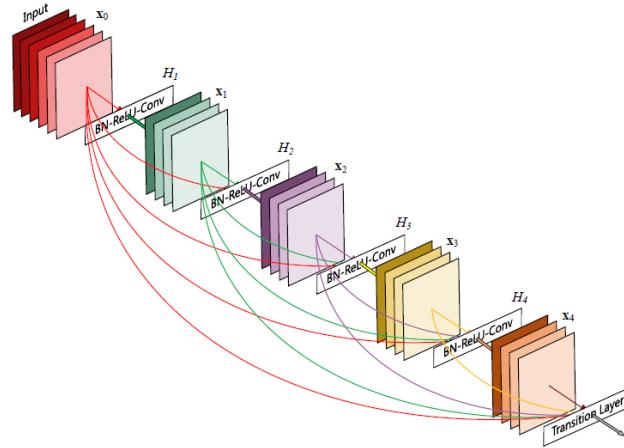


# Densely Connected Convolutional Networks



2018.06.29  
안건이

# Content

- **Convolutional Neural Networks**
  - **History**
- **DenseNet**

# 01 | History LeNet-5

[LeCun et al., 1998]

- Conv filters 5x5, stride 1
- Pooling 2x2, stride 2
- [Conv – Pool – Conv – Pool – FC – FC]

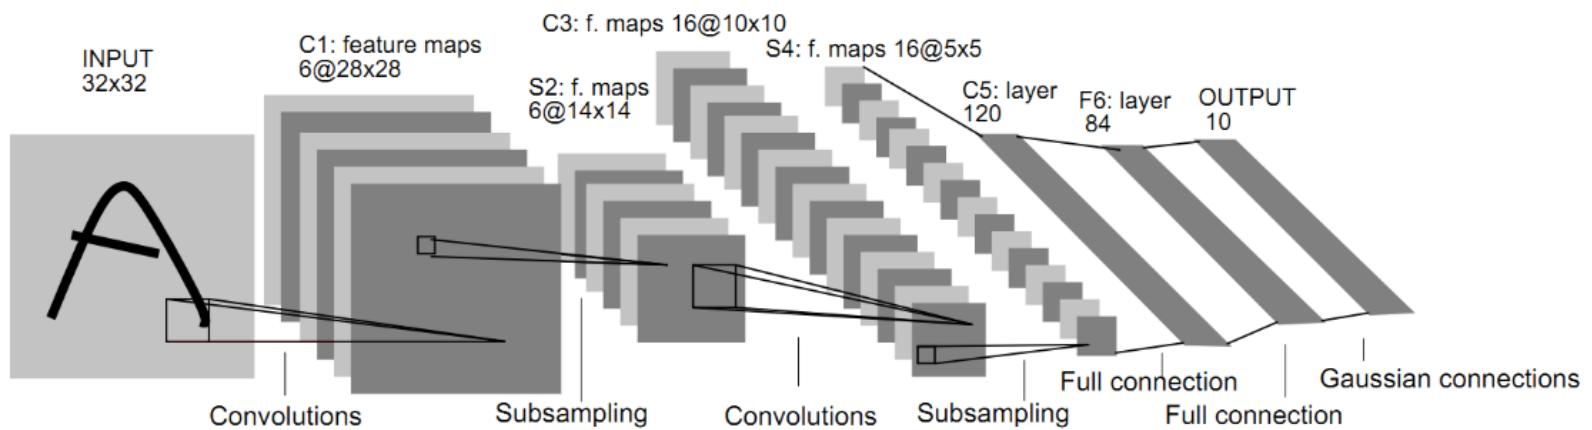
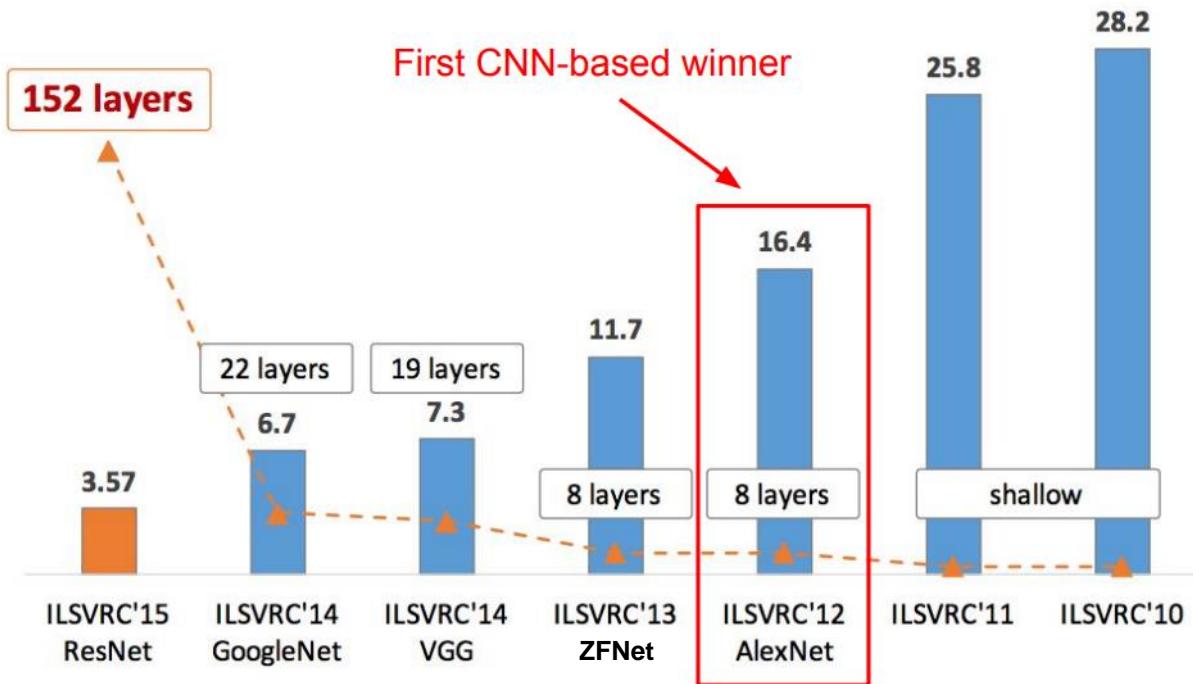


Fig. 2. Architecture of LeNet-5, a Convolutional Neural Network, here for digits recognition. Each plane is a feature map, i.e. a set of units whose weights are constrained to be identical.

# 01

## History ILSVRC Winners

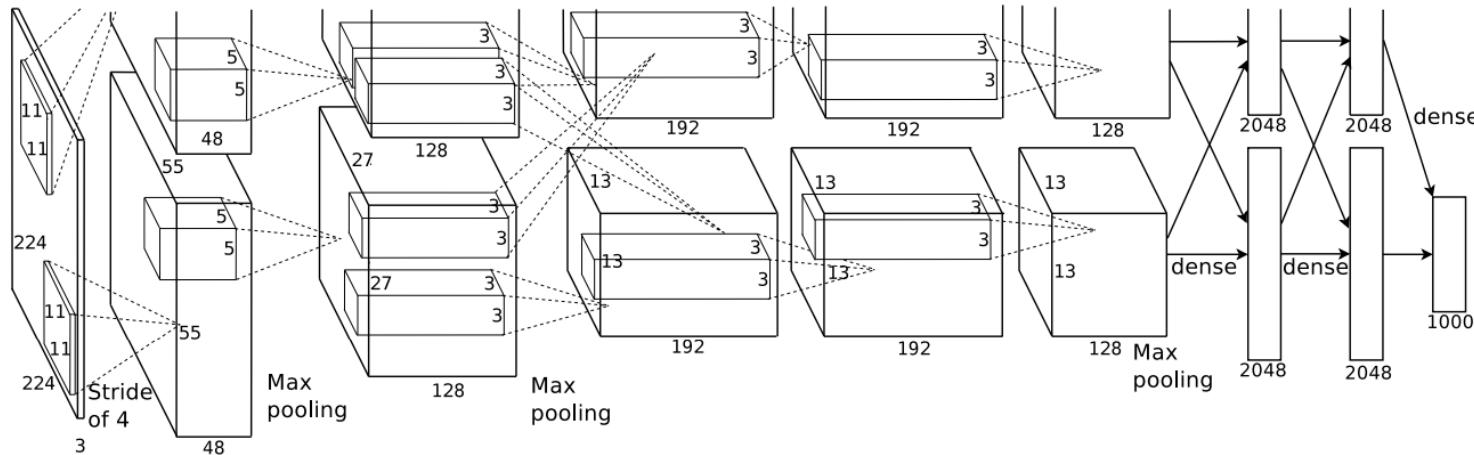
### ImageNet Large Scale Visual Recognition Challenge(ILSVRC) winners



# 01

## History AlexNet

[Krizhevsky et al., 2012]



[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

[6 x 6 x256] MAX POOL3: 3x3 filters at stride 2

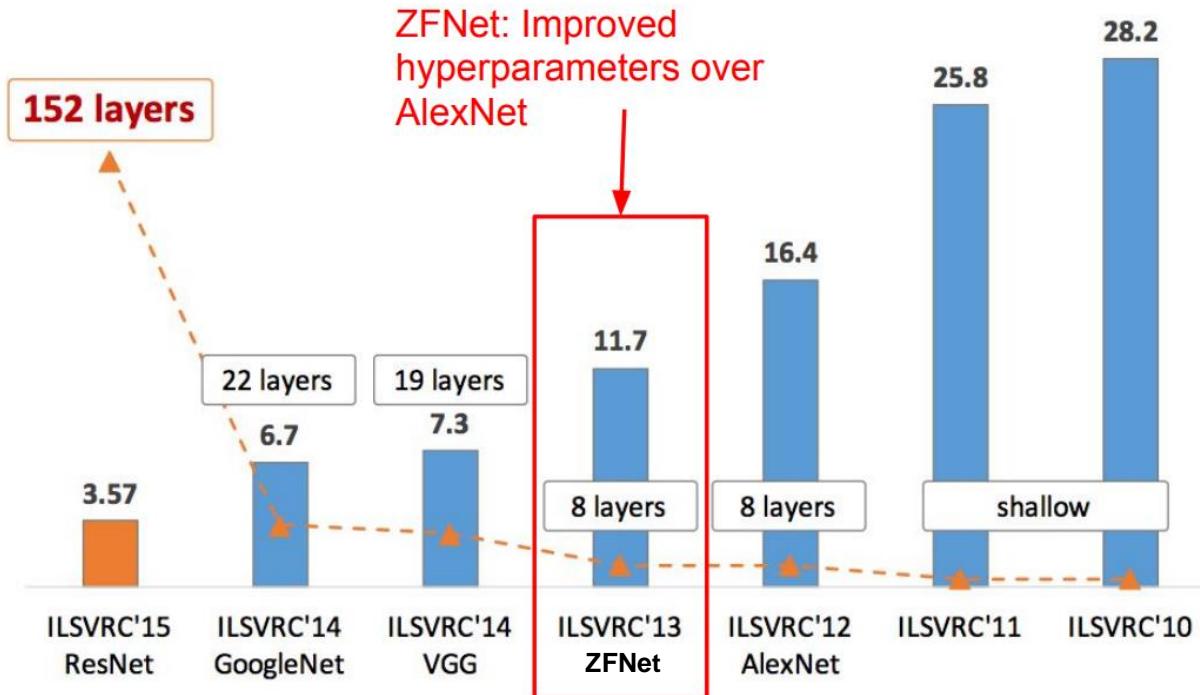
[4096] FC6: 4096 neurons

[4096] FC7: 4096 neurons

[1000] FC8: 1000 neurons (class scores)

# 01 | History ILSVRC Winners

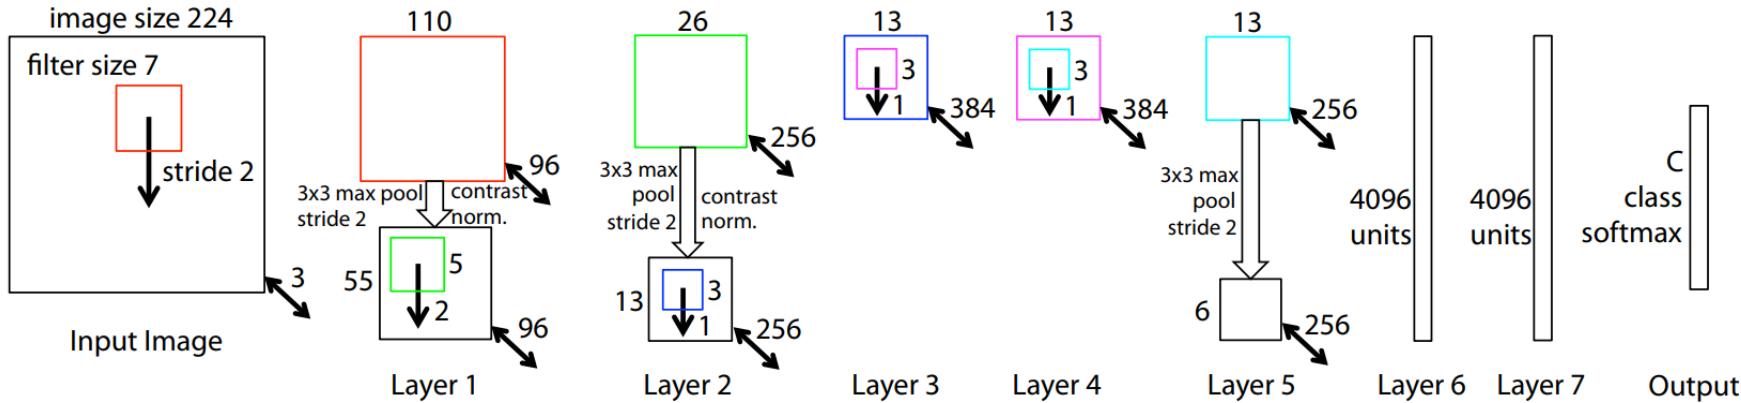
## ImageNet Large Scale Visual Recognition Challenge(ILSVRC) winners



# 01

## History ZFNet

[Zeiler and Fergus, 2013]



AlexNet but:

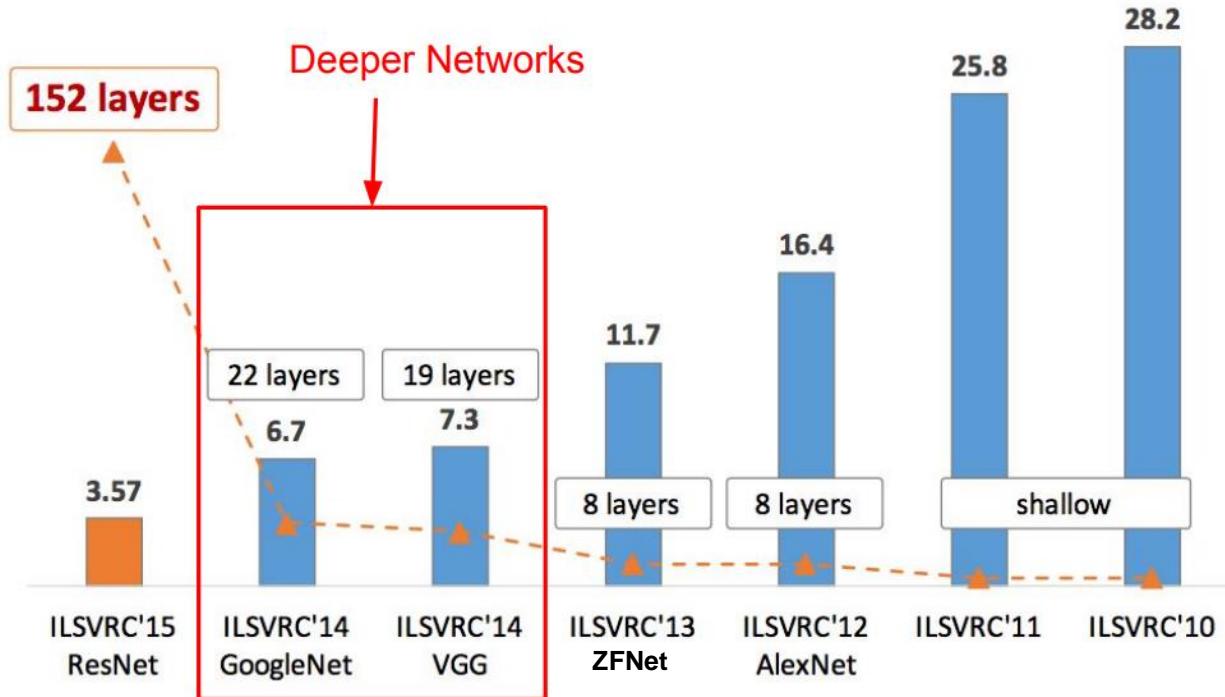
Conv1: change from (11x11 stride 4) to (7x7 stride 2)

