

Deep Learning
Deep Neural Network

Yoon Joong Kim,
Hanbat National University

Deep Learning

CNN(1)

Convolutional Neural Networks

CNN(1) :
Background of Convolutional Neural Networks

CNN, Convolution Neural Network

Case study

CNN(2) :
CNN examples

Yoon Joong Kim

Department of Computer Engineering, Hanbat National University

yjkim@hanbat.ac.kr

Contents

1. Background of Convolutional Neural Networks
2. CNN, Convolution Neural Network
 1. Convolution
 2. Channel
 3. filter, kernel, stride, feature map and activation map
 4. Padding
 5. pooling layer
 6. 출력 레이어의 크기 계산
 7. Fully Connected Layer (FC layer)
4. Case study
 1. LeNet-5
 2. AlexNet
 3. GoogleNet
 4. ResNet
 5. Sentence Classification
 6. AlphaGo
5. CNN examples
 1. CNN 구성예
 2. Mnist digit classifier with CNN
 3. Exercise
 4. ConvNetJS demo: training on CIFAR-10
 5. Exnsemble

Recap

1. Application & tips


1. Learning rate
2. Data processing for the gradient algorithm
3. Overfitting
4. Optimizer
5. Application examples
 1. Dataset normalization
 2. MNIST digit classifier
 3. Application Tips for the MNIST digit classifier

2. Activations(new) and gradient vanishment


Activation Functions

Sigmoid


$$\sigma(x) = 1/(1 + e^{-x})$$




tanh tanh(x)



ReLU max(0,x)




Leaky ReLU max(0.1x, x)



Maxout max(w1^T x + b1, w2^T x + b2)

$$f(x) = \begin{cases} x & \text{if } x > 0 \\ \alpha (\exp(x) - 1) & \text{if } x \leq 0 \end{cases}$$



```
model=Sequential()
model.add(Dense(units=5,input_dim=X.shape[1],
                activation='relu'))
model.add(Dense(units=5,activation='relu'))
model.add(Dense(units=5,activation='relu'))
model.add(Dense(units=5,activation='relu'))
model.add(Dense(units=5,activation='relu'))
model.add(Dense(units=5,activation='relu'))
model.add(Dense(units=5,activation='relu'))
model.add(Dense(units=5,activation='relu'))
model.add(Dense(units=5,activation='relu'))
model.add(Dense(units=1,activation='sigmoid'))
```

```
Epoch 1999/2000
4/4 [=====] - 0s 2ms/step - loss: 0.0107
Epoch 2000/2000
4/4 [=====] - 0s 2ms/step - loss: 0.0107

## Evaluation ##
## Predict :
[[0.14624415]
 [0.99340045]
 [1.
 [0.14624415]]

## Accuracy : 1.0
```

```
model=Sequential()
model.add(dense(units=5, input_dim=2,
                activation='sigmoid')) #1st hidden layer
model.add(dense(units=5, activation='sigmoid'))#2nd hidden layer
model.add(dense(units=5, activation='sigmoid'))#3rd hidden layer
model.add(dense(units=5, activation='sigmoid'))#4th hidden layer
model.add(dense(units=5, activation='sigmoid'))#5th hidden layer
model.add(dense(units=5, activation='sigmoid'))#6th hidden layer
model.add(dense(units=5, activation='sigmoid'))#7th hidden layer
model.add(dense(units=5, activation='sigmoid'))#8th hidden layer
model.add(dense(units=5, activation='sigmoid'))#9th hidden layer
model.add(dense(units=1, activation='sigmoid'))#output layer
```

```
Epoch 1999/2000
4/4 [=====] - 0s 2ms/step - loss: 0.2500
Epoch 2000/2000
4/4 [=====] - 0s 2ms/step - loss: 0.2500

## Evaluation ##
## Predict :
[[0.5]
 [0.5]
 [0.5]
 [0.5]]

## Accuracy : 0.5
```

1. Background of Convolutional Neural Networks

A bit of history:

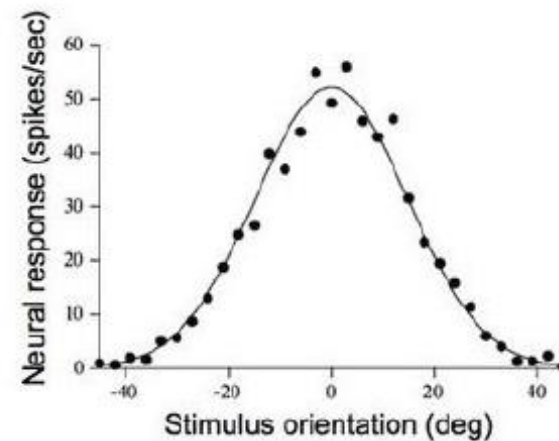


**Hubel & Wiesel,
1959**

RECEPTIVE FIELDS OF SINGLE
NEURONES IN
THE CAT'S STRIATE CORTEX

1962

RECEPTIVE FIELDS, BINOCULAR
INTERACTION
AND FUNCTIONAL ARCHITECTURE IN
THE CAT'S VISUAL CORTEX


1968...



1. Background of Convolutional Neural Networks

Case Study: LeNet-5

[LeCun et al., 1998]

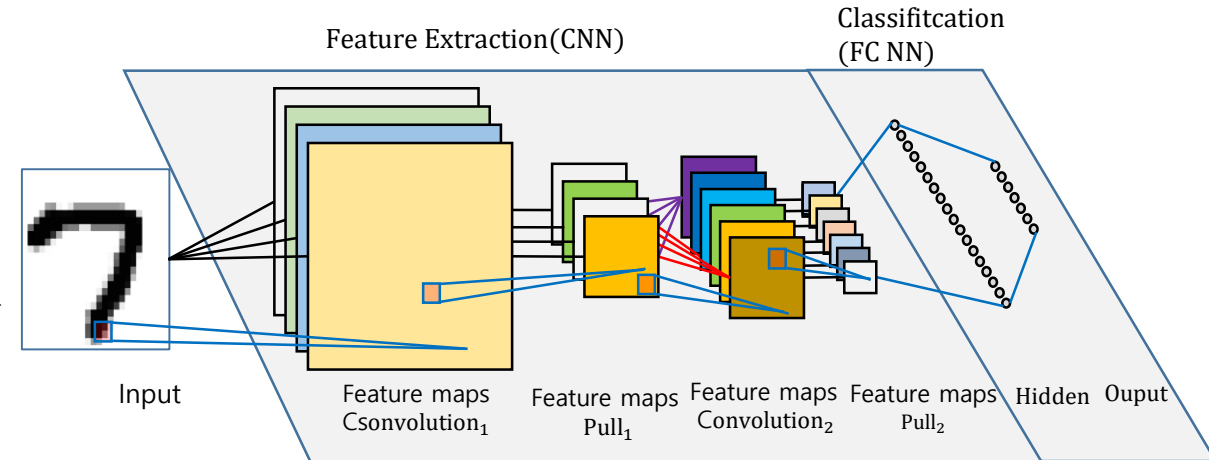


Conv filters were 5x5, applied at stride 1
Subsampling (Pooling) layers were 2x2 applied at stride 2
i.e. architecture is [CONV-POOL-CONV-POOL-CONV-FC]

2. CNN, Convolution Neural Network

- 개요
 - 전연결신경망(FC NN, Fully Connected Neural Network)을 이용한 이미지 분류의 경우 3차원 이미지를 1차원으로 평면화 하여야 한다. 이미지 공간 정보 유실로 인한 정보 부족으로 인공 신경망이 특징을 추출 및 학습이 비효율적이고 정확도를 높이는 데 한계가 있습니다. 이미지의 공간 정보를 유지한 상태로 학습이 가능한 모델이 바로 CNN(Convolutional Neural Network)입니다.

