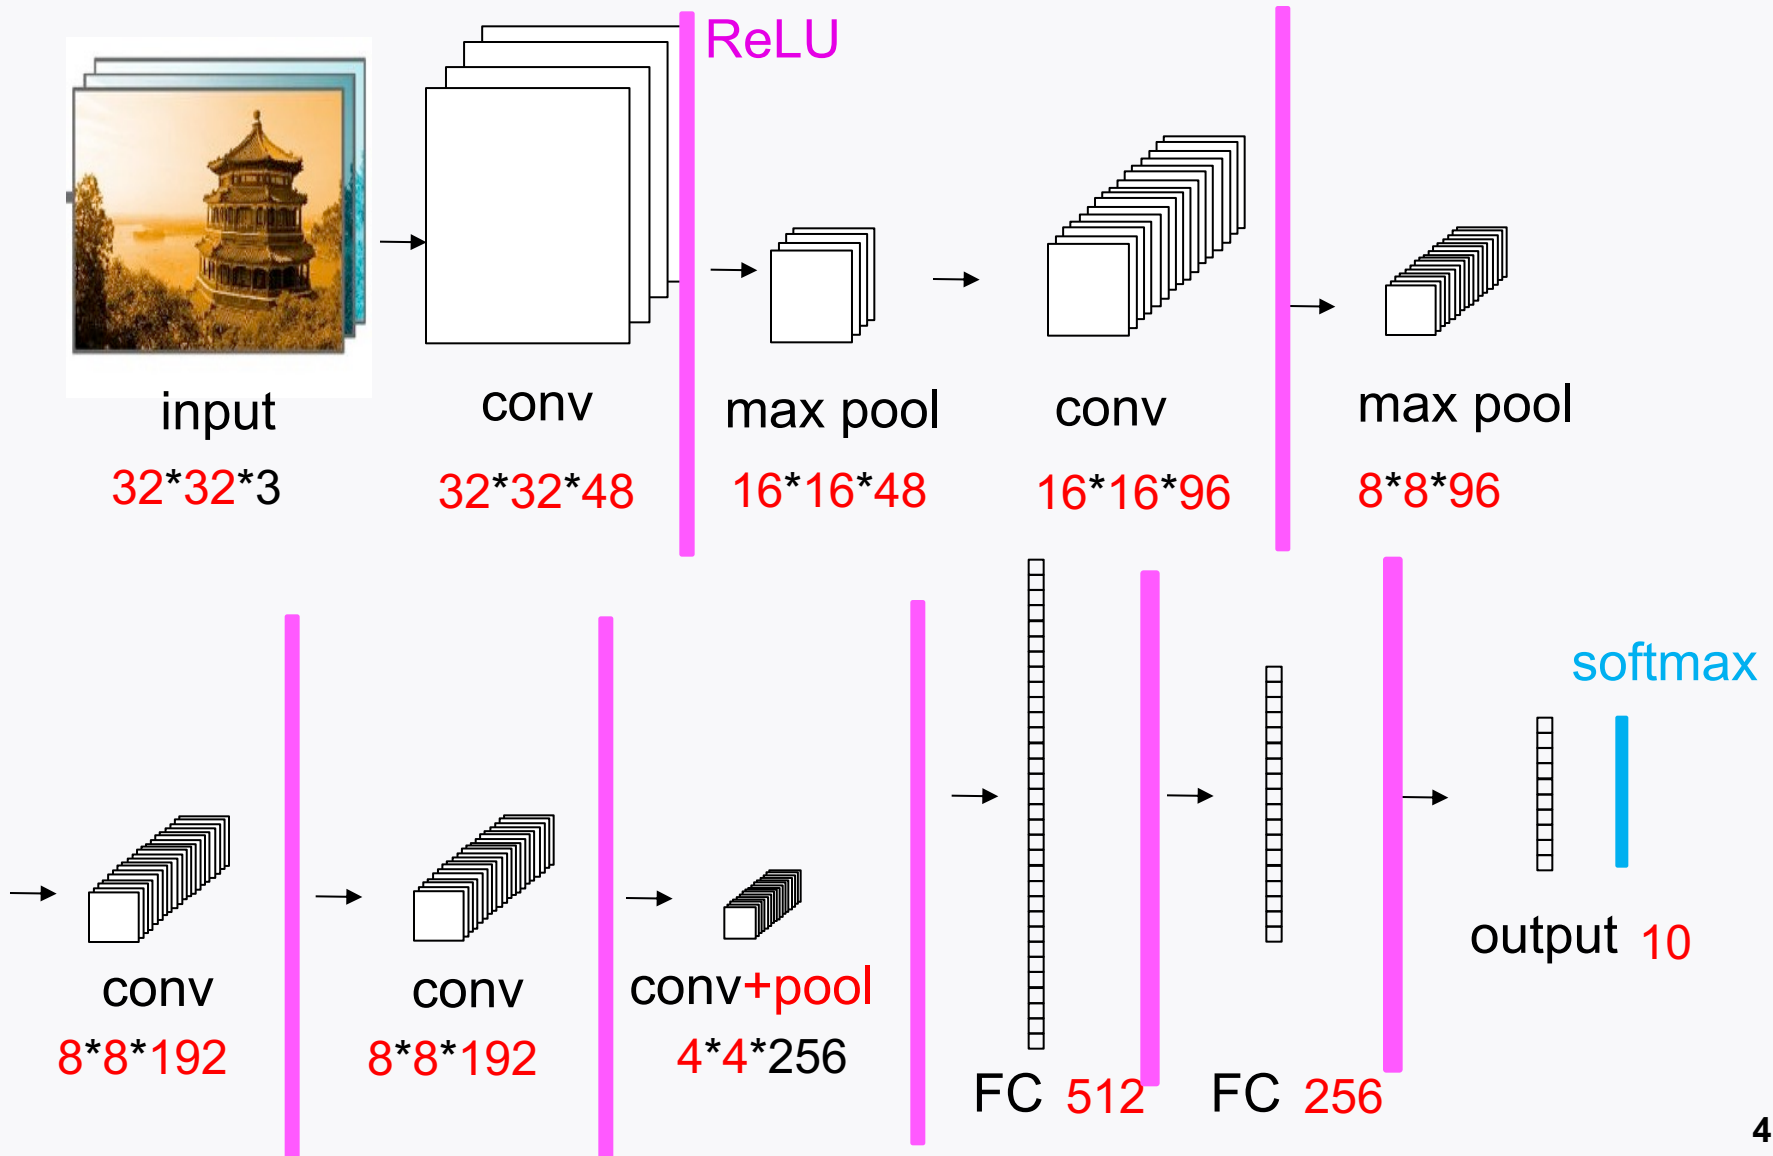
September 29, 2021

AlexNet & ResNet

Outline

Will discuss **two** popular CNN architectures.

1. **AlexNet** (by the Godfather of deep learning)

2. **ResNet** (state of the art)

AlexNet

input
227*227*3

conv
55*55*96

ReLU

max pool
27*27*96

conv
27*27*256

max pool
13*13*256

conv
13*13*384

conv
13*13*384

conv
13*13*256

FC 4096

FC 4096

softmax

output 1000

Simplified AlexNet for CIFAR10 (Check in PS)