Conv3,4,5: instead of 384, 384, 256 filters use 512, 1024, 512

ImageNet top 5 error: 16.4% → 11.7%

# 01 | History ILSVRC Winners

## ImageNet Large Scale Visual Recognition Challenge(ILSVRC) winners



# 01 | History VGGNet

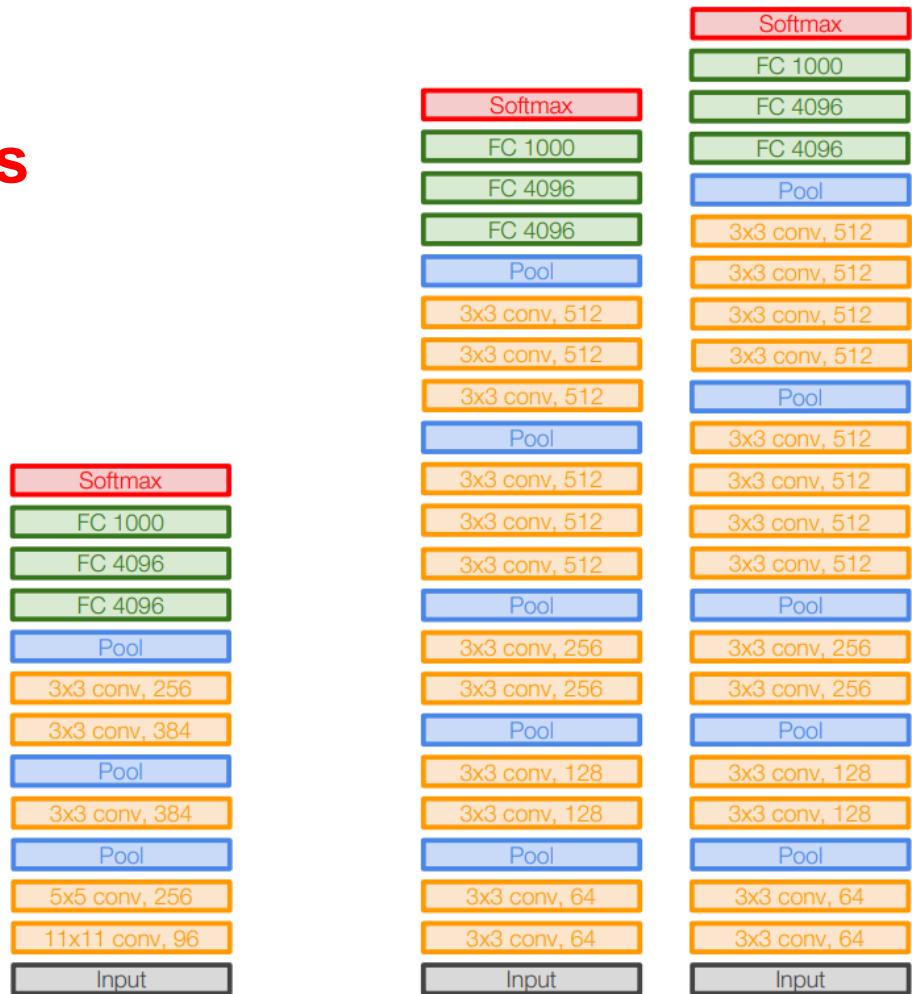
[Simonyan and Zisserman, 2014]

## Small Filters, Deeper networks

AlexNet: 8 layers

VGGNet: 16-19 layers

Only 3x3 Conv stride 1, pad 1  
2x2 MAX POOL stride 2



ImageNet top 5 error: 11.7% → 7.3%

AlexNet

VGG16

VGG19

# 01 | History VGGNet

[Simonyan and Zisserman, 2014]

## Small Filters, Deeper networks

AlexNet: 8 layers

VGGNet: 16-19 layers

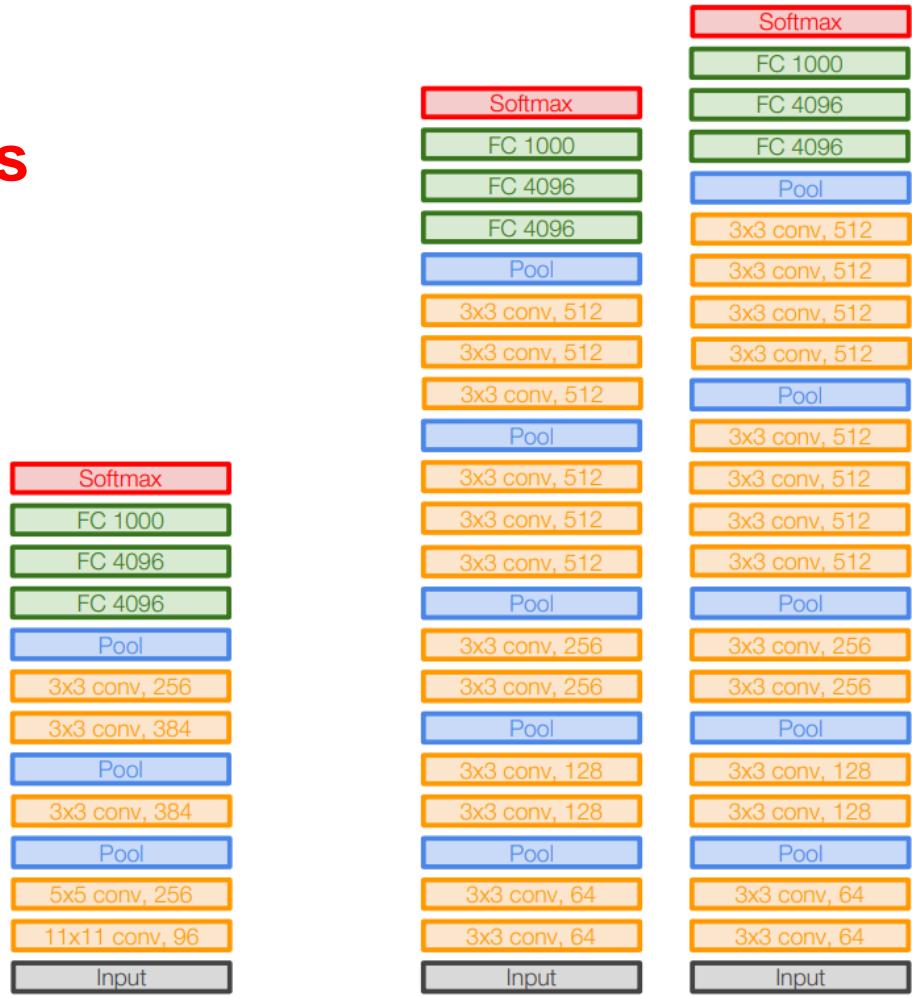
Only 3x3 Conv stride 1, pad 1  
2x2 MAX POOL stride 2

-Stack of three 3x3 conv  
# of param:  $3 * (3^2 * C^2) = 27 * C^2$

-Stack of one 7x7 conv  
# of param:  $7^2 * C^2 = 49 * C^2$

C: # of channel

ImageNet top 5 error: 11.7% → 7.3%



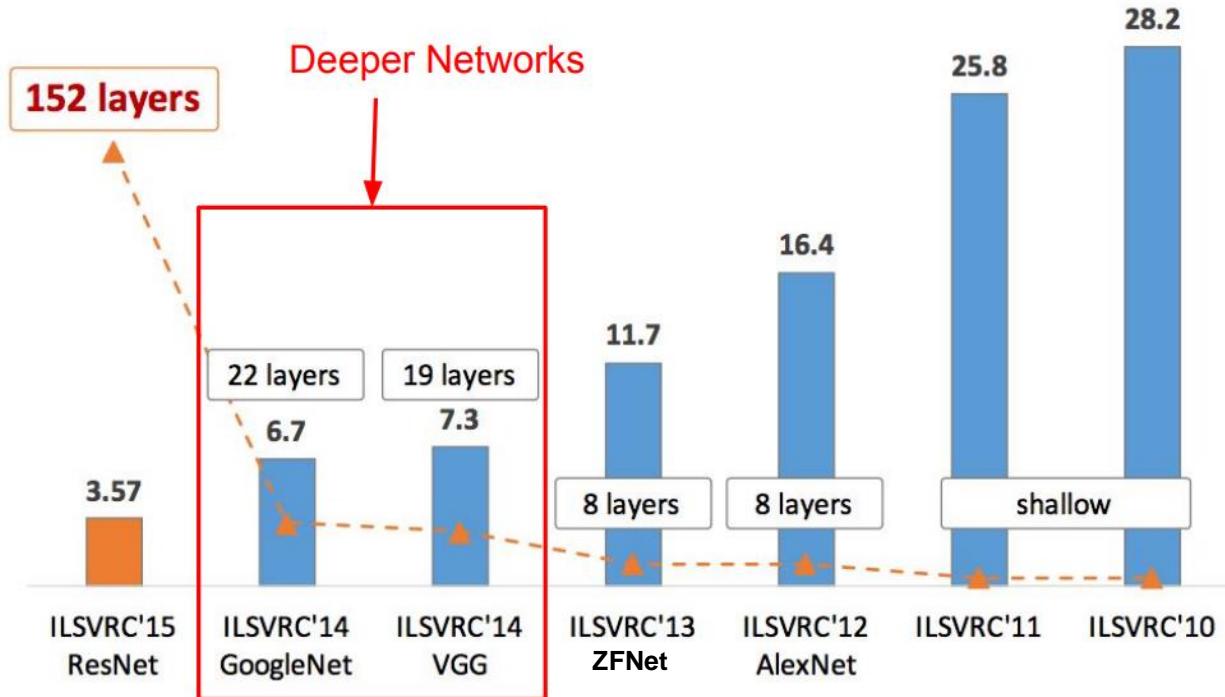
AlexNet

VGG16

VGG19

# 01 | History ILSVRC Winners

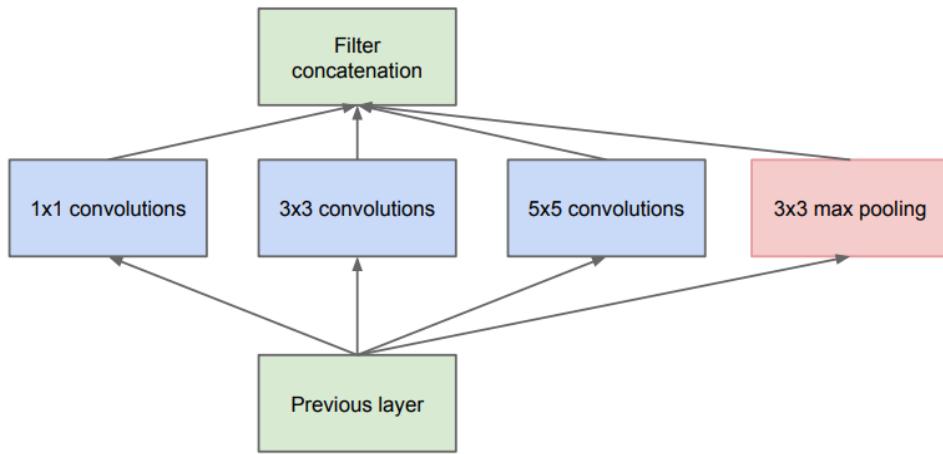
## ImageNet Large Scale Visual Recognition Challenge(ILSVRC) winners



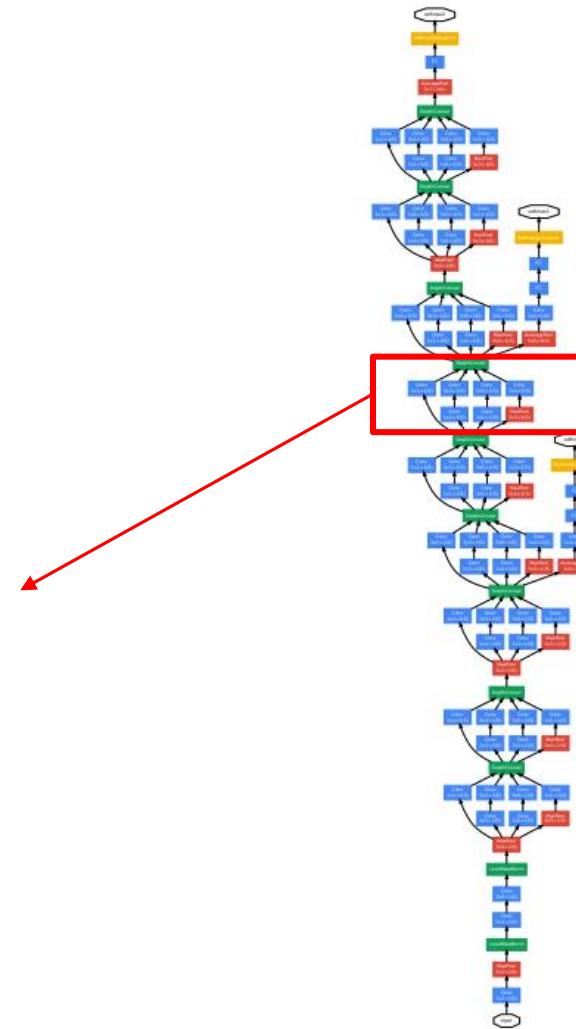
## Deeper networks, with computational efficiency

VGGNet: 16-19 layers

GoogLeNet: 22 layers

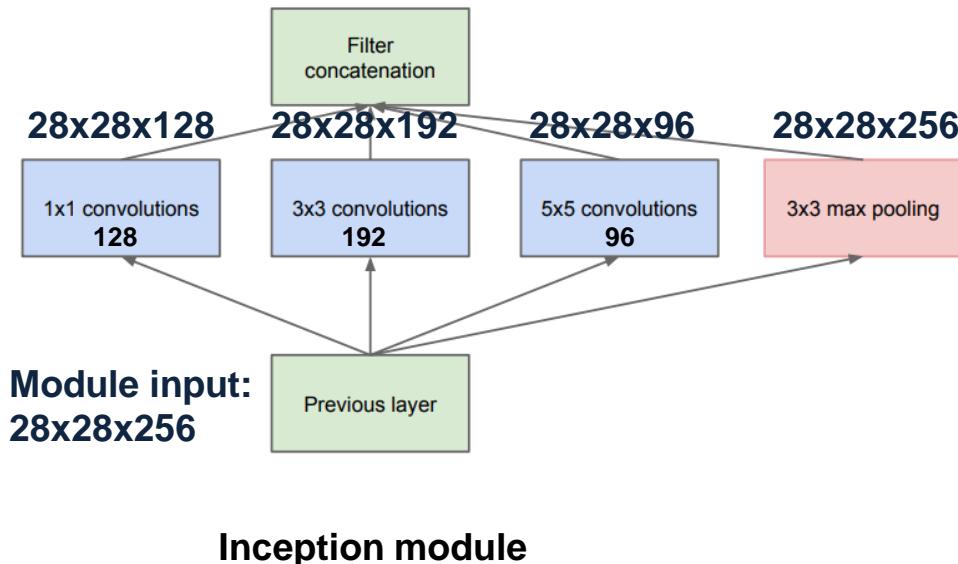


Inception module



## Deeper networks, with computational efficiency

$$28 \times 28 \times (128 + 192 + 96 + 256) = 28 \times 28 \times 672$$



**Conv Ops:**

[1x1 conv, 128]  $(28 \times 28 \times 128) \times (1 \times 1) \times 256$

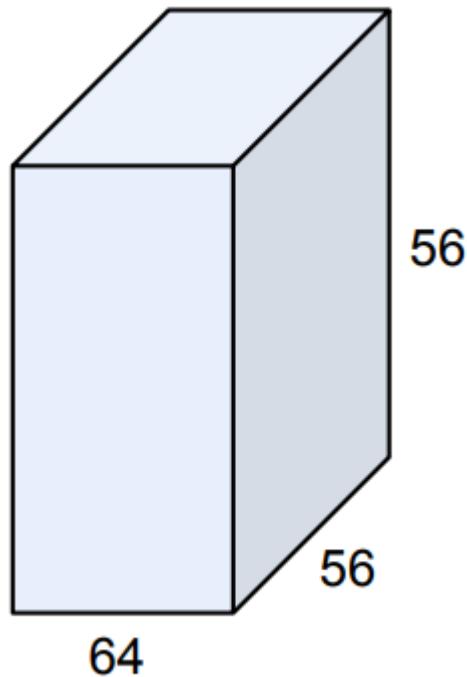
[3x3 conv, 192]  $(28 \times 28 \times 192) \times (3 \times 3) \times 256$