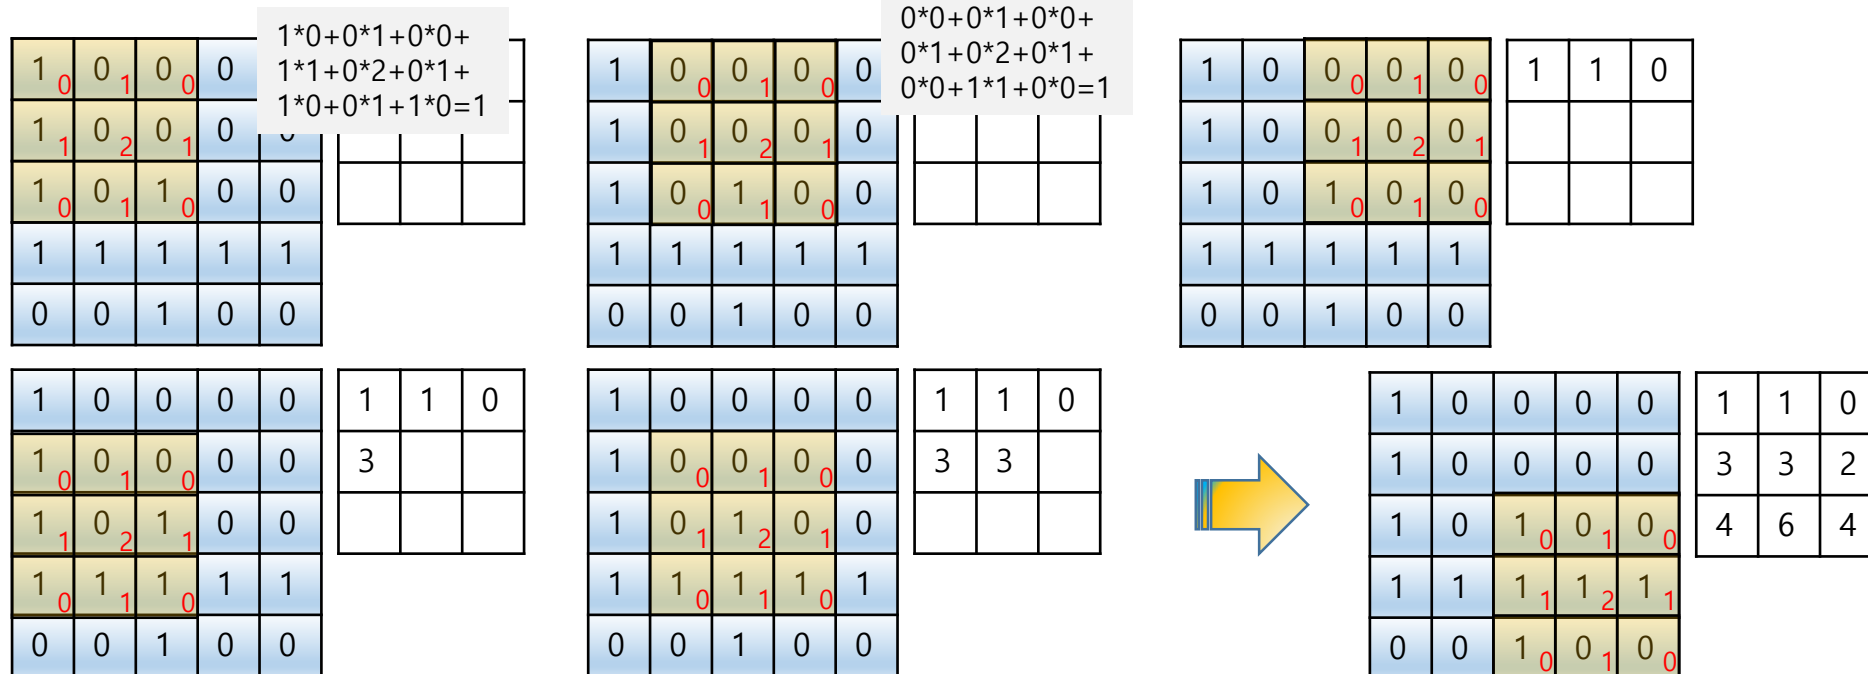
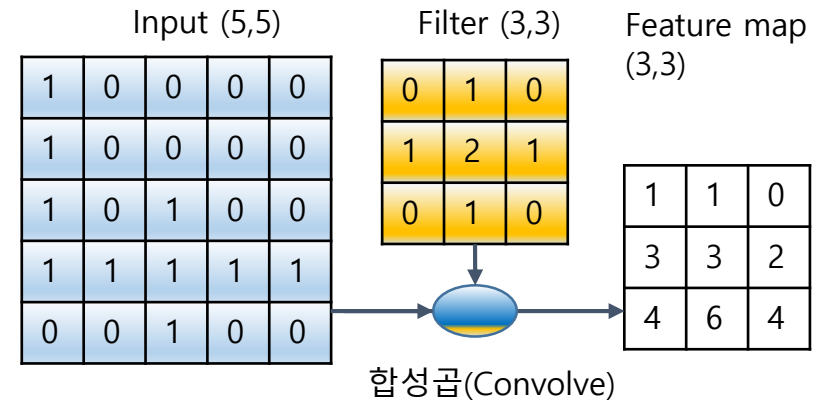
- FC NN에 대한 CNN의 차이점
 - 각 레이어의 입출력 데이터의 형상 유지
 - 이미지의 공간 정보를 유지하면서 인접 이미지와의 특징을 효과적으로 인식
 - 복수의 필터로 이미지의 특징 추출 및 학습
 - 추출한 이미지의 특징을 모으고 강화하는 Pooling 레이어
 - 필터를 공유 파라미터로 사용하기 때문에, 일반 인공 신경망과 비교하여 학습 파라미터가 매우 적음



- 1. 주요용어
 - 합성곱 (Convolution)
 - 채널 (Channel)
 - 필터 (Filter)
 - 커널 (Kernel)
 - 스트라이드 (Stride)
 - 패딩 (Padding)
 - 피쳐맵 (Feature map)
 - 액티베이션 맵 (Activation Map)
 - 풀링 레이어 (Pooling layer)