input
32*32*3

conv
32*32*48

ReLU

max pool
16*16*48

conv
16*16*96

max pool
8*8*96

conv
8*8*192

conv
8*8*192

conv+pool
4*4*256

FC 512

FC 256

softmax

output 10
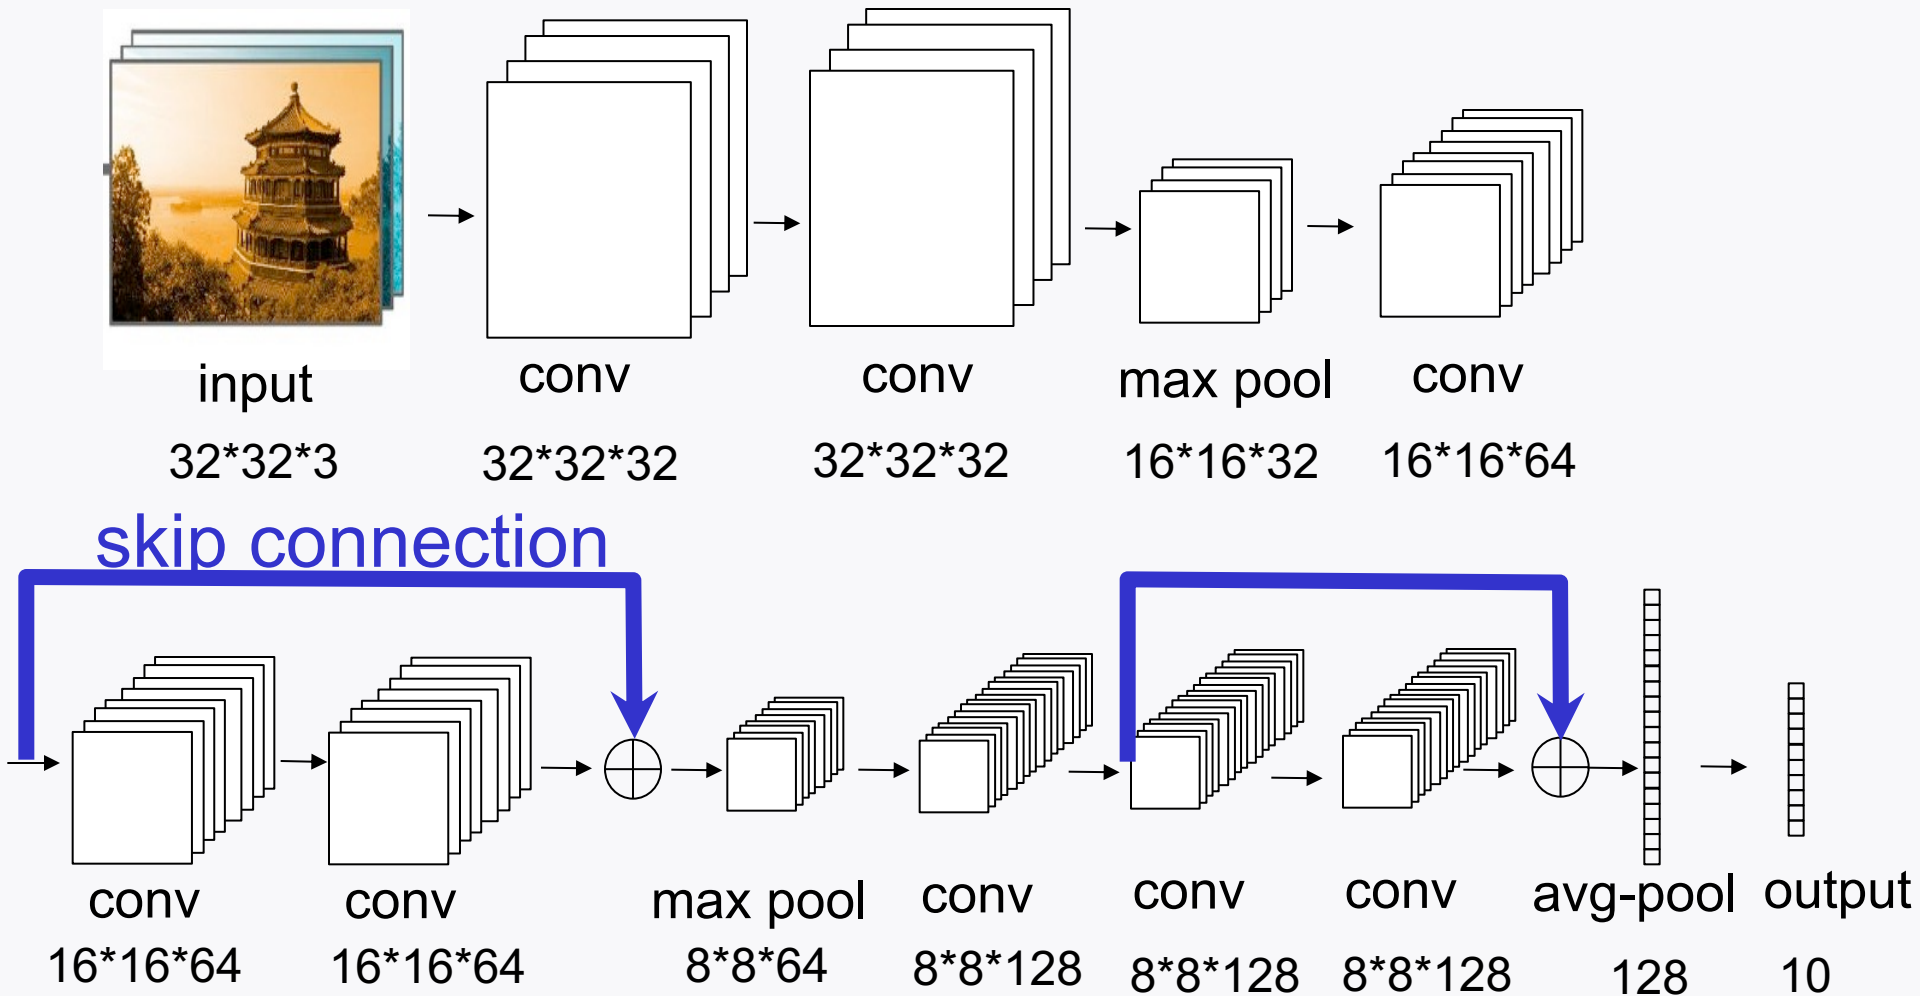
4

ResNet
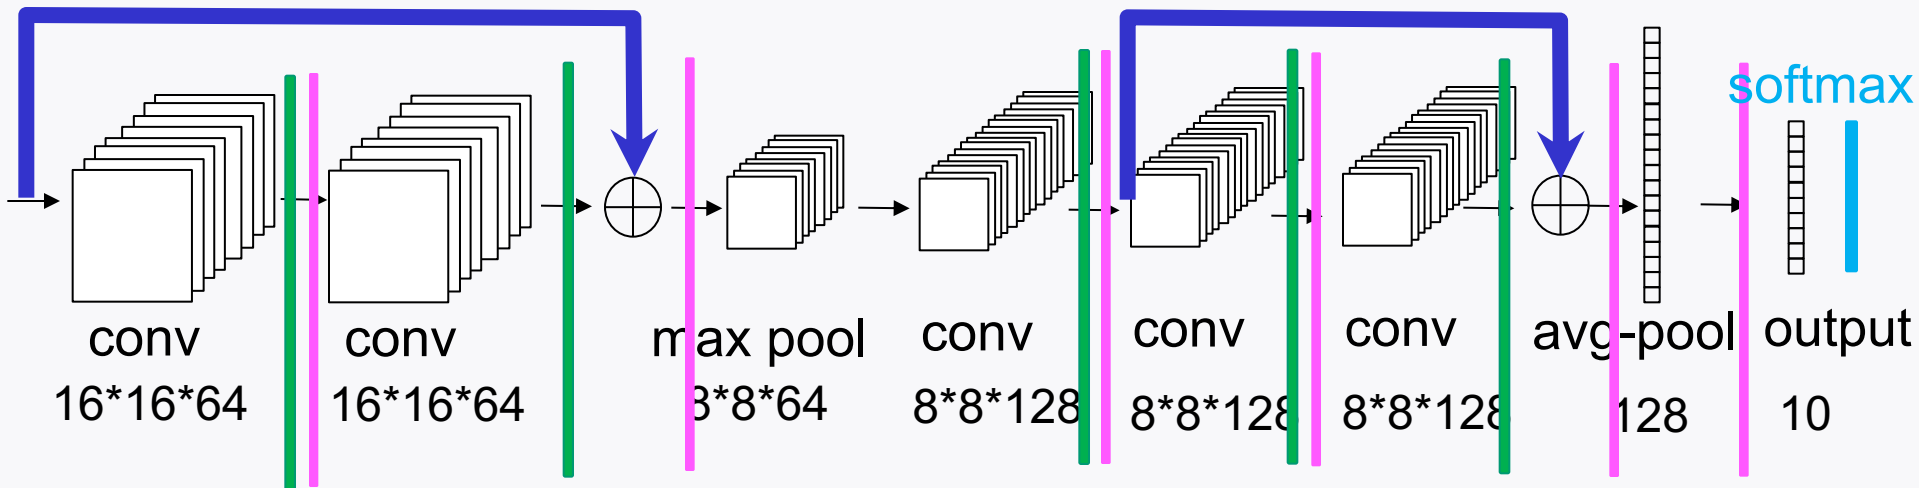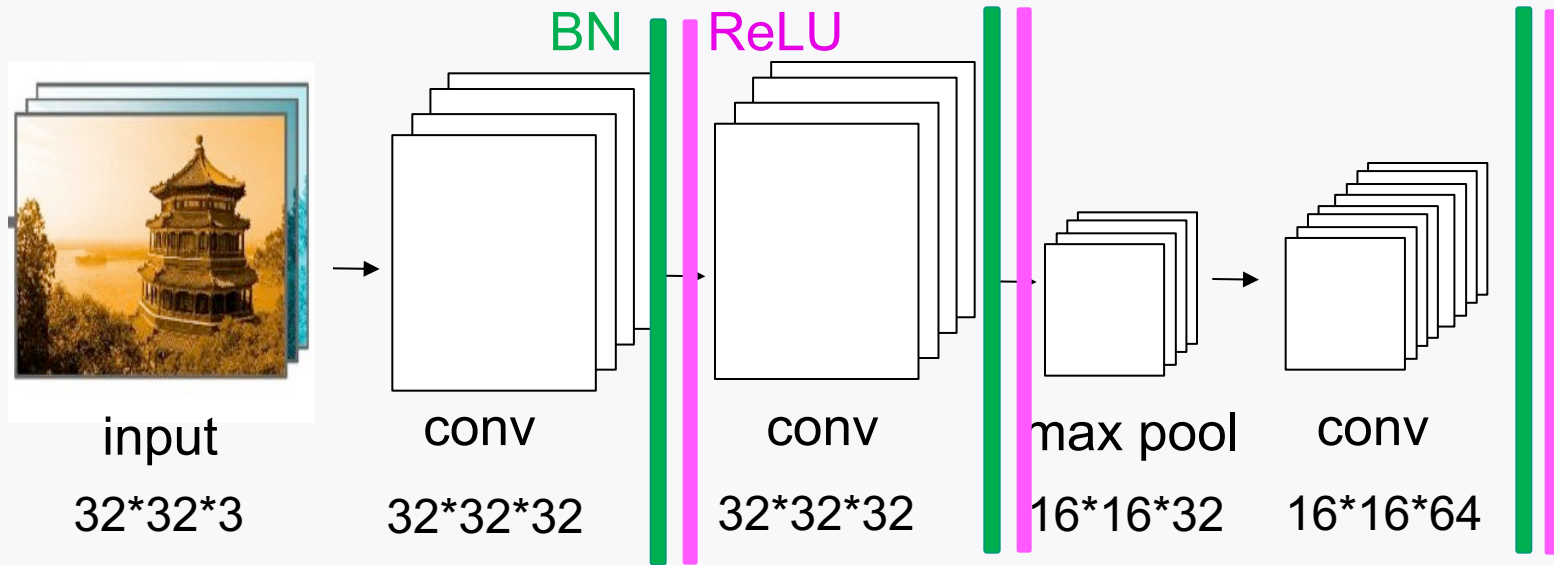
The initial version that won the 2015 ImageNet competition is complex with 152 layers.

Hence: Will explain it via a simplified version that we will also explore in PS with CIFAR10 dataset.

ResNet

input
32*32*3

conv
32*32*32

conv
32*32*32

max pool
16*16*32

conv
16*16*64

skip connection

conv
16*16*64

conv
16*16*64

max pool
8*8*64

conv
8*8*128

conv
8*8*128

conv
8*8*128

avg-pool
128

output
10

A modified ResNet that will be used in PS

Turns out …

Skip-connection plays a crucial role to enable stable & fast training!

What you will check in PS:

ResNet offers higher test accuracy than AlexNet.

Applications of CNNs

Image recognition

Image inpainting

Object detection

Coloring

Defect detection

Style transfer

Medical diagnosis
(e.g., cancer detection)

Super-resolution image
synthesis

Any **decision** or **manipulation** w.r.t. image data

Limitations

Not well applicable to time series data.

This is where recurrent neural networks (RNNs) kick in.

Outline of Day 4 lectures

1. Talk about RNN's applications and history.

2. Study two key building blocks of RNNs.

   Recurrent neurons

   A **memory cell**

3. Investigate basic RNNs.

4. Study LSTM (Long Short-Term Memory) cells.