[5x5 conv, 96]  $(28 \times 28 \times 96) \times (5 \times 5) \times 256$

**Total: 854M ops**

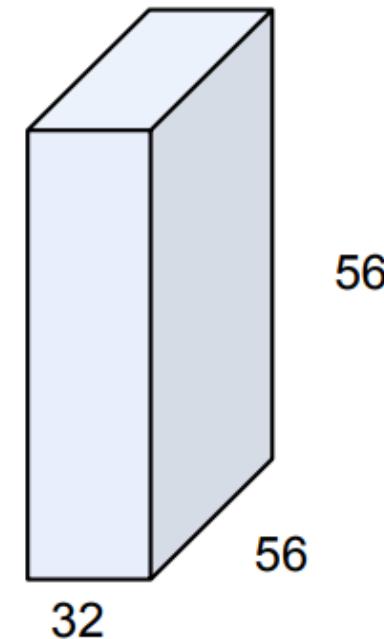
**Very expensive compute !!!**

## 1x1 conv “bottleneck” layers



1x1 Conv  
With 32 filters

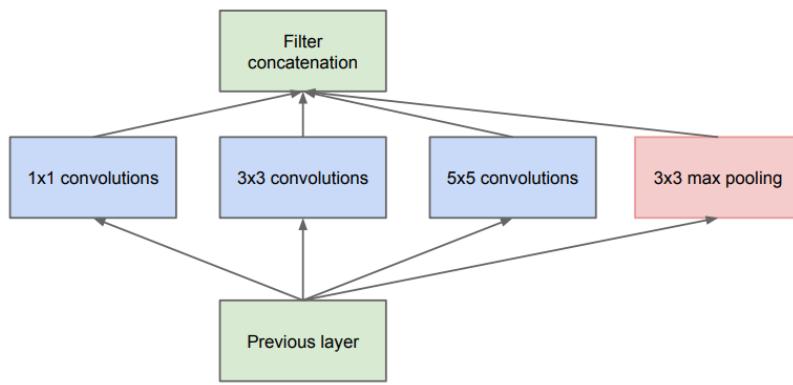
Preserves spatial dim  
Reduces depth dim



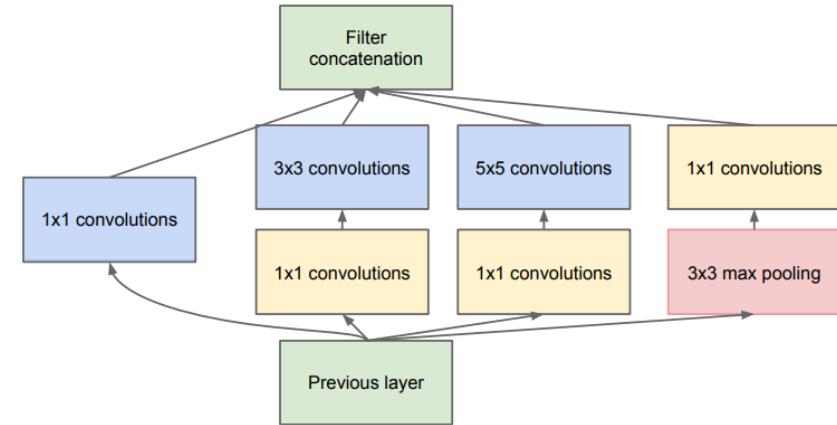
# 01

# History GoogLeNet

[Szegedy et al., 2014]



**Inception module**



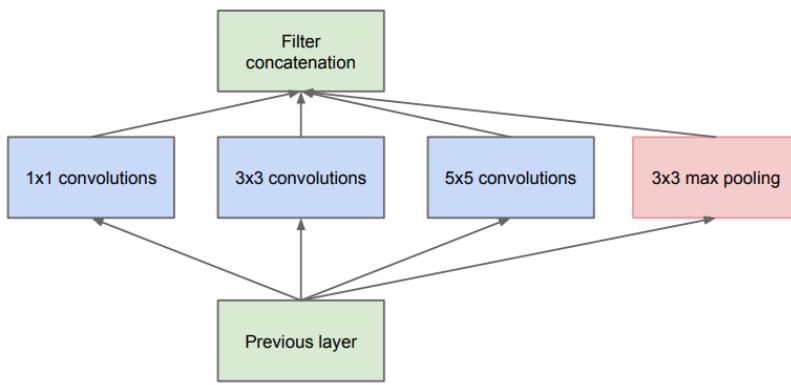
**Inception module with dimension reduction**

# 01

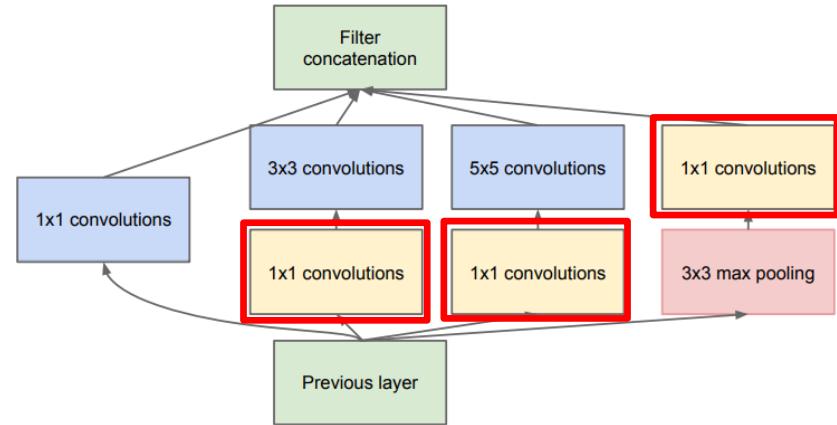
## History GoogLeNet

[Szegedy et al., 2014]

1x1 conv “bottleneck” layers



Inception module



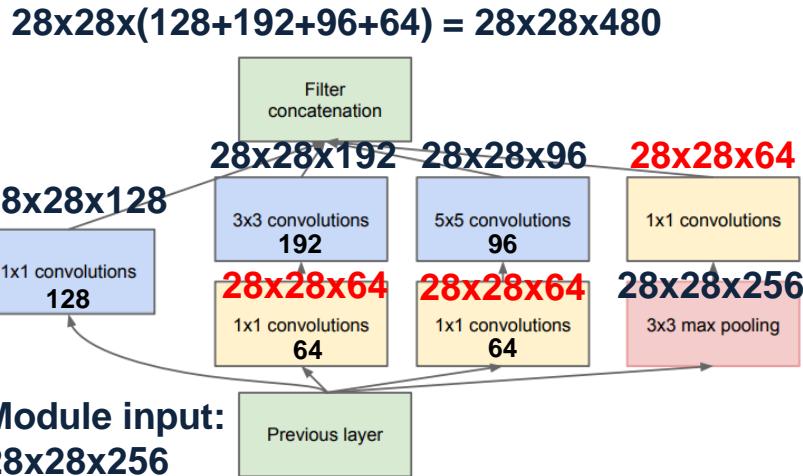
Inception module with dimension reduction

# 01

## History GoogLeNet

[Szegedy et al., 2014]

64, 1x1 conv “bottleneck” layers



Conv Ops:

- [1x1 conv, 64]  $(28 \times 28 \times 64) \times (1 \times 1) \times 256$
- [1x1 conv, 64]  $(28 \times 28 \times 64) \times (1 \times 1) \times 256$
- [1x1 conv, 128]  $(28 \times 28 \times 128) \times (1 \times 1) \times 256$
- [3x3 conv, 192]  $(28 \times 28 \times 192) \times (3 \times 3) \times 64$
- [5x5 conv, 96]  $(28 \times 28 \times 96) \times (5 \times 5) \times 64$
- [1x1 conv, 64]  $(28 \times 28 \times 64) \times (5 \times 5) \times 256$

Total: 358M ops

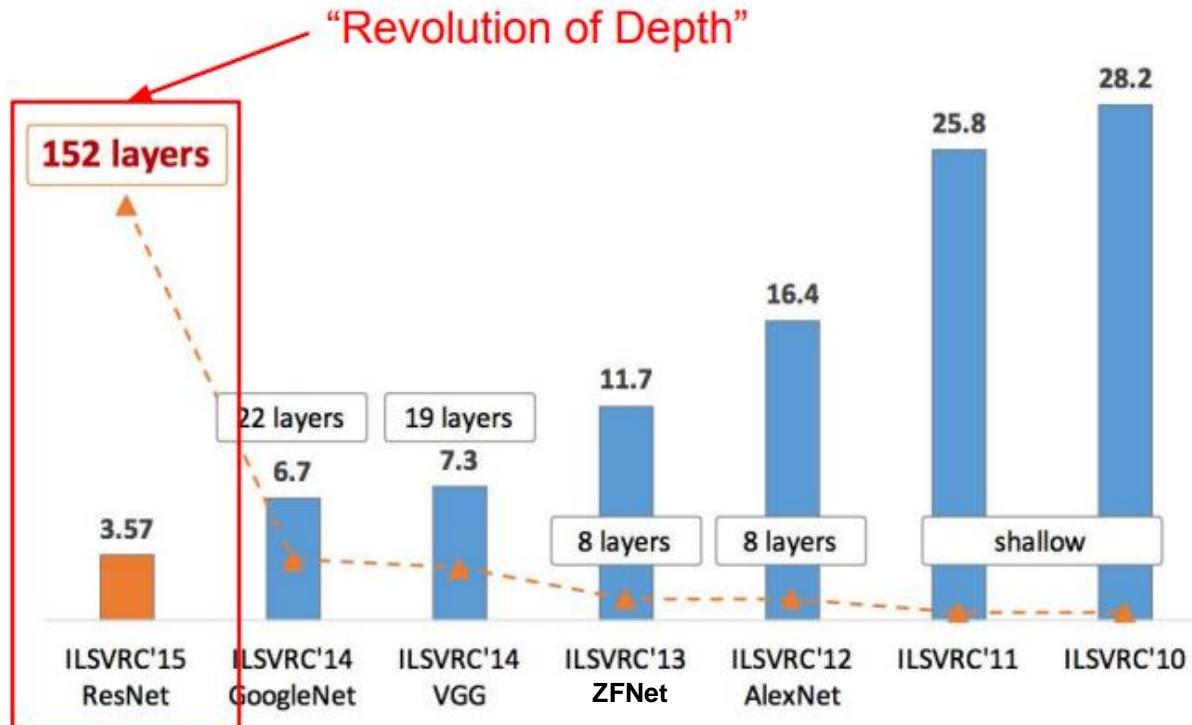
854M → 358M

Inception module with dimension reduction

ImageNet top 5 error: 7.3% → 6.7%

# 01 | History ILSVRC Winners

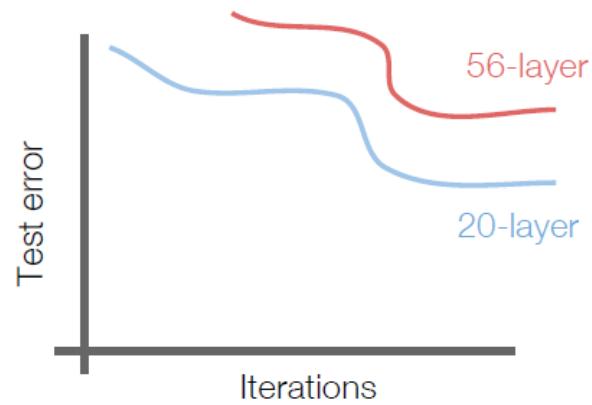
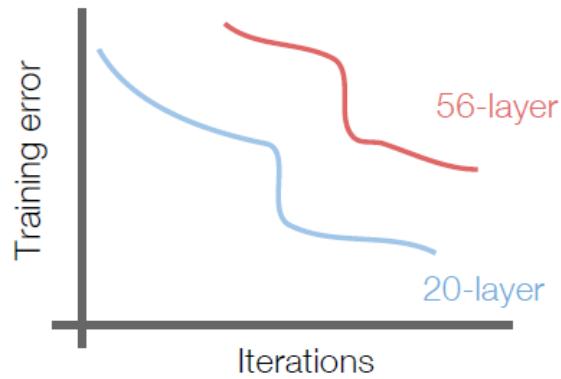
## ImageNet Large Scale Visual Recognition Challenge(ILSVRC) winners



# 01 | History ResNet

[He et al., 2015]

56-layer model performs worse on both training and test error  
The deeper model performs worse, but it's not caused by overfitting

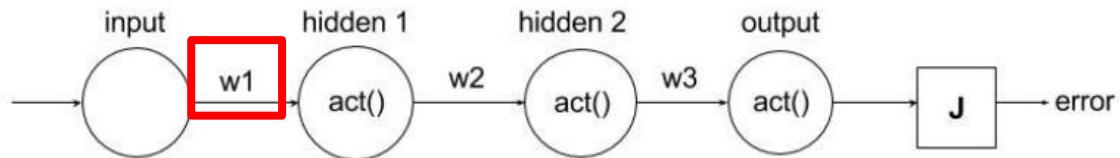
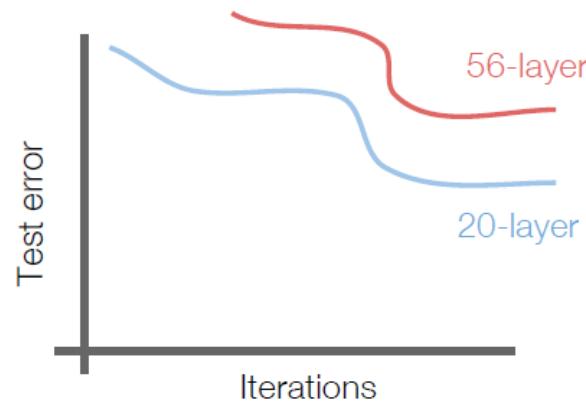
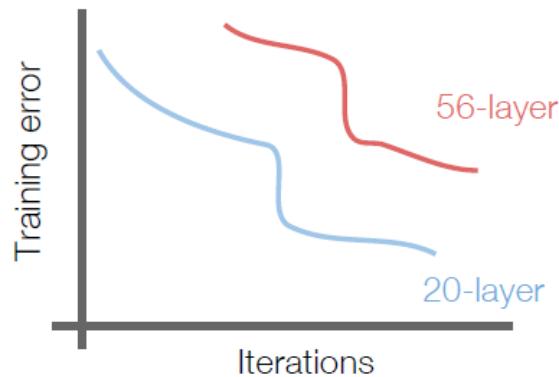


# 01

## History ResNet

[He et al., 2015]

56-layer model performs worse on both training and test error  
The deeper model performs worse, but it's not caused by overfitting



$$\frac{\partial \text{output}}{\partial \text{hidden2}} \frac{\partial \text{hidden2}}{\partial \text{hidden1}} = \frac{\frac{\partial \text{Sigmoid}(z_1)}{\partial z_1} < 1/4}{w_3} w_2 \frac{\frac{\partial \text{Sigmoid}(z_2)}{\partial z_2} < 1}{w_2} < 1/4$$