2.1 convolution

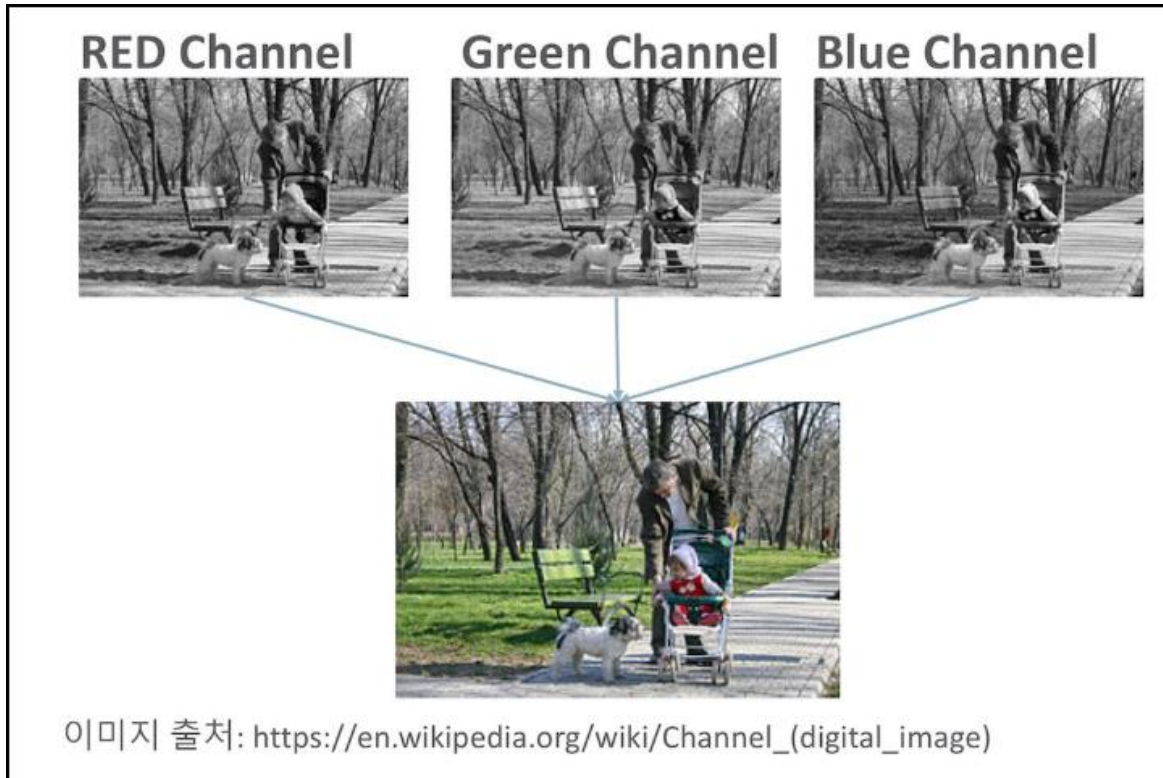
- 2.1 합성곱(Convolution)
 - 합성곱은 하나의 함수와 또 다른 함수를 반전 이동한 값을 곱한 다음, 구간에 대해 적분하여 새로운 함수를 구하는 수학 연산자이다. [\[link\]](#)