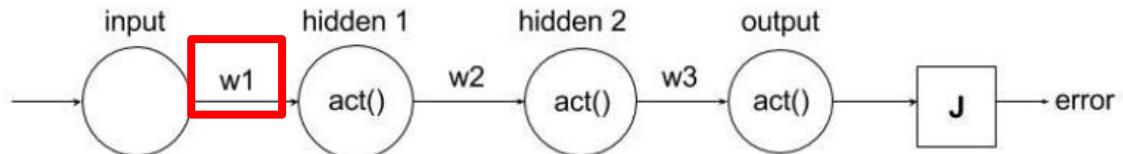
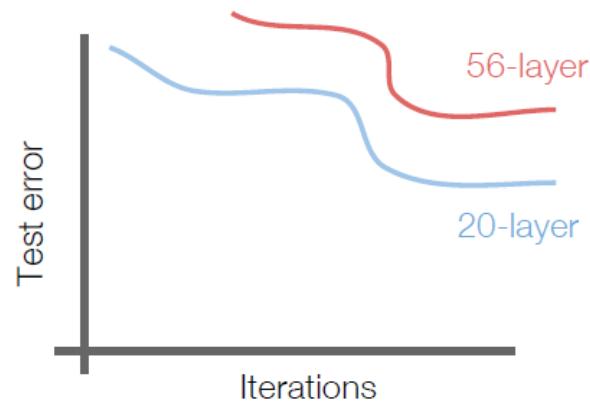
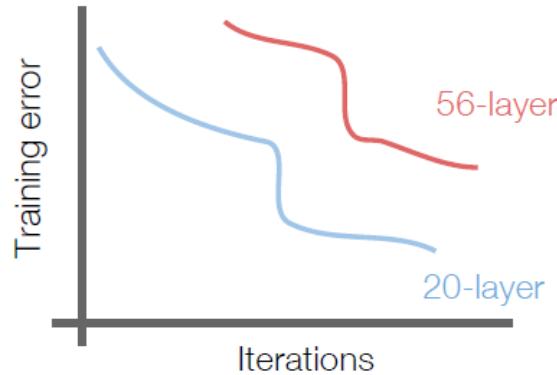
$$0.02625 = 0.25 \times 0.6 \times 0.25 \times 0.7$$

# 01

## History ResNet

[He et al., 2015]

56-layer model performs worse on both training and test error  
The deeper model performs worse, but it's not caused by overfitting



Caused by Gradient vanishing !!!

$$\frac{\partial \text{output}}{\partial \text{hidden2}} \frac{\partial \text{hidden2}}{\partial \text{hidden1}} = \frac{\partial \text{Sigmoid}(z_1)}{\partial z_1} w_3 * \frac{\partial \text{Sigmoid}(z_2)}{\partial z_2} w_2$$

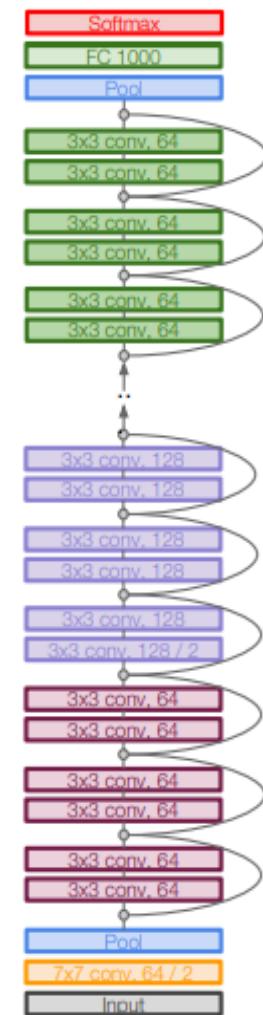
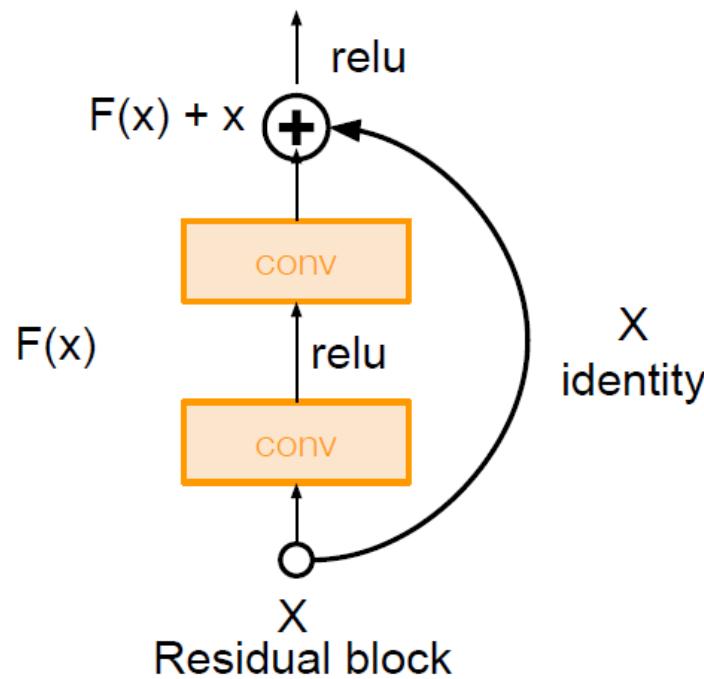
$< 1/4$        $< 1$        $< 1/4$        $< 1$

$$0.02625 = 0.25 \times 0.6 \times 0.25 \times 0.7$$

## Very deep networks using residual connections

GoogLeNet: 22 layers

ResNet: 152 layers



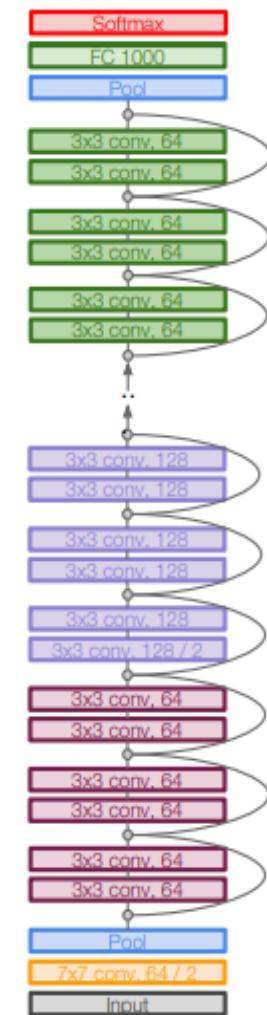
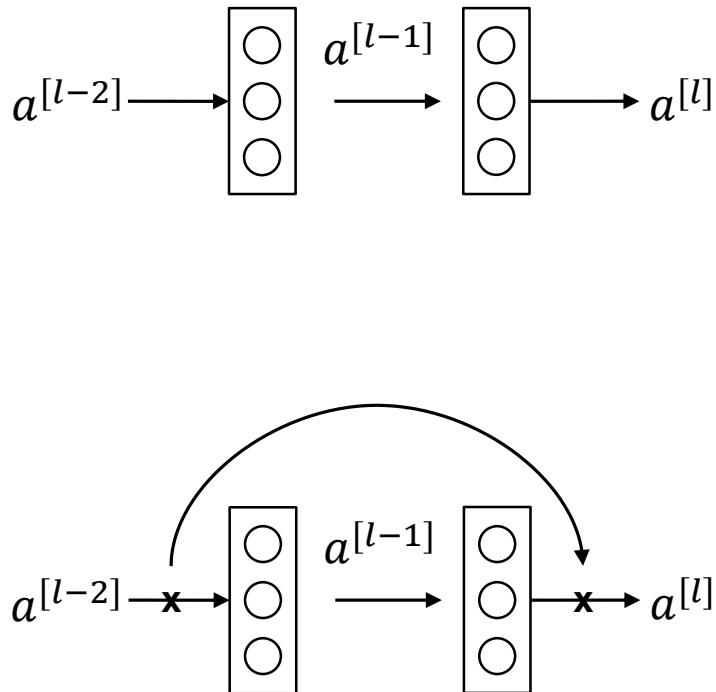
# 01 | History ResNet

[He et al., 2015]

## Very deep networks using residual connections

## GoogLeNet: 22 layers

# ResNet: 152 layers



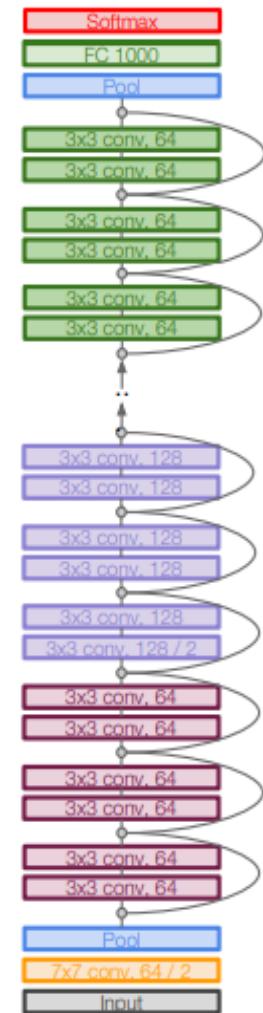
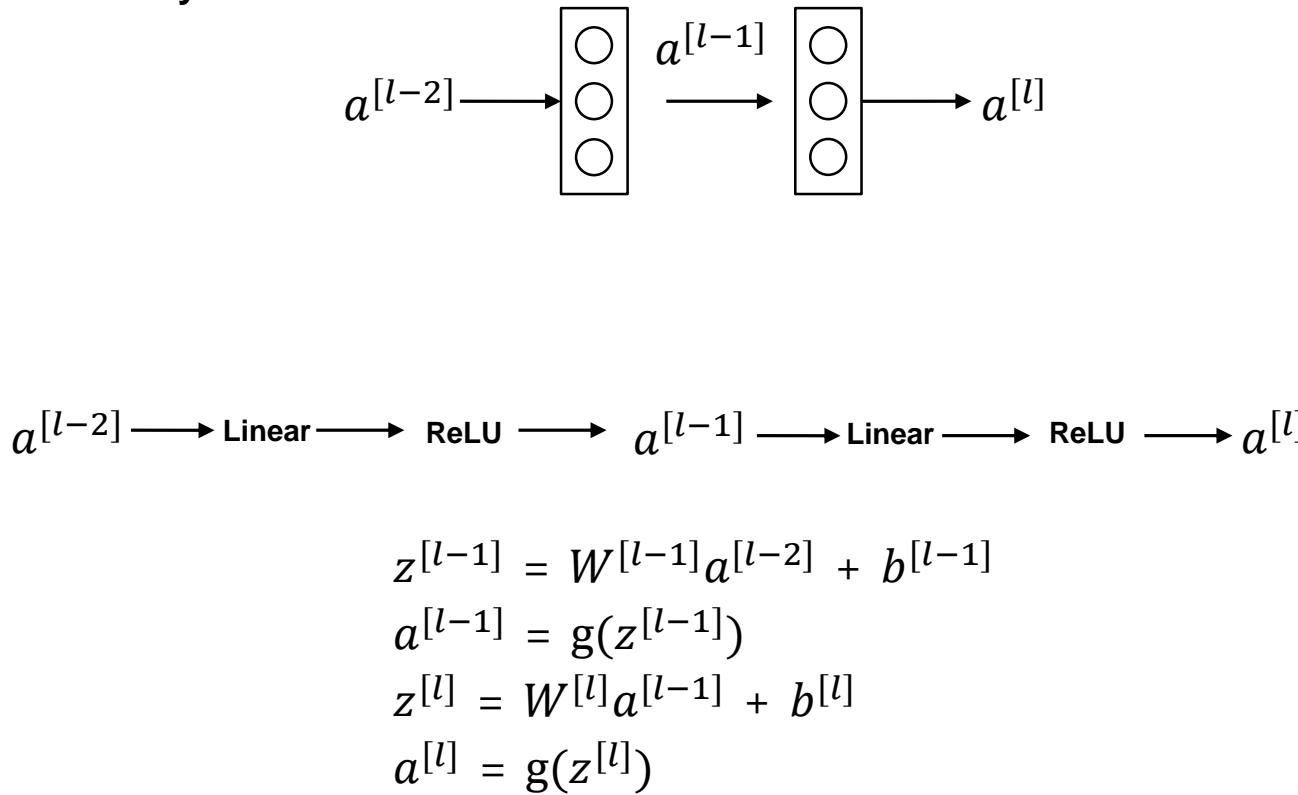
# 01 | History ResNet

[He et al., 2015]