2.2 channel

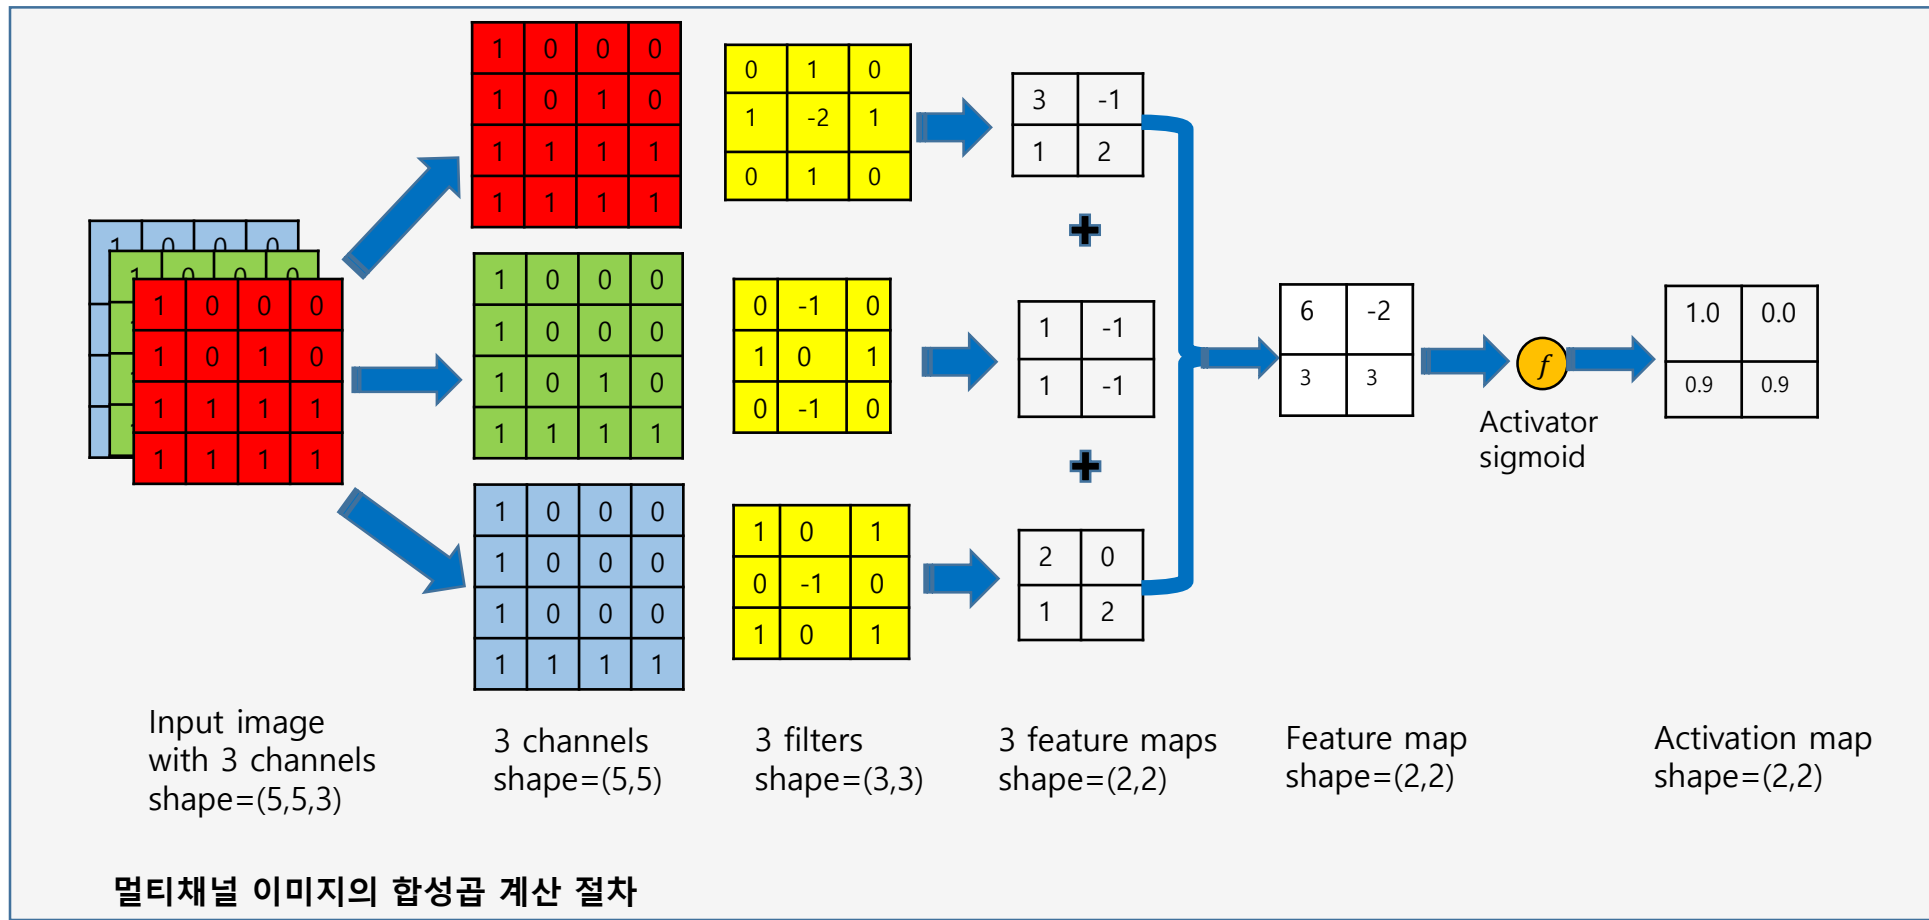
- 2.2 채널, channel

- 컬러 이미지 픽셀의 색상은 Red, Green, Blue의 크기로 합성됩니다. 따라서 컬러사의 이미지는 3장의 채널로 구성된다. 컬러 이미지의 shape=(32,32,3) 흑백 이미지는 shape=(32,32,1)이 된다.



2.3 filter, kernel, stride, feature map and activation map

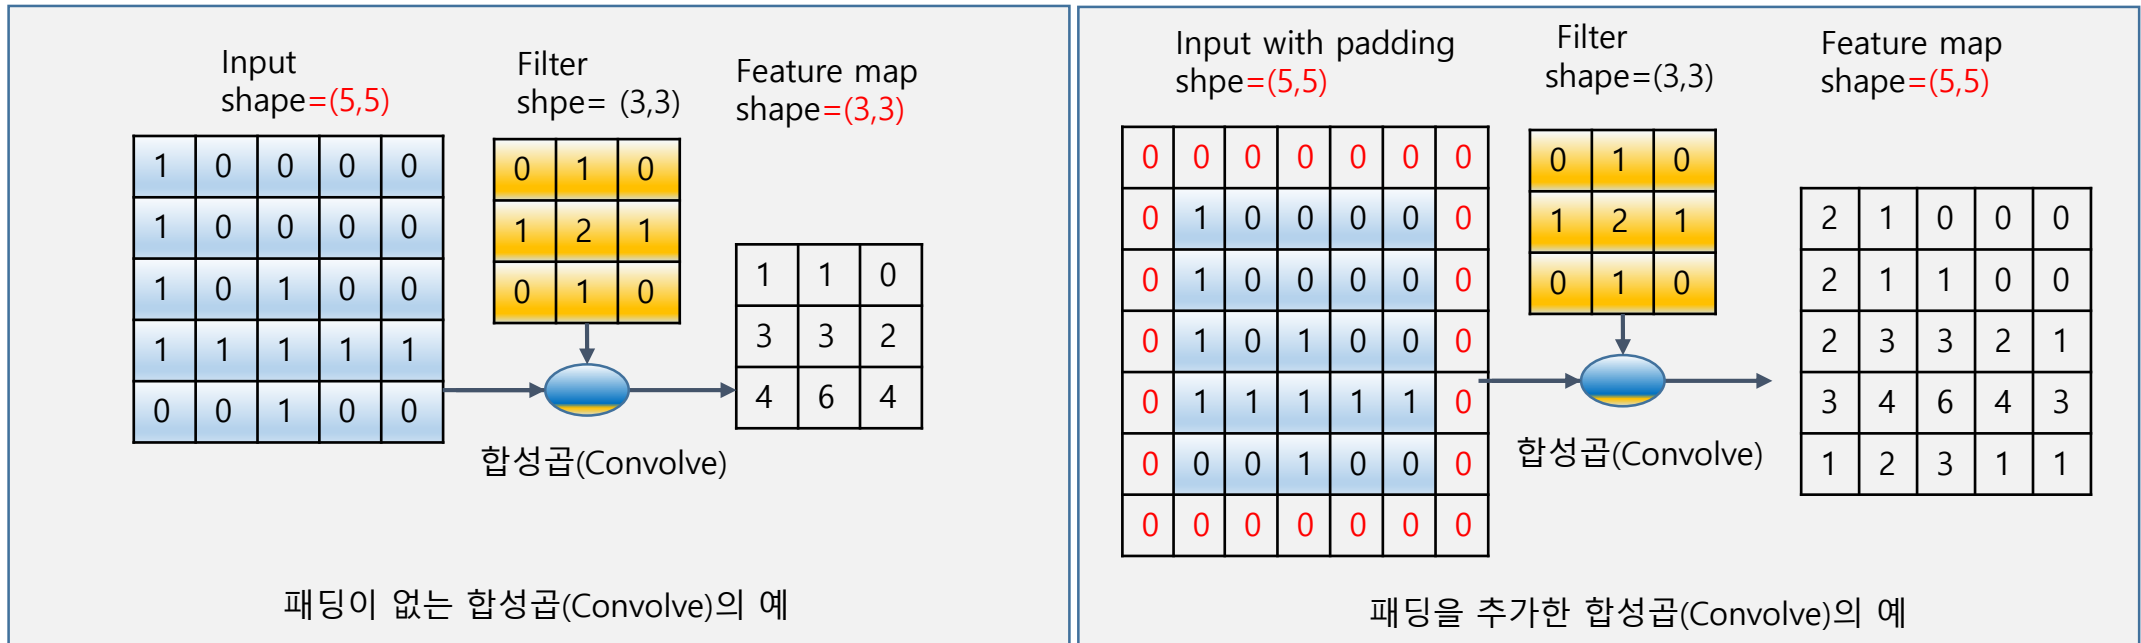
- 2.3 합성곱 : 필터, 스트라이드, 피쳐맵(feature map), 커널(kernel)
 - 커널 : 필터와 동일, 스트라이드는 필터의 이동 픽셀 수, 피쳐맵 : 합성곱의 계산결과 이미지
 - 액티베이션 맵 : 피쳐맵에 액티베이션을 적용한 결과



2.4 padding


- 2.4 패딩

- 합성곱을 하면 입력에 비하여 피쳐맵의 사이이즈가 작아진다. 같은 크기가 되도록 입력의 가장자리에 0을 채우는 과정을 말한다.