## Very deep networks using residual connections

GoogLeNet: 22 layers

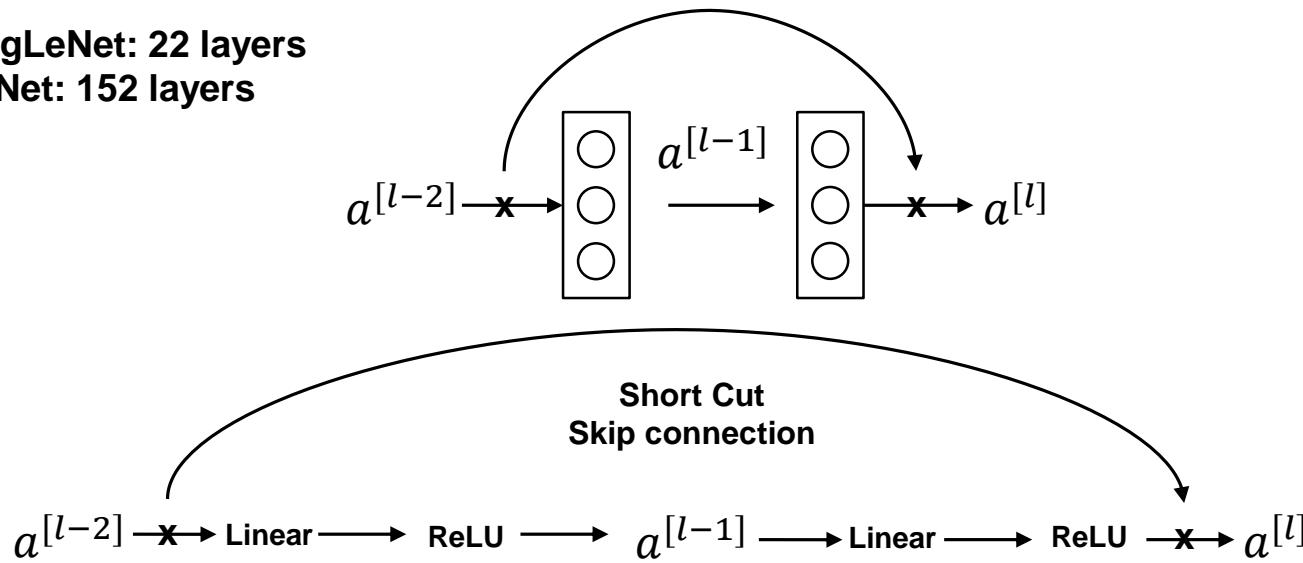
ResNet: 152 layers



## Very deep networks using residual connections

GoogLeNet: 22 layers

ResNet: 152 layers

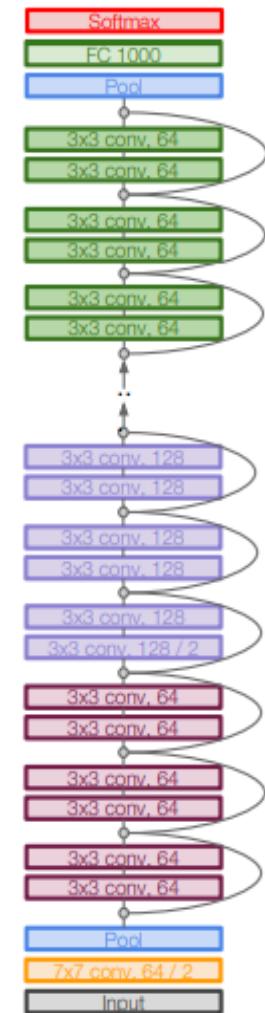


$$z^{[l-1]} = W^{[l-1]}a^{[l-2]} + b^{[l-1]}$$

$$a^{[l-1]} = g(z^{[l-1]})$$

$$z^{[l]} = W^{[l]}a^{[l-1]} + b^{[l]}$$

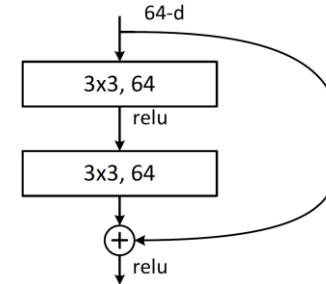
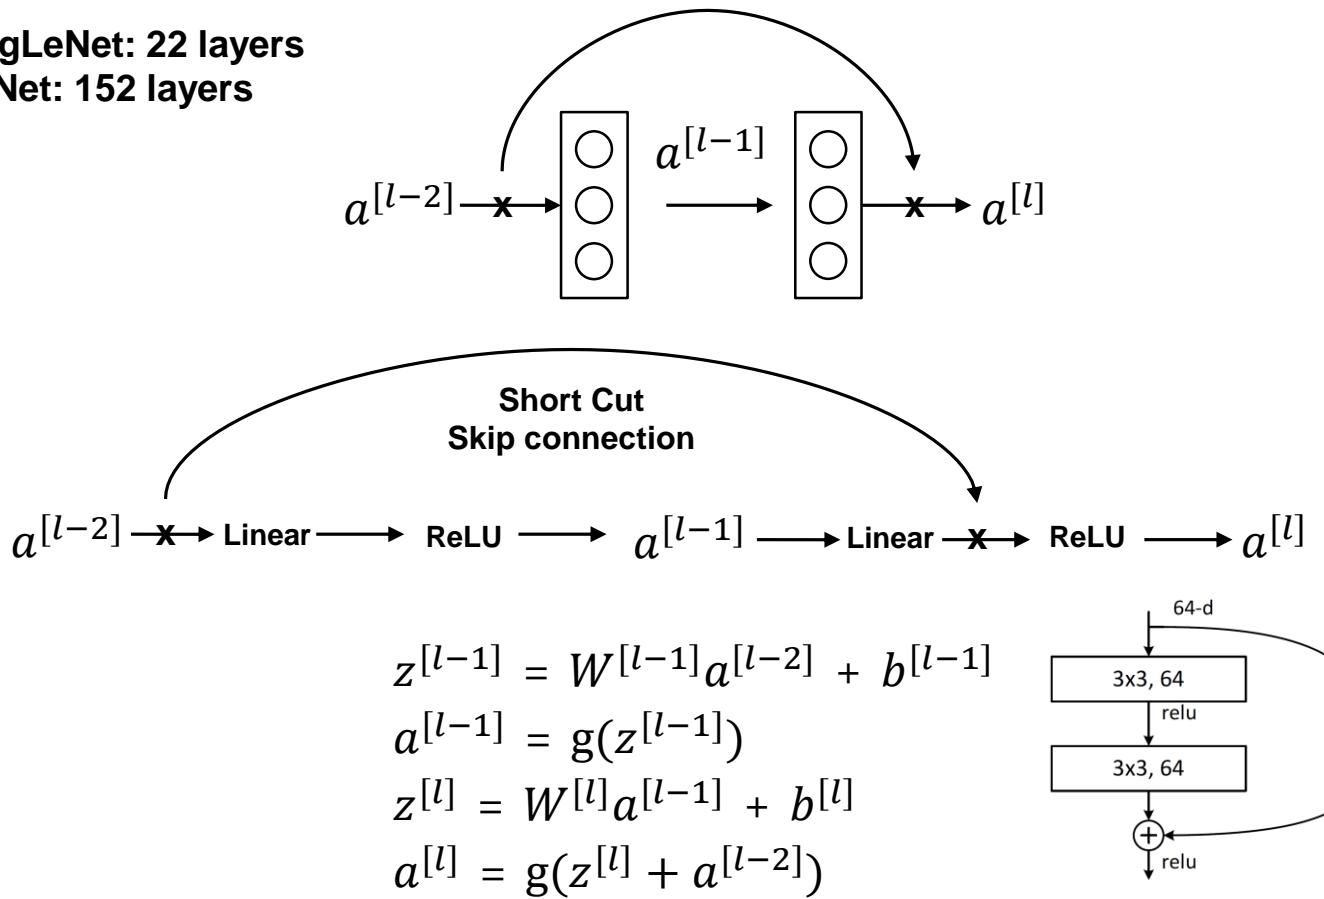
$$a^{[l]} = g(z^{[l]}) + a^{[l-2]}$$



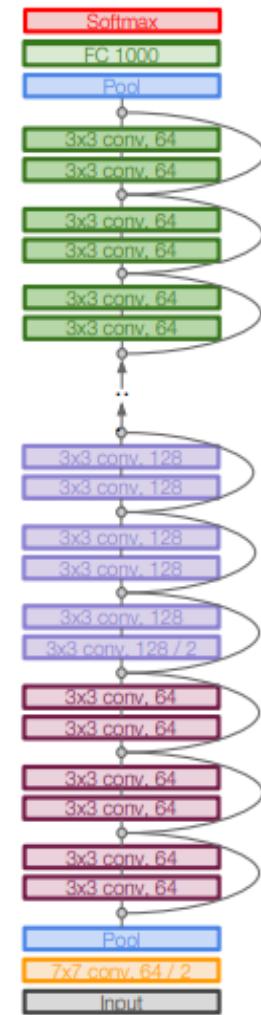
## Very deep networks using residual connections

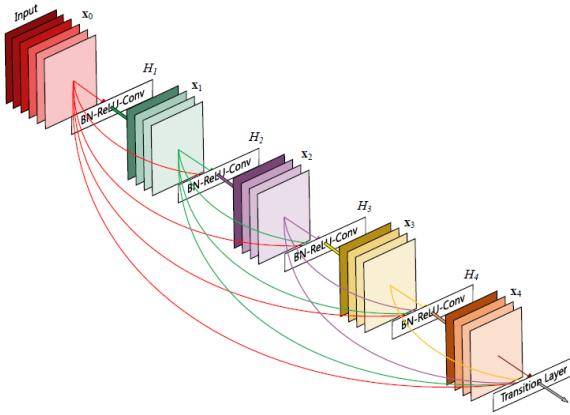
GoogLeNet: 22 layers

ResNet: 152 layers

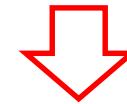


ImageNet top 5 error: 6.7% → 3.57%





**Skip-Connection**



**Dense Connectivity**

### DenseNet Advantage

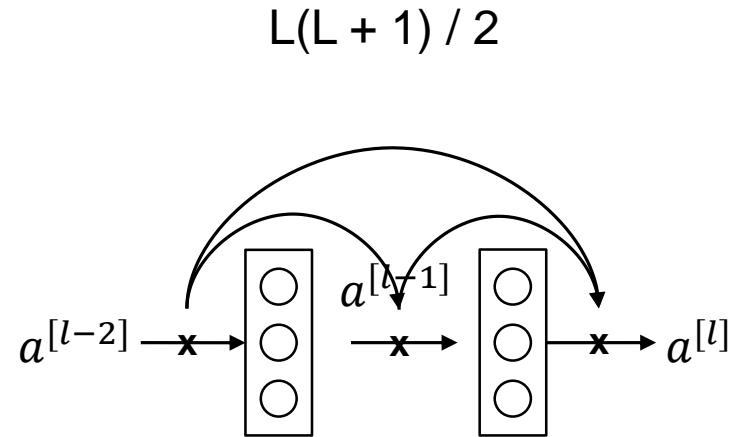
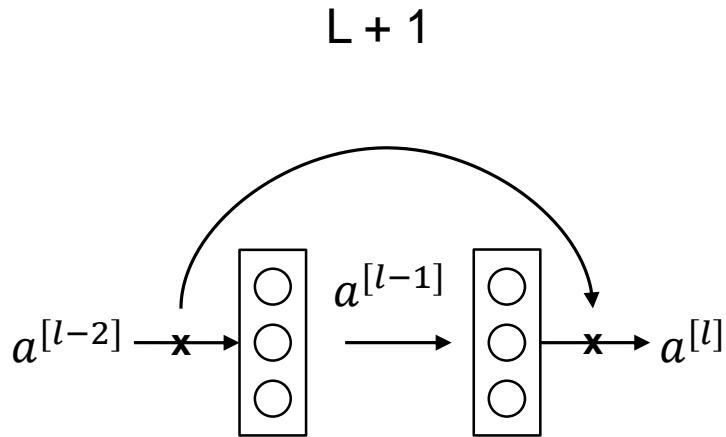
1. They alleviate the vanishing-gradient problem
2. Strengthen feature propagation
3. Encourage feature reuse
4. Substantially reduce the number of parameters

**Skip-Connection  
(ResNet)**

$$\mathbf{a}^{[l]} = \mathbf{H}_l(\mathbf{a}^{[l-1]}) + \mathbf{a}^{[l-2]}$$

**Dense Connectivity  
(DenseNet)**

$$\mathbf{a}^{[l]} = \mathbf{H}_l([\mathbf{a}^{[0]}, \mathbf{a}^{[1]}, \dots, \mathbf{a}^{[l-1]}])$$

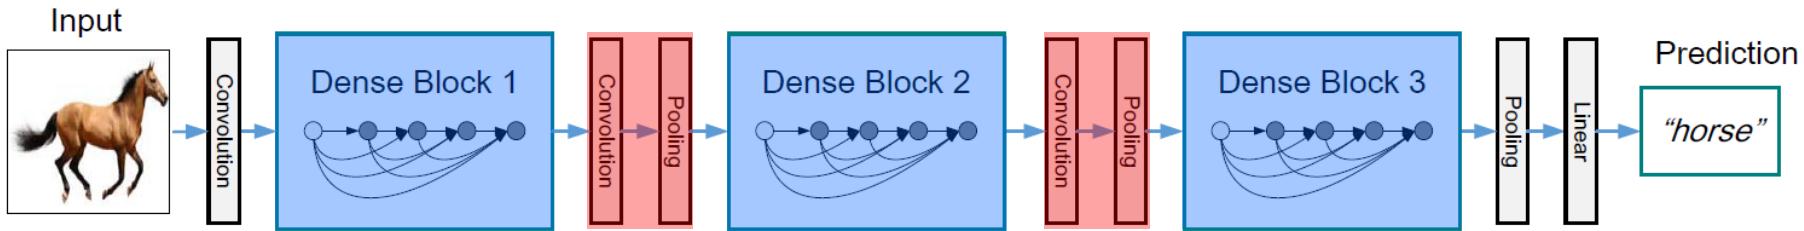


## 02

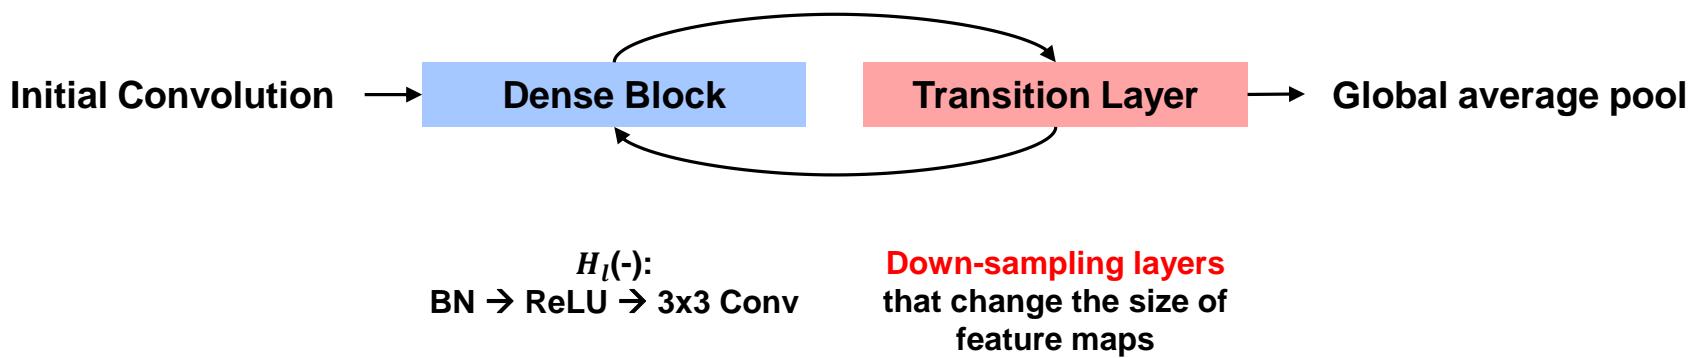
# Densely Connected Convolutional Networks

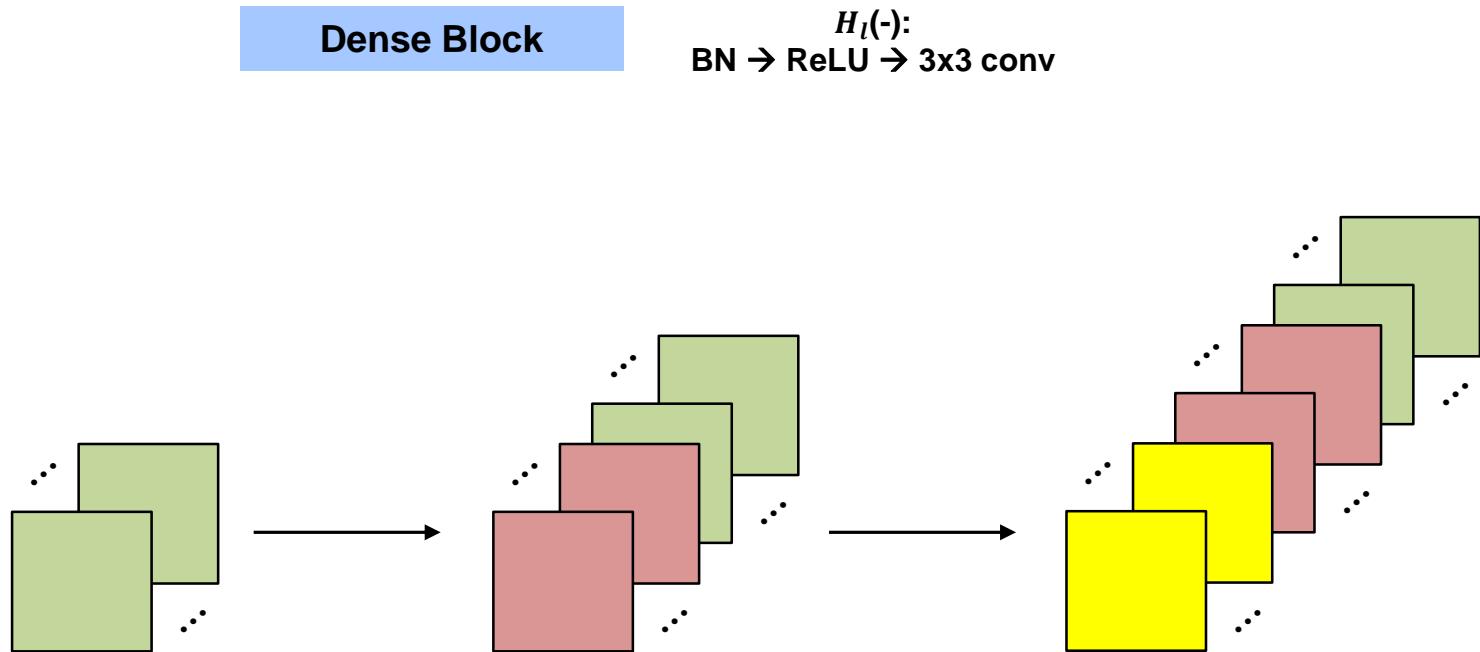
## DenseNet

[Huang et al, 2016]  
Best Paper in CVPR 2017



**Figure 2:** A deep DenseNet with three dense blocks. The layers between two adjacent blocks are referred to as transition layers and change feature-map sizes via convolution and pooling.





**Not viable when size of feature maps changes**

## Simple DenseNet

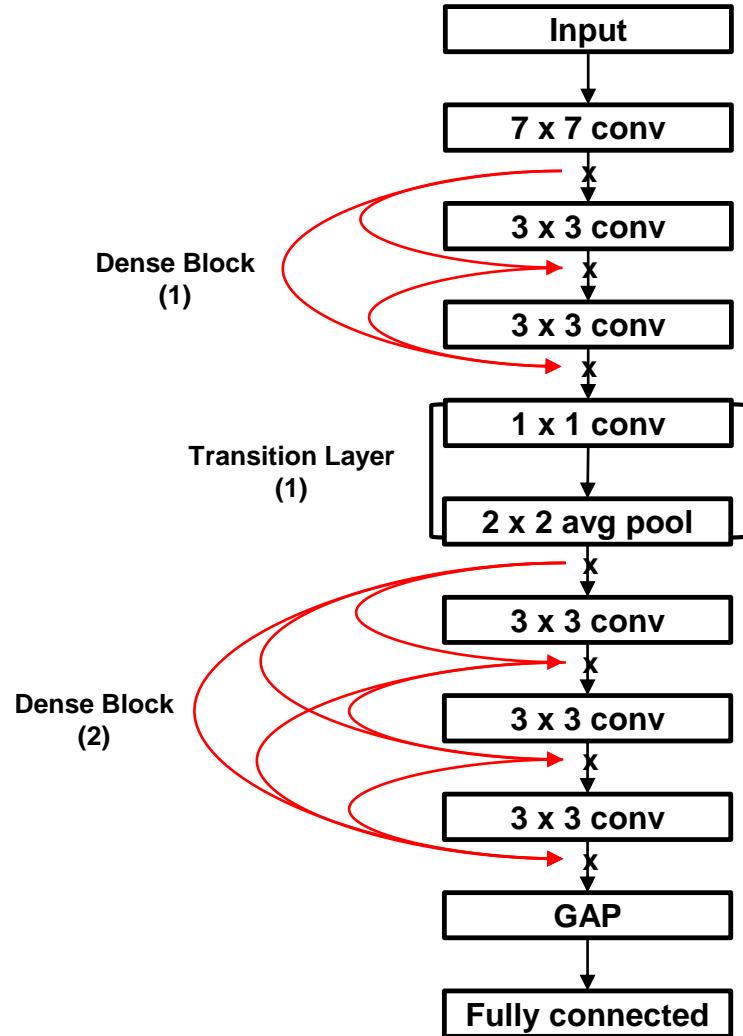
**Convolution** :  $7 \times 7$  conv

**Dense Block** :  $[3 \times 3 \text{ conv}] \times 2$   
(1)

**Transition Layer** :  $1 \times 1$  conv  
(1)  $2 \times 2$  average pool

**Dense Block** :  $[3 \times 3 \text{ conv}] \times 3$   
(2)

**Classification Layer** : global average pool  
2D fully-connected, softmax



## Simple DenseNet

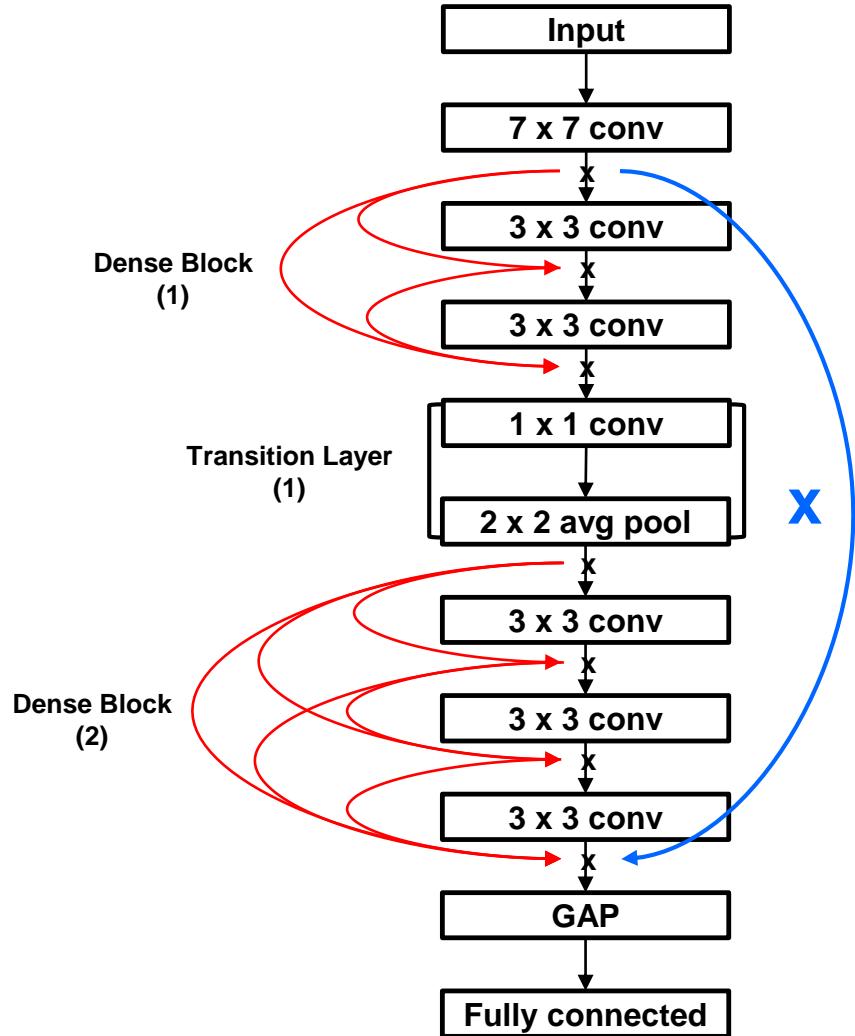
**Convolution** :  $7 \times 7$  conv

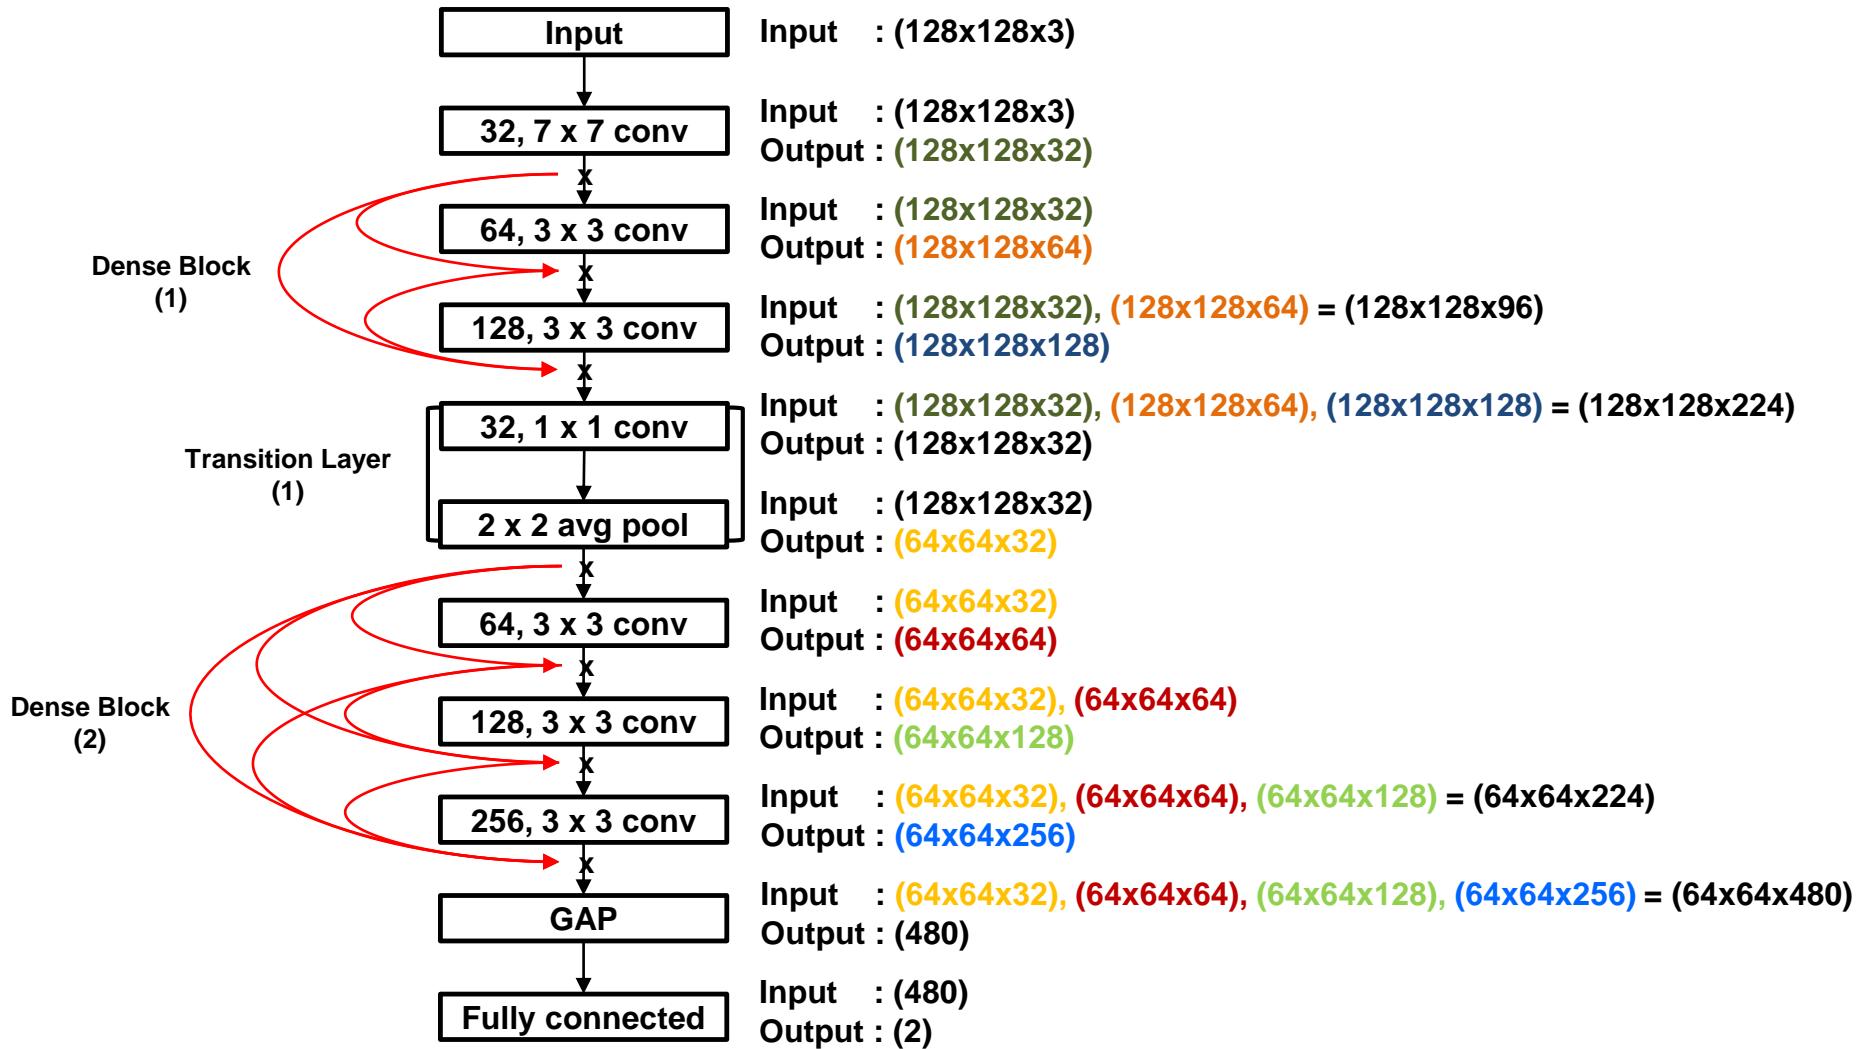
**Dense Block** :  $[3 \times 3 \text{ conv}] \times 2$   
(1)