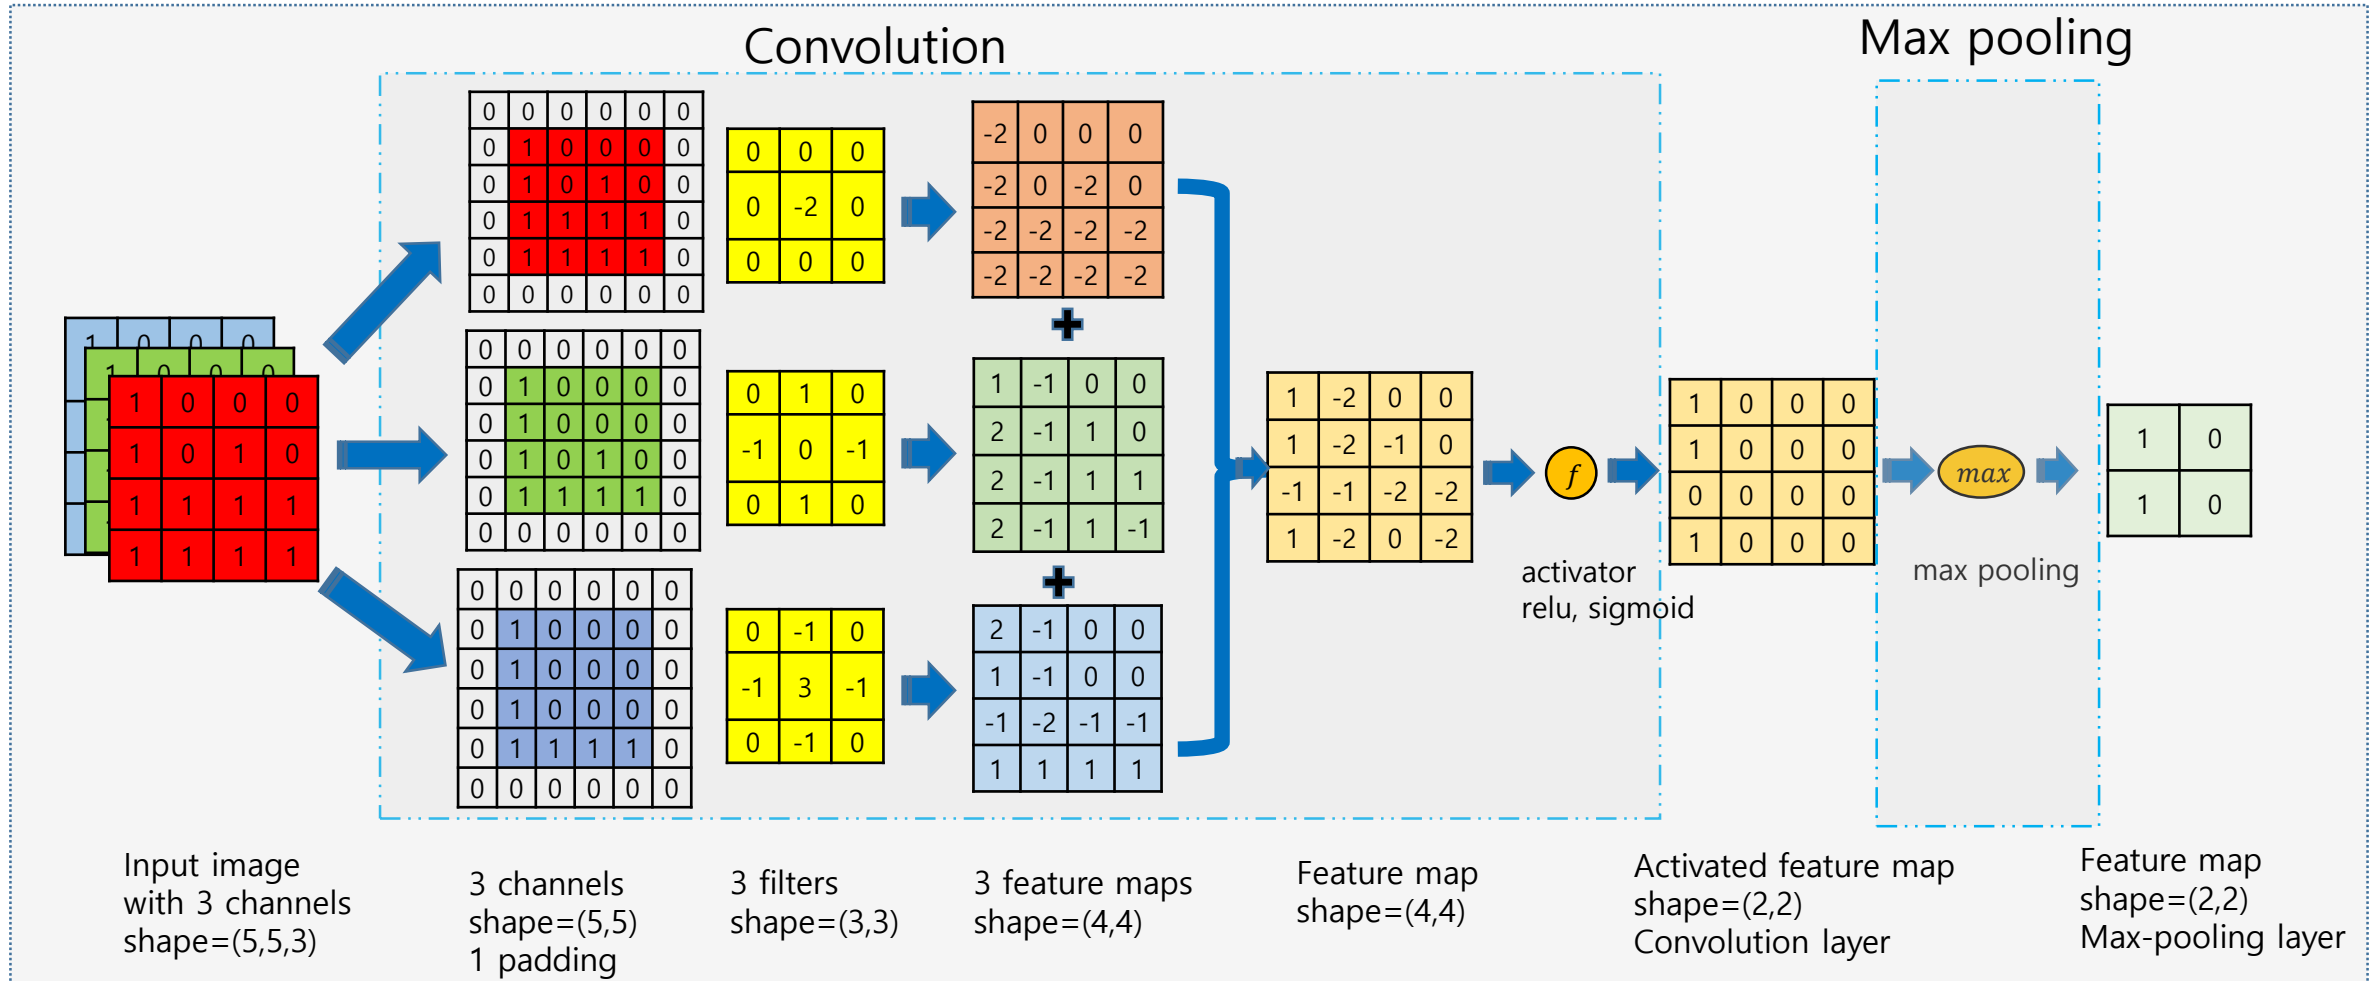


2.5 pooling layer

- 2.5 pooling layer
 - Convolution layer의 출력(activation map)의 일정 값은 강조하거나 크기를 줄이는 목적으로 사용된다. Pooling 방법은 max, min, average pooling이 있다. 일반적으로 pooling size와 stride size로 한다.
 - Pooling 특징
 - 학습대상 파라미터가 없음
 - Pooling 레이어를 통과하면 행렬의 크기 감소
 - Pooling 레이어를 통해서 채널 수 변경 없음



2.5 pooling layer



멀티채널 이미지의 합성곱 layer 및 max pooling layer 계산 예

2.6 출력 레이어의 크기 계산

- 2.6 레이어의 출력 크기 계산

- 입력 데이터 높이: H
- 입력 데이터 폭: W
- 필터 높이 : FH
- 필터 폭 : FW
- Stride 크기 : S
- 패딩 사이즈 : P

- Convolution layer 출력의 크기

- $OutputHeight = oh = \frac{H+2P-FH}{S} + 1$

- $OutputWidth = ow = \frac{H+2P-FW}{S} + 1$

- Pooling layer 의 출력 크기

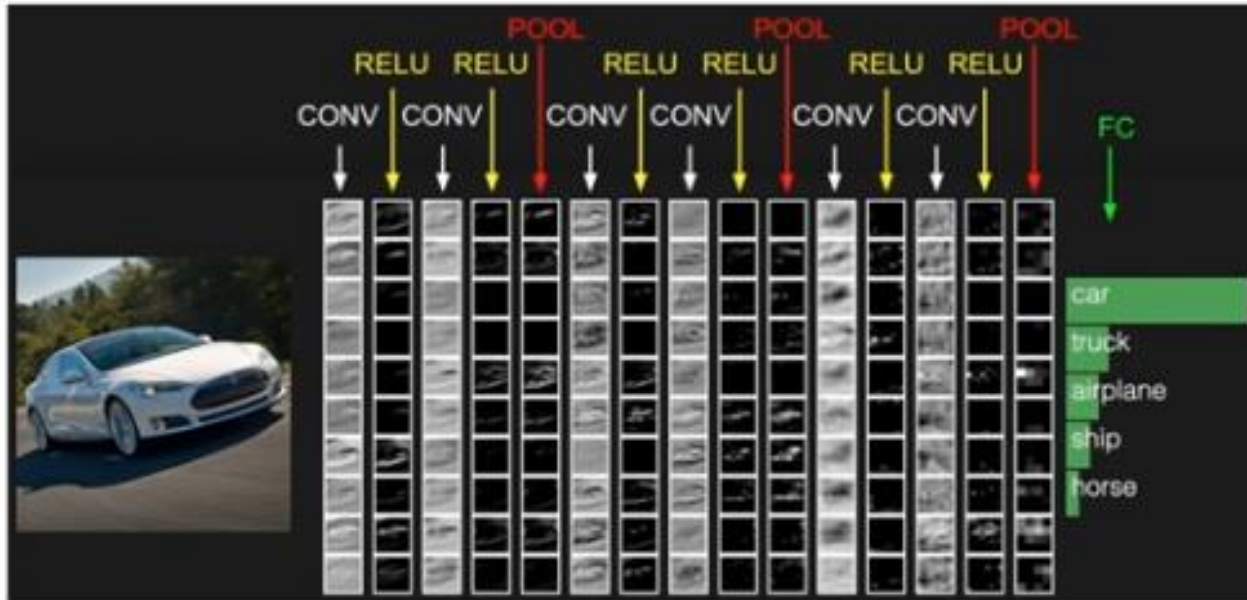
- $OutputRowSize = \frac{InputRowSize}{PoolingSize}$

- $OutputColumnSize = \frac{InputColumnSize}{PoolingSize}$

2.7 Fully Connected Layer (FC layer)

Fully Connected Layer (FC layer)

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




3. Case study

1. LeNet-5
2. AlexNet
3. GoogleNet
4. ResNet
5. AlphaGo

3.1 LeNet-5

Case Study: LeNet-5

[LeCun et al., 1998]



Conv filters were 5×5 , applied at stride 1
Subsampling (Pooling) layers were 2×2 applied at stride 2
i.e. architecture is [CONV-POOL-CONV-POOL-CONV-FC]


3.1 LeNet-5 (cont.)

Case Study: AlexNet

[Krizhevsky et al. 2012]

ImageNet large Scale Visual Recognition Challenge

ILSVRC 2012 top-5 error 15.3%
10.8% 개선



Input: 227x227x3 images

First layer (CONV1): 96 11x11 filters applied at stride 4

=>


Output volume [55x55x96]

Parameters: $(11*11*3)*96 = 35K$

3.2 AlexNet

Case Study: AlexNet

[Krizhevsky et al. 2012]




Input: 227x227x3 images
After CONV1: 55x55x96

Second layer (POOL1): 3x3 filters applied at stride 2
Output volume: 27x27x96
Parameters: 0!

3.2 AlexNet(cont.)

Case Study: AlexNet

[Krizhevsky et al. 2012]



Input: 227x227x3 images

After CONV1: 55x55x96

After POOL1: 27x27x96

...

3.2 AlexNet (cont.)

Case Study: 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)



3.3 AlexNet(cont.)

Case Study: 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%

3.3 GoogleNet

Case Study: GoogLeNet

[Szegedy et al., 2014]




Inception module

ILSVRC 2014 winner (6.7% top 5 error)

3.4 ResNet


Case Study: ResNet [He et al., 2015] ILSVRC 2015 winner (3.6% top 5 error)



MSRA @ ILSVRC & COCO 2015 Competitions

- **1st places in all five main tracks**
 - ImageNet Classification: “Ultra-deep” (quote Yann) **152-layer nets**
 - ImageNet Detection: **16%** better than 2nd
 - ImageNet Localization: **27%** better than 2nd
 - COCO Detection: **11%** better than 2nd
 - COCO Segmentation: **12%** better than 2nd