**Transition Layer** :  $1 \times 1$  conv  
(1)  $2 \times 2$  average pool

**Dense Block** :  $[3 \times 3 \text{ conv}] \times 3$   
(2)

**Classification Layer** : global average pool  
2D fully-connected, softmax

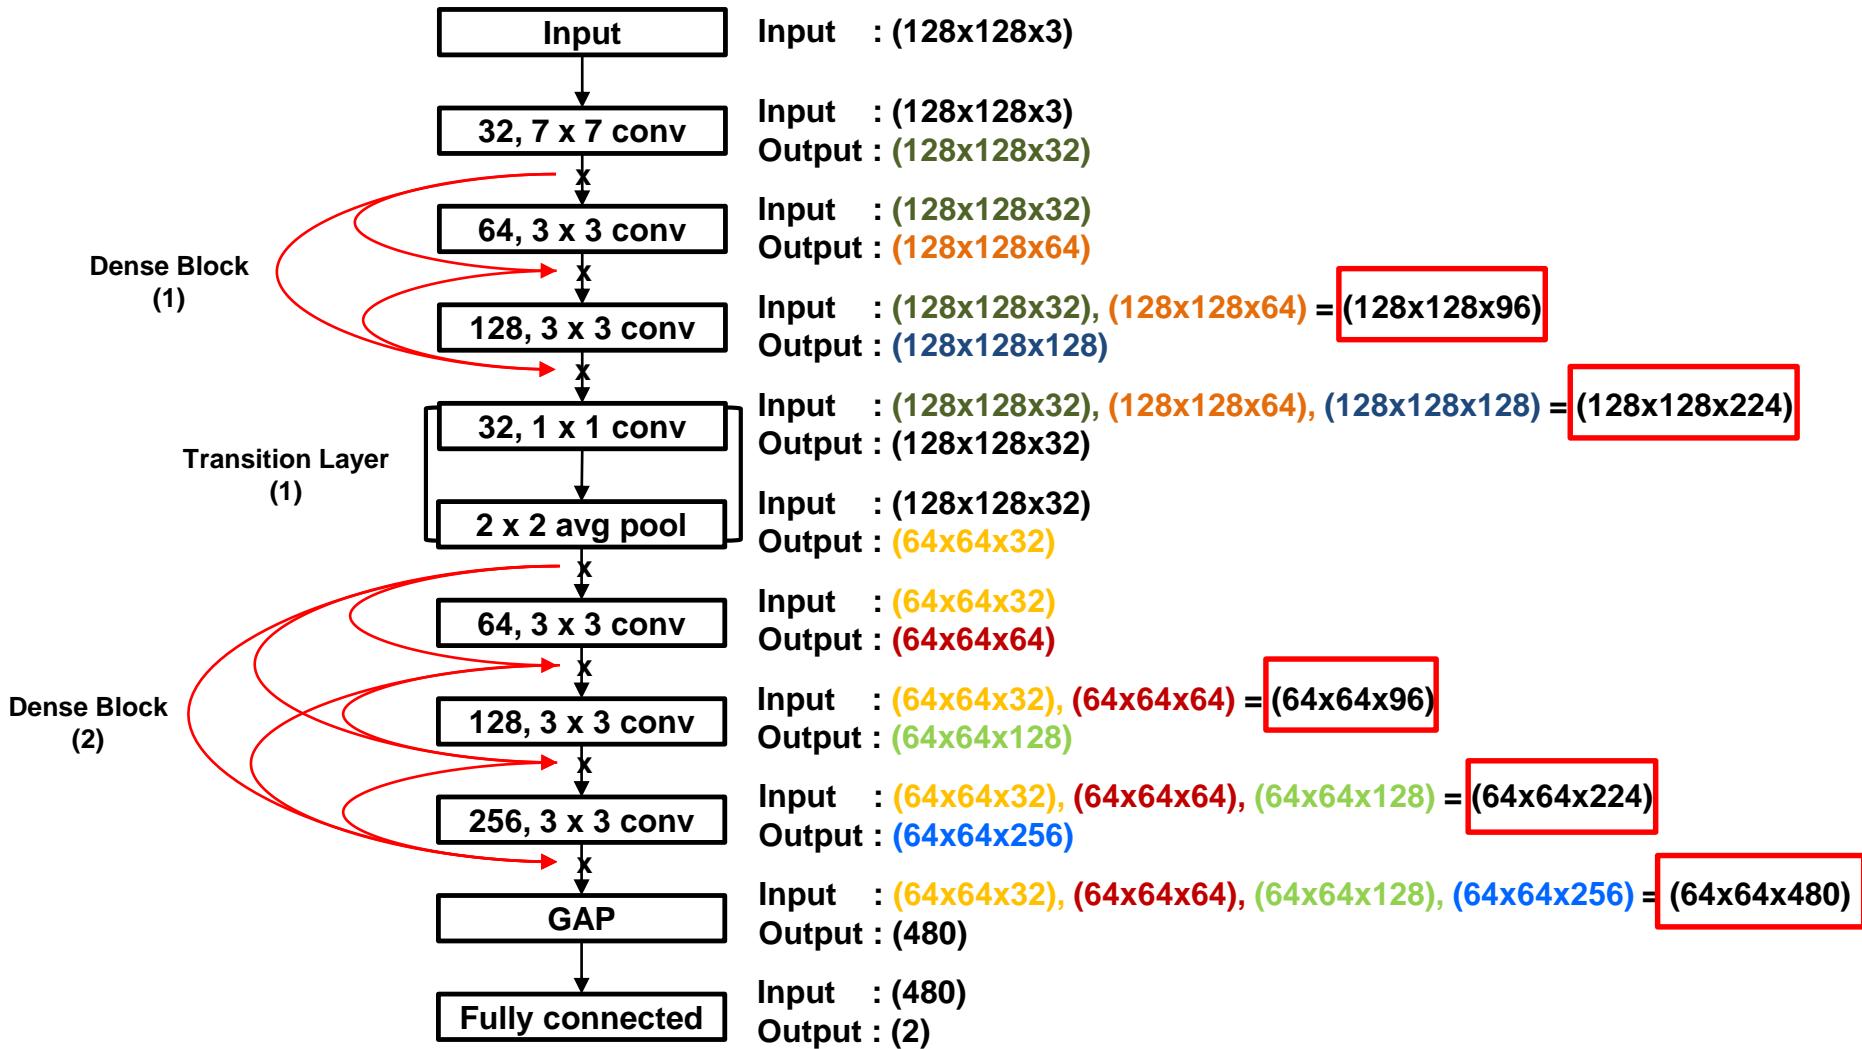




## 02

Densely Connected Convolutional Networks  
DenseNet

[Huang et al, 2016]  
Best Paper in CVPR 2017



**Simple DenseNet**

**Convolution** :  $7 \times 7$  conv

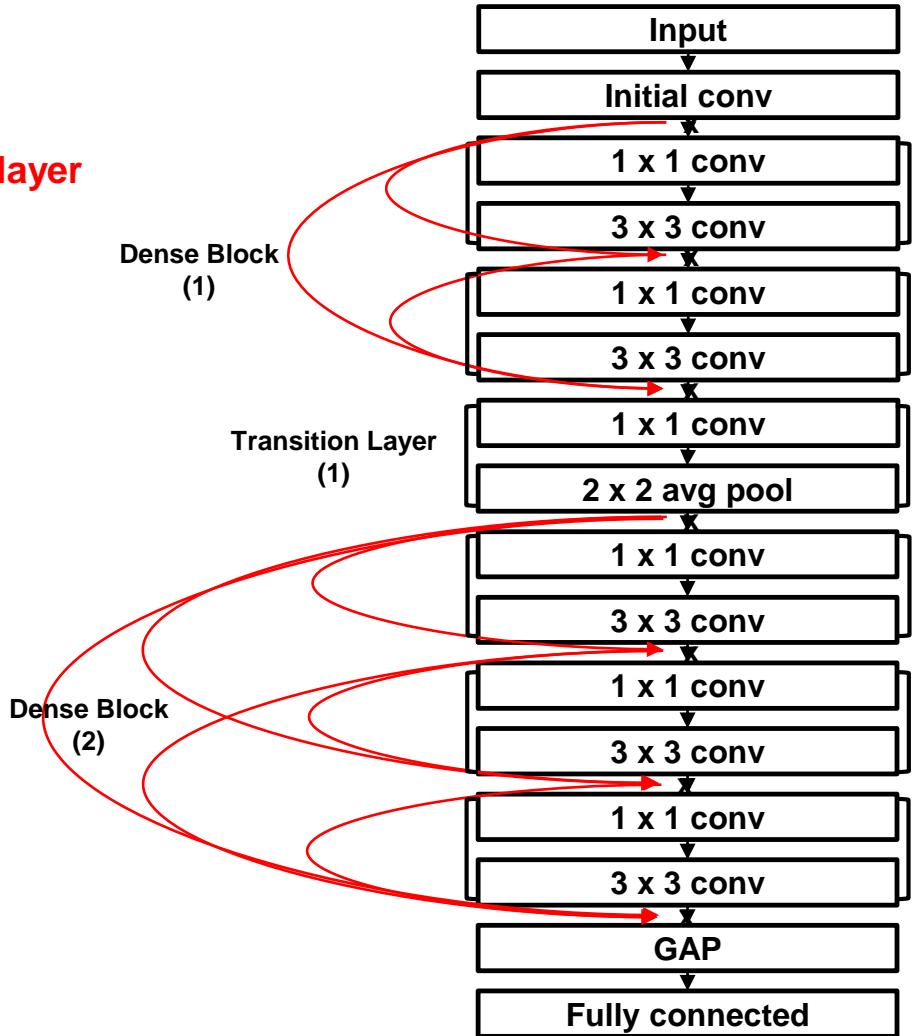
**Dense Block (1)** :  $\begin{bmatrix} 1 \times 1 \text{ conv} \\ 3 \times 3 \text{ conv} \end{bmatrix} \times 2$

**Transition Layer (1)** :  $1 \times 1$  conv  
 $2 \times 2$  average pool

**Dense Block (2)** :  $\begin{bmatrix} 1 \times 1 \text{ conv} \\ 3 \times 3 \text{ conv} \end{bmatrix} \times 3$

**Classification Layer** : global average pool  
 2D fully-connected, softmax

**Bottleneck layer**



## 02

Densely Connected Convolutional Networks  
DenseNet

[Huang et al, 2016]  
Best Paper in CVPR 2017

Layers	Output Size	DenseNet-121	DenseNet-169	DenseNet-201	DenseNet-264
Convolution	$112 \times 112$		$7 \times 7$ conv, stride 2		
Pooling	$56 \times 56$		$3 \times 3$ max pool, stride 2		
Dense Block (1)	$56 \times 56$	$\begin{bmatrix} 1 \times 1 \text{ conv} \\ 3 \times 3 \text{ conv} \end{bmatrix} \times 6$	$\begin{bmatrix} 1 \times 1 \text{ conv} \\ 3 \times 3 \text{ conv} \end{bmatrix} \times 6$	$\begin{bmatrix} 1 \times 1 \text{ conv} \\ 3 \times 3 \text{ conv} \end{bmatrix} \times 6$	$\begin{bmatrix} 1 \times 1 \text{ conv} \\ 3 \times 3 \text{ conv} \end{bmatrix} \times 6$
Transition Layer (1)	$56 \times 56$			$1 \times 1$ conv	
	$28 \times 28$			$2 \times 2$ average pool, stride 2	
Dense Block (2)	$28 \times 28$	$\begin{bmatrix} 1 \times 1 \text{ conv} \\ 3 \times 3 \text{ conv} \end{bmatrix} \times 12$	$\begin{bmatrix} 1 \times 1 \text{ conv} \\ 3 \times 3 \text{ conv} \end{bmatrix} \times 12$	$\begin{bmatrix} 1 \times 1 \text{ conv} \\ 3 \times 3 \text{ conv} \end{bmatrix} \times 12$	$\begin{bmatrix} 1 \times 1 \text{ conv} \\ 3 \times 3 \text{ conv} \end{bmatrix} \times 12$
Transition Layer (2)	$28 \times 28$			$1 \times 1$ conv	
	$14 \times 14$			$2 \times 2$ average pool, stride 2	
Dense Block (3)	$14 \times 14$	$\begin{bmatrix} 1 \times 1 \text{ conv} \\ 3 \times 3 \text{ conv} \end{bmatrix} \times 24$	$\begin{bmatrix} 1 \times 1 \text{ conv} \\ 3 \times 3 \text{ conv} \end{bmatrix} \times 32$	$\begin{bmatrix} 1 \times 1 \text{ conv} \\ 3 \times 3 \text{ conv} \end{bmatrix} \times 48$	$\begin{bmatrix} 1 \times 1 \text{ conv} \\ 3 \times 3 \text{ conv} \end{bmatrix} \times 64$
Transition Layer (3)	$14 \times 14$			$1 \times 1$ conv	
	$7 \times 7$			$2 \times 2$ average pool, stride 2	
Dense Block (4)	$7 \times 7$	$\begin{bmatrix} 1 \times 1 \text{ conv} \\ 3 \times 3 \text{ conv} \end{bmatrix} \times 16$	$\begin{bmatrix} 1 \times 1 \text{ conv} \\ 3 \times 3 \text{ conv} \end{bmatrix} \times 32$	$\begin{bmatrix} 1 \times 1 \text{ conv} \\ 3 \times 3 \text{ conv} \end{bmatrix} \times 32$	$\begin{bmatrix} 1 \times 1 \text{ conv} \\ 3 \times 3 \text{ conv} \end{bmatrix} \times 48$
Classification Layer	$1 \times 1$			$7 \times 7$ global average pool	
				1000D fully-connected, softmax	

**Table 1:** DenseNet architectures for ImageNet. The growth rate for all the networks is  $k = 32$ . Note that each “conv” layer shown in the table corresponds the sequence BN-ReLU-Conv.

## 02

Densely Connected Convolutional Networks  
DenseNet

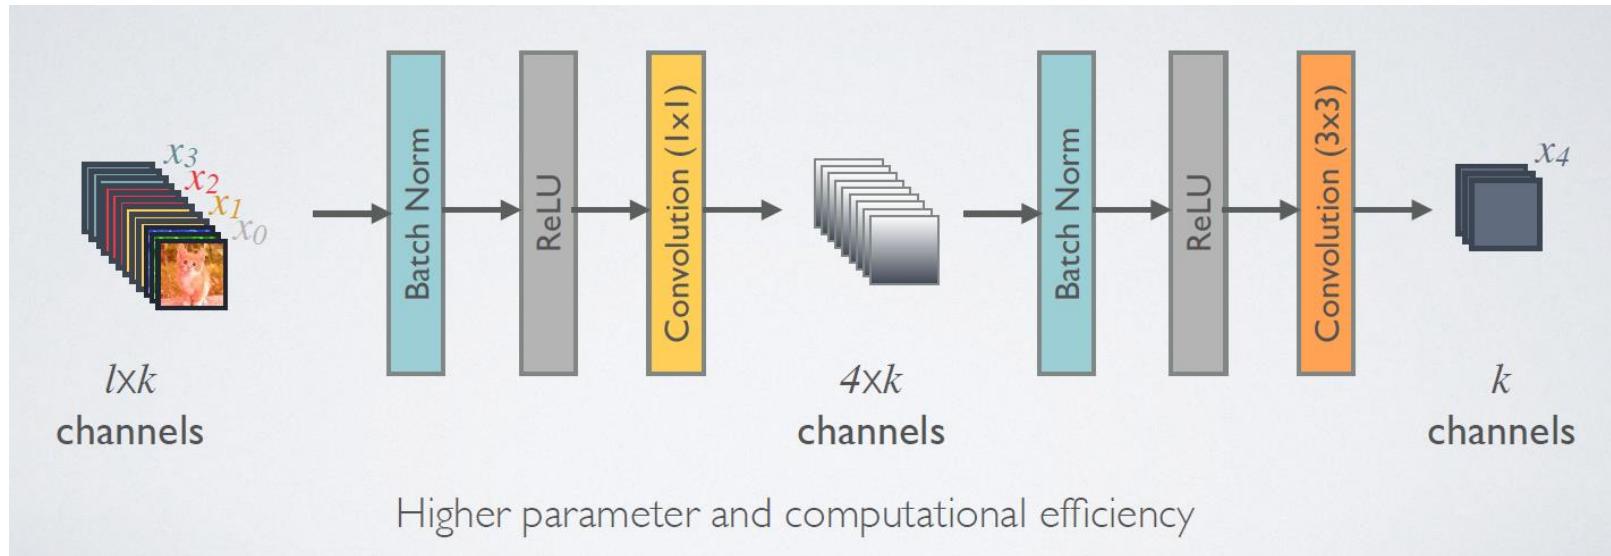
[Huang et al, 2016]  
Best Paper in CVPR 2017

Layers	Output Size	DenseNet-121	DenseNet-169	DenseNet-201	DenseNet-264
Convolution	$112 \times 112$		$7 \times 7$ conv, stride 2		
Pooling	$56 \times 56$		$3 \times 3$ max pool, stride 2		
Dense Block (1)	$56 \times 56$	$\begin{bmatrix} 1 \times 1 \text{ conv} \\ 3 \times 3 \text{ conv} \end{bmatrix} \times 6$	$\begin{bmatrix} 1 \times 1 \text{ conv} \\ 3 \times 3 \text{ conv} \end{bmatrix} \times 6$	$\begin{bmatrix} 1 \times 1 \text{ conv} \\ 3 \times 3 \text{ conv} \end{bmatrix} \times 6$	$\begin{bmatrix} 1 \times 1 \text{ conv} \\ 3 \times 3 \text{ conv} \end{bmatrix} \times 6$
Transition Layer (1)	$56 \times 56$			$1 \times 1$ conv	
	$28 \times 28$			$2 \times 2$ average pool, stride 2	
Dense Block (2)	$28 \times 28$	$\begin{bmatrix} 1 \times 1 \text{ conv} \\ 3 \times 3 \text{ conv} \end{bmatrix} \times 12$	$\begin{bmatrix} 1 \times 1 \text{ conv} \\ 3 \times 3 \text{ conv} \end{bmatrix} \times 12$	$\begin{bmatrix} 1 \times 1 \text{ conv} \\ 3 \times 3 \text{ conv} \end{bmatrix} \times 12$	$\begin{bmatrix} 1 \times 1 \text{ conv} \\ 3 \times 3 \text{ conv} \end{bmatrix} \times 12$
Transition Layer (2)	$28 \times 28$			$1 \times 1$ conv	
	$14 \times 14$			$2 \times 2$ average pool, stride 2	
Dense Block (3)	$14 \times 14$	$\begin{bmatrix} 1 \times 1 \text{ conv} \\ 3 \times 3 \text{ conv} \end{bmatrix} \times 24$	$\begin{bmatrix} 1 \times 1 \text{ conv} \\ 3 \times 3 \text{ conv} \end{bmatrix} \times 32$	$\begin{bmatrix} 1 \times 1 \text{ conv} \\ 3 \times 3 \text{ conv} \end{bmatrix} \times 48$	$\begin{bmatrix} 1 \times 1 \text{ conv} \\ 3 \times 3 \text{ conv} \end{bmatrix} \times 64$
Transition Layer (3)	$14 \times 14$			$1 \times 1$ conv	
	$7 \times 7$			$2 \times 2$ average pool, stride 2	
Dense Block (4)	$7 \times 7$	$\begin{bmatrix} 1 \times 1 \text{ conv} \\ 3 \times 3 \text{ conv} \end{bmatrix} \times 16$	$\begin{bmatrix} 1 \times 1 \text{ conv} \\ 3 \times 3 \text{ conv} \end{bmatrix} \times 32$	$\begin{bmatrix} 1 \times 1 \text{ conv} \\ 3 \times 3 \text{ conv} \end{bmatrix} \times 32$	$\begin{bmatrix} 1 \times 1 \text{ conv} \\ 3 \times 3 \text{ conv} \end{bmatrix} \times 48$
Classification Layer	$1 \times 1$			$7 \times 7$ global average pool	
				1000D fully-connected, softmax	

**Table 1:** DenseNet architectures for ImageNet. The growth rate for all the networks is  $k = 32$ . Note that each “conv” layer shown in the table corresponds the sequence BN-ReLU-Conv.