*Improvements are relative numbers



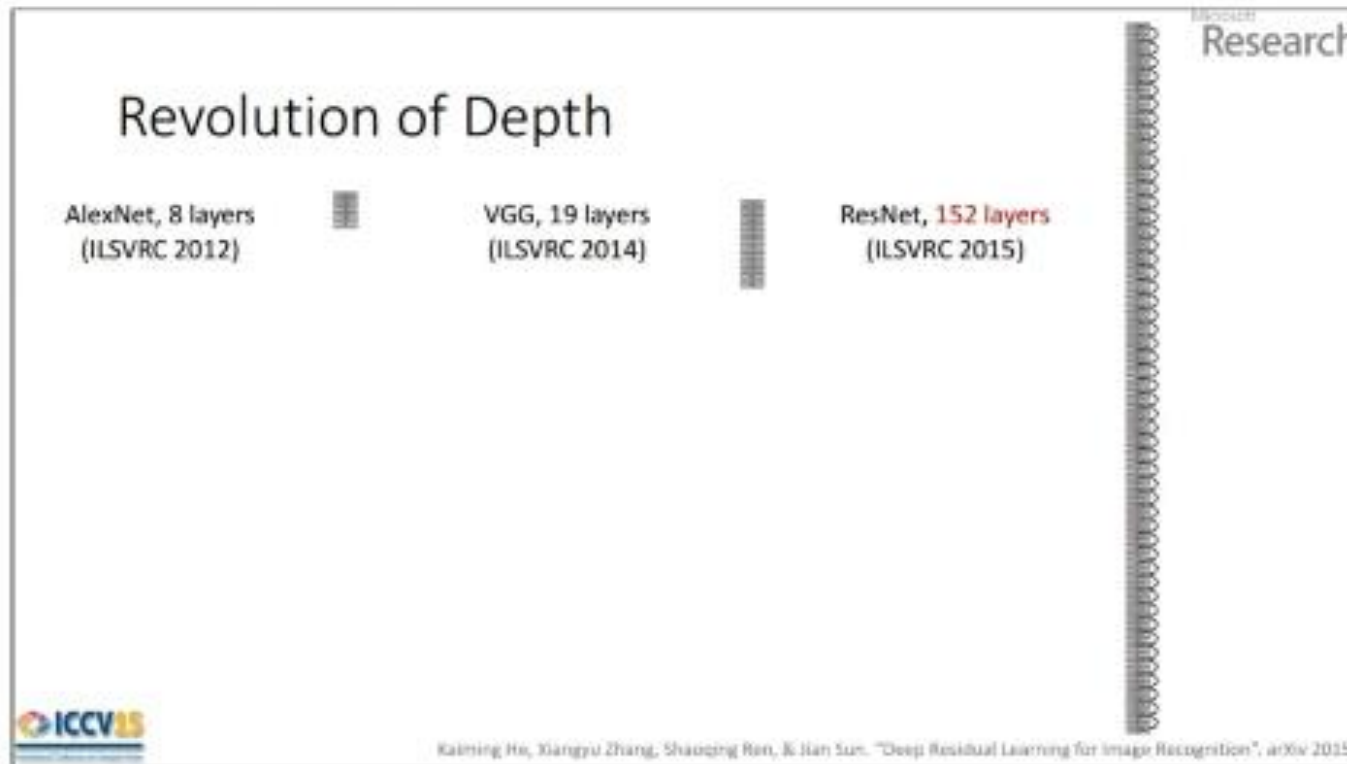
Kaiming He, Xiangyu Zhang, Shaoqing Ren, & Jian Sun. “Deep Residual Learning for Image Recognition”. arXiv 2015.

Slide from Kaiming He’s recent presentation <https://www.youtube.com/watch?v=1PGLj-uKT1w>

3.4 ResNet (cont.)

Case Study: ResNet [He et al., 2015]

ILSVRC 2015 winner (3.6% top 5 error)



2-3 weeks of training on 8 GPU machine


at runtime: faster than a VGGNet! (even though it has 8x more layers)

(slide from Kaiming He's recent presentation)

3.4 ResNet (cont.)


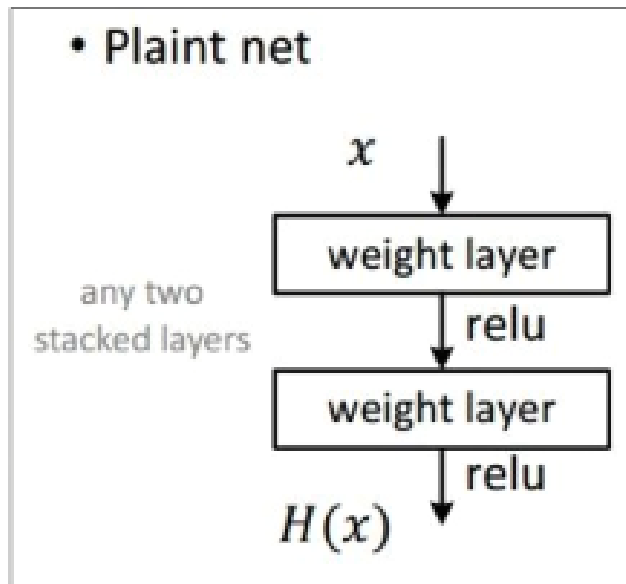
Case Study: ResNet

[He et al., 2015]



3.4 ResNet (cont.)

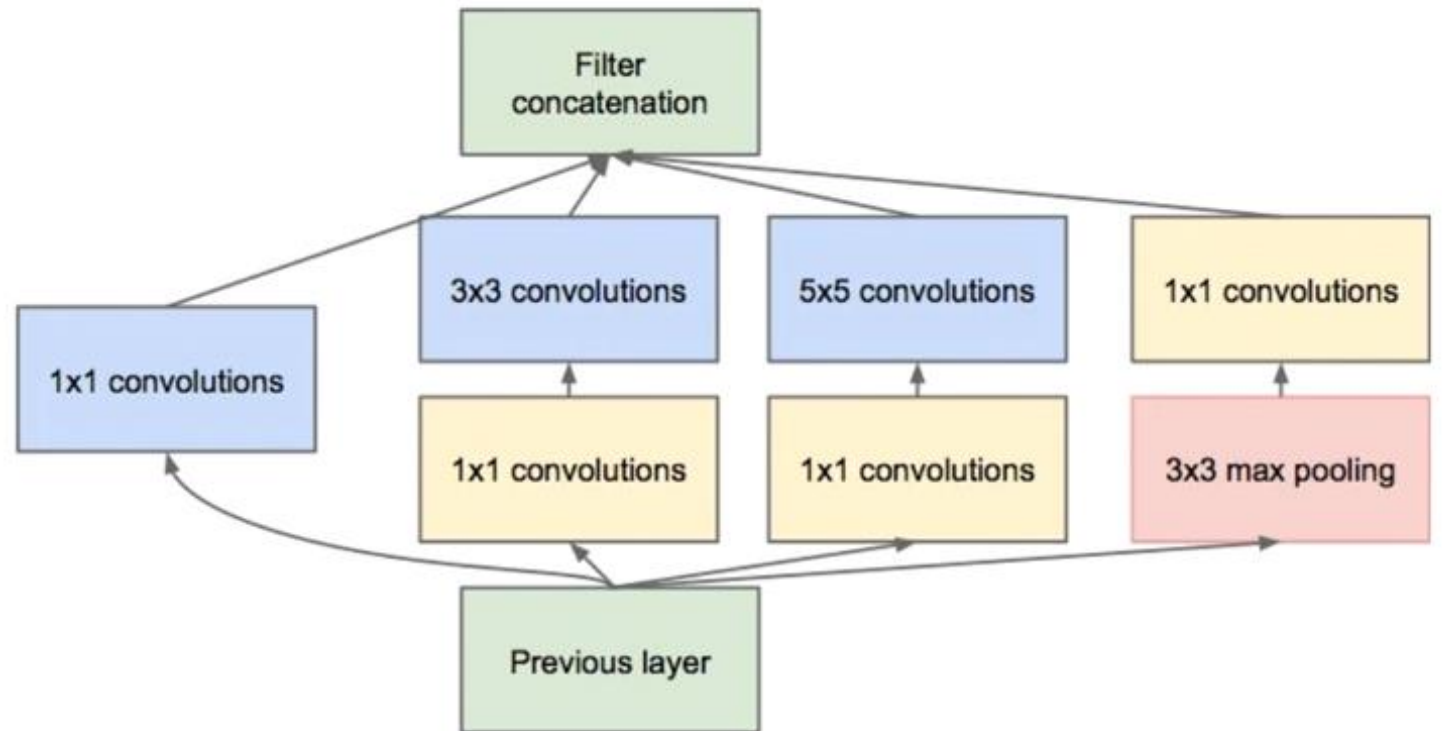

Case Study: ResNet [He et al., 2015]

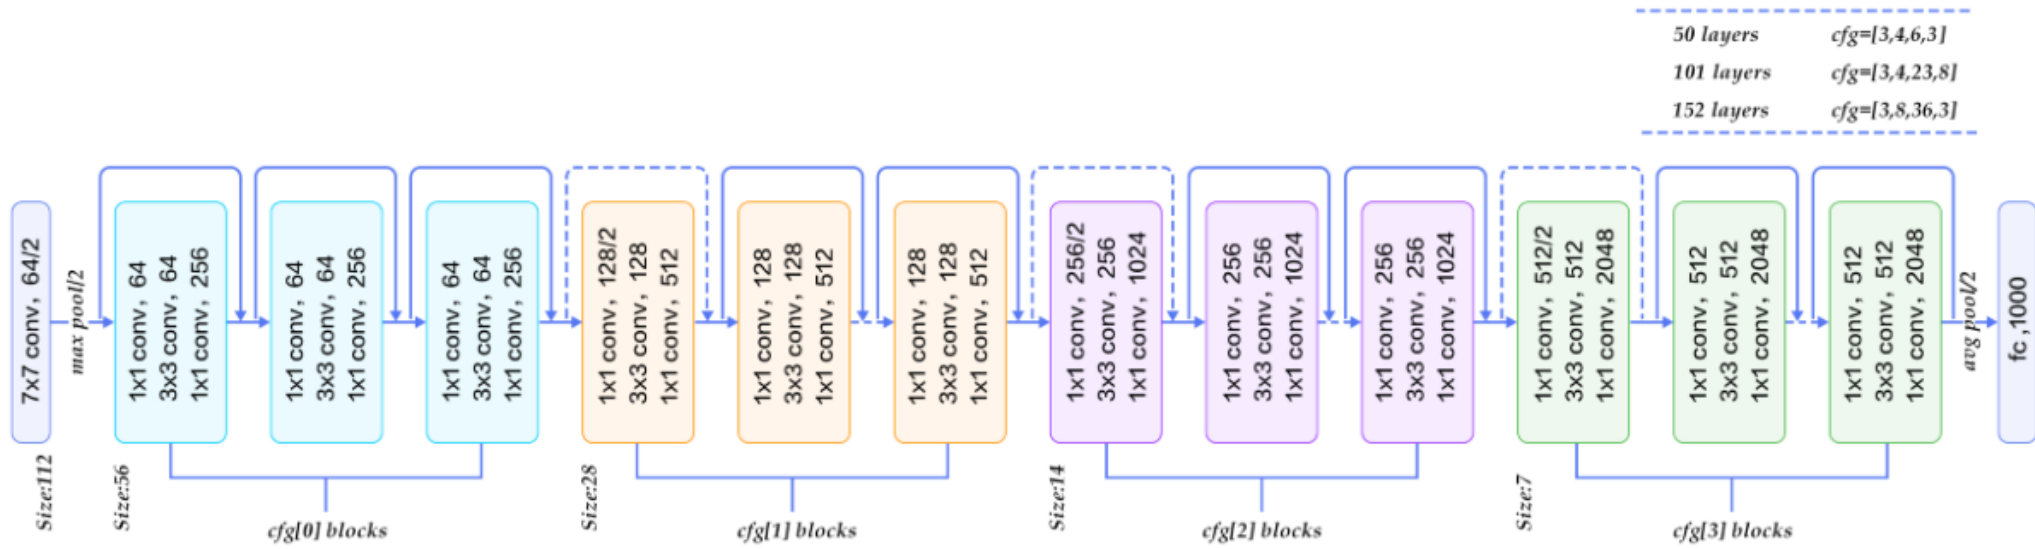


3.4 ResNet (cont.)

Case Study: ResNet



[He et al., 2015]





Summary of CNN case

Year	CNN	Developed by	Place	Top-5 error rate	No. of parameters
1998	LeNet(8)	Yann LeCun et al			60 thousand
2012	AlexNet(7)	Alex Krizhevsky, Geoffrey Hinton, Ilya Sutskever	1st	15.3%	60 million
2013	ZFNet()	Matthew Zeiler and Rob Fergus	1st	14.8%	
2014	GoogLeNet(19)	Google	1st	6.67%	4 million
2014	VGG Net(16)	Simonyan, Zisserman	2nd	7.3%	138 million
2015	ResNet(152)	Kaiming He	1st	3.6%	



3.5 Sentence Classification

- Convolution neural networks for sentence classification [Yoon Kim, 2014]





Figure 1: Model architecture with two channels for an example sentence.

3.6 AlphaGo

Case Study Bonus: DeepMind's AlphaGo



3.6 AlphaGo (cont.)

The input to the policy network is a $19 \times 19 \times 48$ image stack consisting of 48 feature planes. The first hidden layer zero pads the input into a 23×23 image, then convolves k filters of kernel size 5×5 with stride 1 with the input image and applies a rectifier nonlinearity. Each of the subsequent hidden layers 2 to 12 zero pads the respective previous hidden layer into a 21×21 image, then convolves k filters of kernel size 3×3 with stride 1, again followed by a rectifier nonlinearity. The final layer convolves 1 filter of kernel size 1×1 with stride 1, with a different bias for each position, and applies a softmax function. The match version of AlphaGo used $k = 192$ filters; [Fig. 2b](#) and [Extended Data Table 3](#) additionally show the results of training with $k = 128, 256$ and 384 filters.

policy network:


[19x19x48] Input

CONV1: 192 5x5 filters , stride 1, pad 2 => [19x19x192]

CONV2..12: 192 3x3 filters, stride 1, pad 1 => [19x19x192]


CONV: 1 1x1 filter, stride 1, pad 0 => [19x19] (*probability map of promising moves*)

'The only limit is your imagination'



Recap - CNN

1. Background of Convolutional Neural Networks
2. CNN, Convolution Neural Network
 1. Convolution
 2. Channel
 3. filter, kernel, stride, feature map and activation map
 4. Padding
 5. pooling layer
 6. 출력 레이어의 크기 계산
 7. Fully Connected Layer (FC layer)
4. Case study
 1. LeNet-5
 2. AlexNet
 3. GoogleNet
 4. ResNet
 5. Sentence Classification
 6. AlphaGo
5. CNN examples
 1. CNN 구성예
 2. Mnist digit classifier with CNN
 3. Exercise
 4. ConvNetJS demo: training on CIFAR-10
 5. Exnsenble



Deep Learning Deep Neural Network

Yoon Joong Kim,
Hanbat National University