**$K$  : Growth Rate  
(Hyperparameter)**

**DenseNet can have very narrow layers**



$H_l(\cdot)$ :  
BN → ReLU → 3x3 conv

$H_l'(\cdot)$ :  
BN → ReLU → 1x1 conv → BN → ReLU → 3x3 conv

Computational efficiency ↑

## Simple DenseNet

**Convolution** :  $7 \times 7$  conv

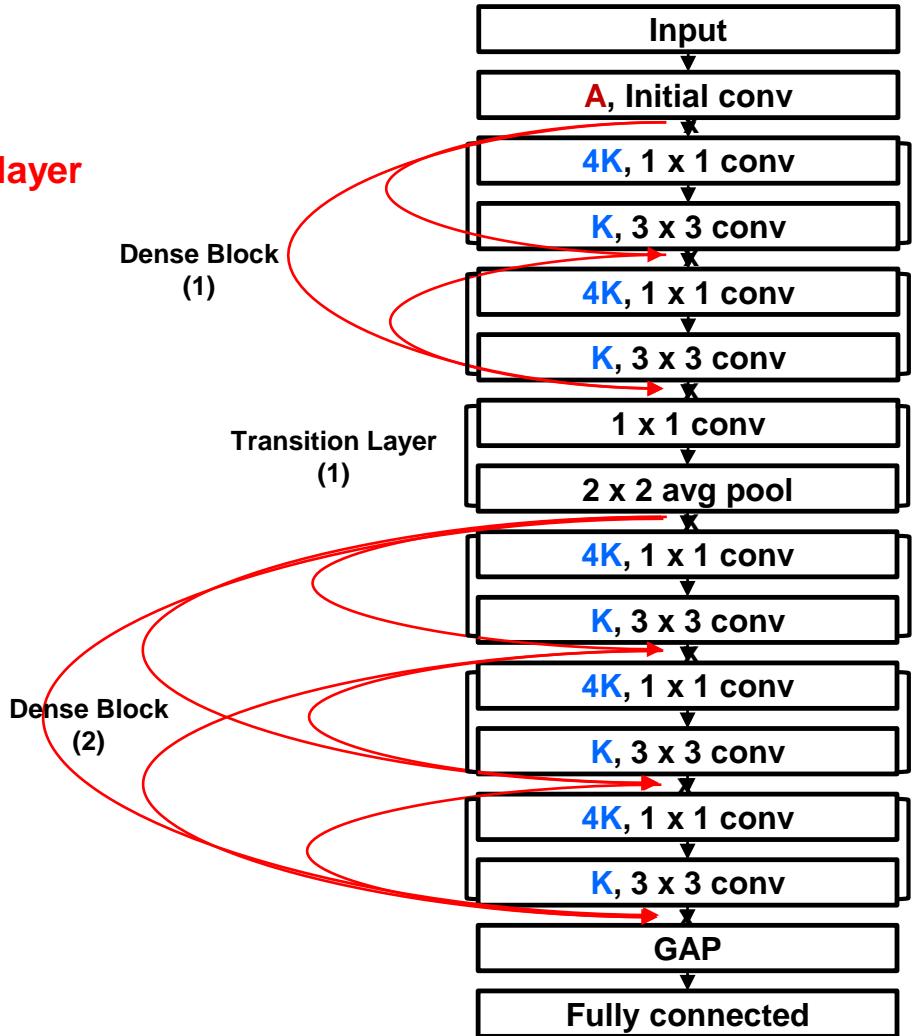
**Dense Block (1)** :  $\begin{bmatrix} 1 \times 1 \text{ conv} \\ 3 \times 3 \text{ conv} \end{bmatrix} \times 2$

**Transition Layer (1)** :  $1 \times 1$  conv  
 $2 \times 2$  average pool

**Dense Block (2)** :  $\begin{bmatrix} 1 \times 1 \text{ conv} \\ 3 \times 3 \text{ conv} \end{bmatrix} \times 3$

**Classification Layer** : global average pool  
 2D fully-connected, softmax

**Bottleneck layer**



## Simple DenseNet

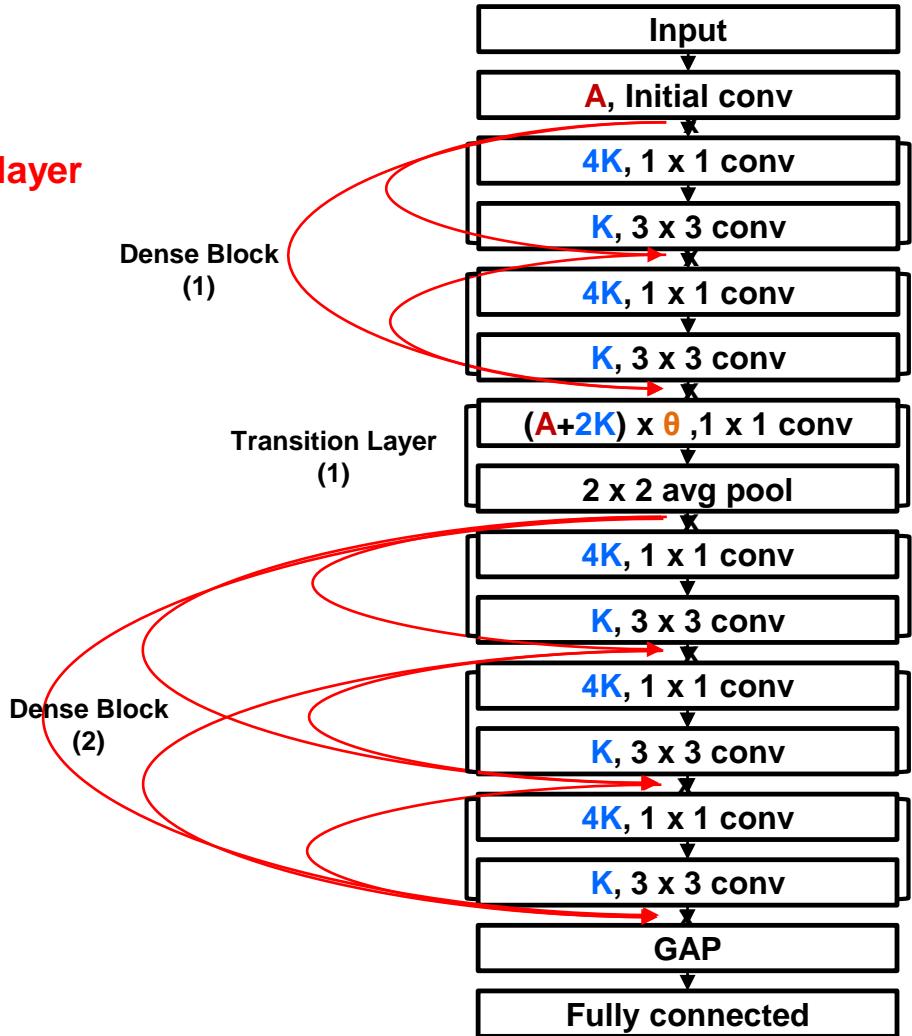
**Convolution** :  $7 \times 7$  conv

**Dense Block (1)** :  $\begin{bmatrix} 1 \times 1 \text{ conv} \\ 3 \times 3 \text{ conv} \end{bmatrix} \times 2$

**Transition Layer (1)** :  $1 \times 1$  conv  
 $2 \times 2$  average pool

**Dense Block (2)** :  $\begin{bmatrix} 1 \times 1 \text{ conv} \\ 3 \times 3 \text{ conv} \end{bmatrix} \times 3$

**Classification Layer** : global average pool  
 2D fully-connected, softmax



## Simple DenseNet

**Convolution** :  $7 \times 7$  conv

**Dense Block (1)** :  $\begin{bmatrix} 1 \times 1 \text{ conv} \\ 3 \times 3 \text{ conv} \end{bmatrix} \times 2$

**Transition Layer (1)** :  $1 \times 1$  conv  
 $2 \times 2$  average pool

**Dense Block (2)** :  $\begin{bmatrix} 1 \times 1 \text{ conv} \\ 3 \times 3 \text{ conv} \end{bmatrix} \times 3$

**Classification Layer** : global average pool  
 2D fully-connected, softmax

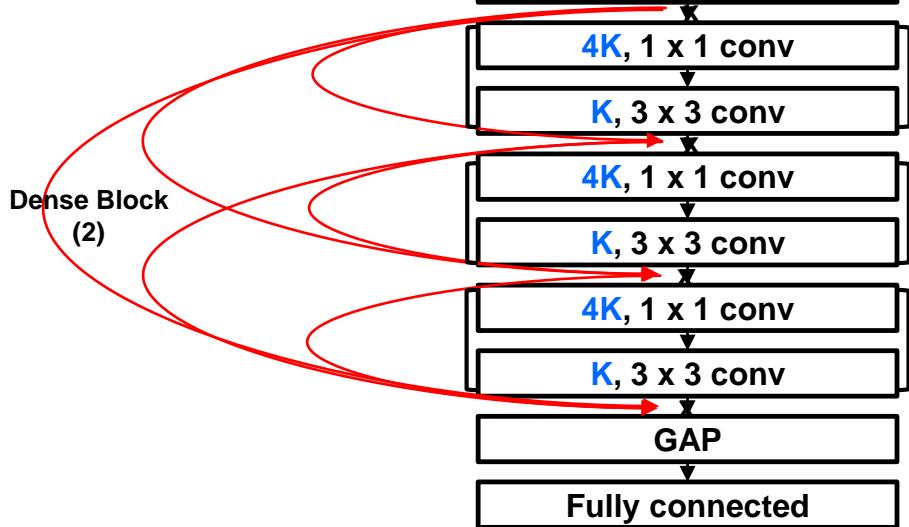
**Compression  $\theta$**   
 $(\theta = 0.5)$

Dense Block (1)

Transition Layer (1)

Dense Block (2)

Bottleneck layer



Method	Depth	Params	C10	C10+	C100	C100+	SVHN
Network in Network [22]	-	-	10.41	8.81	35.68	-	2.35
All-CNN [32]	-	-	9.08	7.25	-	33.71	-
Deeply Supervised Net [20]	-	-	9.69	7.97	-	34.57	1.92
Highway Network [34]	-	-	-	7.72	-	32.39	-
FractalNet [17]	21	38.6M	10.18	5.22	35.34	23.30	2.01
with Dropout/Drop-path	21	38.6M	7.33	4.60	28.20	23.73	1.87
ResNet [11]	110	1.7M	-	6.61	-	-	-
ResNet (reported by [13])	110	1.7M	13.63	6.41	44.74	27.22	2.01
ResNet with Stochastic Depth [13]	110	1.7M	11.66	5.23	37.80	24.58	1.75
	1202	10.2M	-	4.91	-	-	-
Wide ResNet [42]	16	11.0M	-	4.81	-	22.07	-
	28	36.5M	-	4.17	-	20.50	-
with Dropout	16	2.7M	-	-	-	-	1.64
ResNet (pre-activation) [12]	164	1.7M	11.26*	5.46	35.58*	24.33	-
	1001	10.2M	10.56*	4.62	33.47*	22.71	-
DenseNet ( $k = 12$ )	40	1.0M	<b>7.00</b>	5.24	<b>27.55</b>	24.42	1.79
DenseNet ( $k = 12$ )	100	7.0M	<b>5.77</b>	<b>4.10</b>	<b>23.79</b>	<b>20.20</b>	1.67
DenseNet ( $k = 24$ )	100	27.2M	<b>5.83</b>	<b>3.74</b>	<b>23.42</b>	<b>19.25</b>	<b>1.59</b>
DenseNet-BC ( $k = 12$ )	100	0.8M	<b>5.92</b>	4.51	<b>24.15</b>	22.27	1.76
DenseNet-BC ( $k = 24$ )	250	15.3M	<b>5.19</b>	<b>3.62</b>	<b>19.64</b>	<b>17.60</b>	1.74
DenseNet-BC ( $k = 40$ )	190	25.6M	-	<b>3.46</b>	-	<b>17.18</b>	-

**Table 2:** Error rates (%) on CIFAR and SVHN datasets.  $k$  denotes network's growth rate. Results that surpass all competing methods are **bold** and the overall best results are **blue**. "+" indicates standard data augmentation (translation and/or mirroring). \* indicates results run by ourselves. All the results of DenseNets without data augmentation (C10, C100, SVHN) are obtained using Dropout. DenseNets achieve lower error rates while using fewer parameters than ResNet. Without data augmentation, DenseNet performs better by a large margin.

CIFAR-10 : Training : 50,000, Testing: 10,000, Class : 10

CIFAR-100 : Training : 50,000, Testing: 10,000, Class : 100

CIFAR-10+ : CIFAR-10 + Mirroring/Shifting (Data augmentation)

CIFAR-100+ : CIFAR-100 + Mirroring/Shifting (Data augmentation)

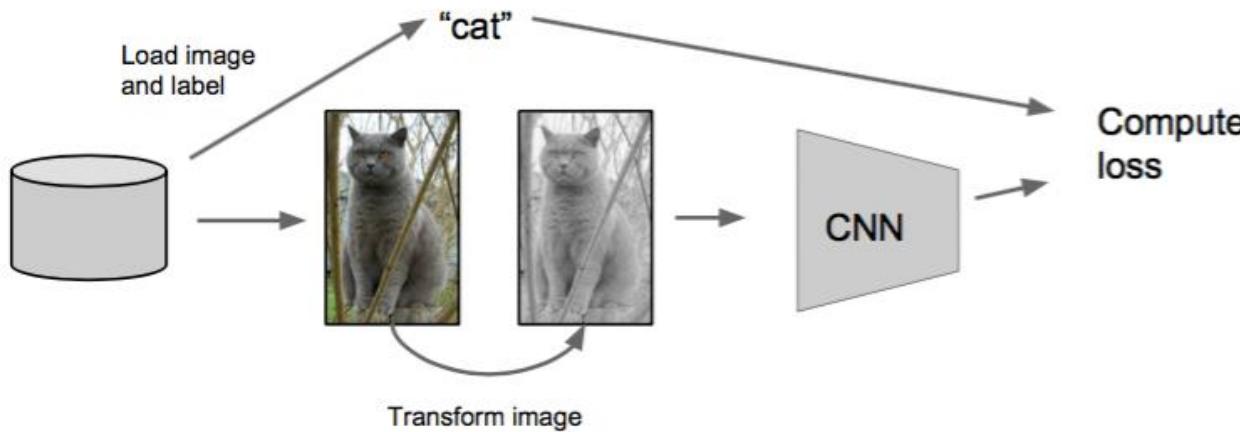
SVHN(Street View House Numbers) : Training : 73,257, Testing : 26,032

# 02

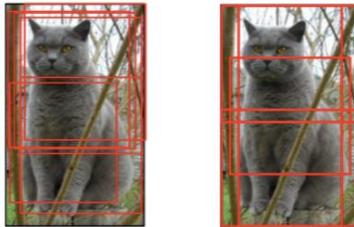
## Densely Connected Convolutional Networks DenseNet

[Huang et al, 2016]  
Best Paper in CVPR 2017

### Data augmentation

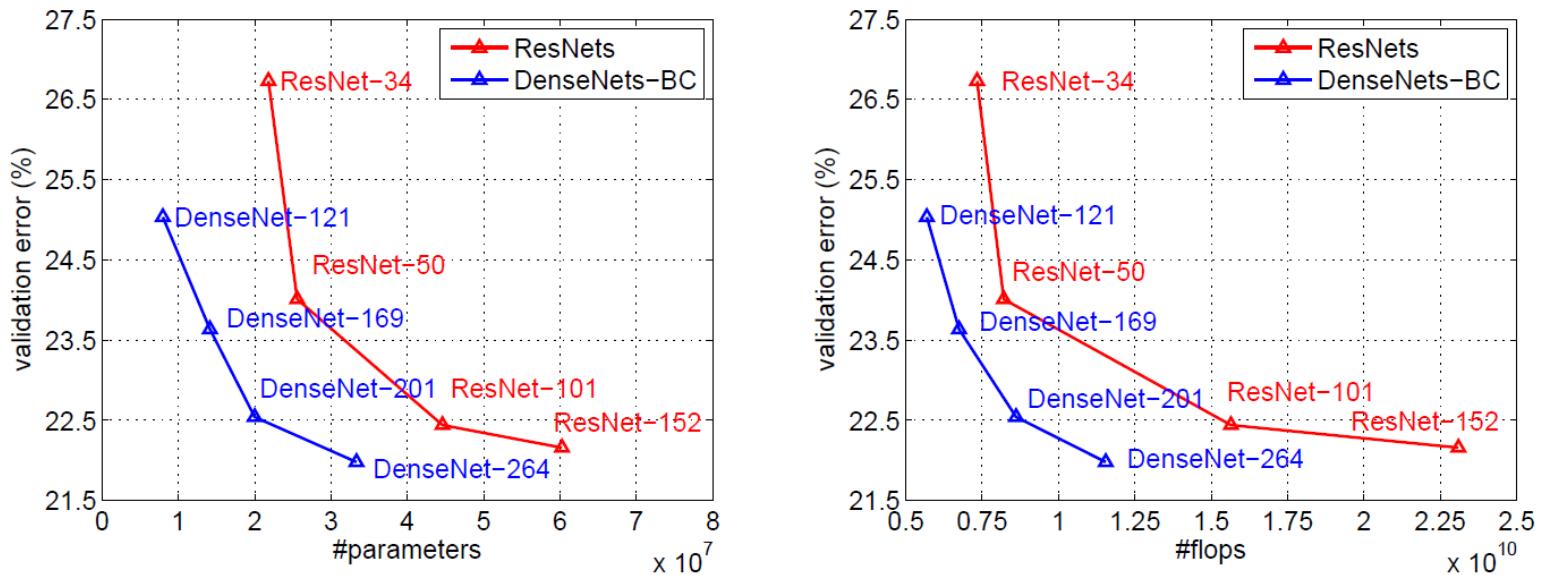


### Random crops



이미지를 랜덤하게 잘라서 샘플링에 사용하는 방식

[Huang et al, 2016]  
Best Paper in CVPR 2017

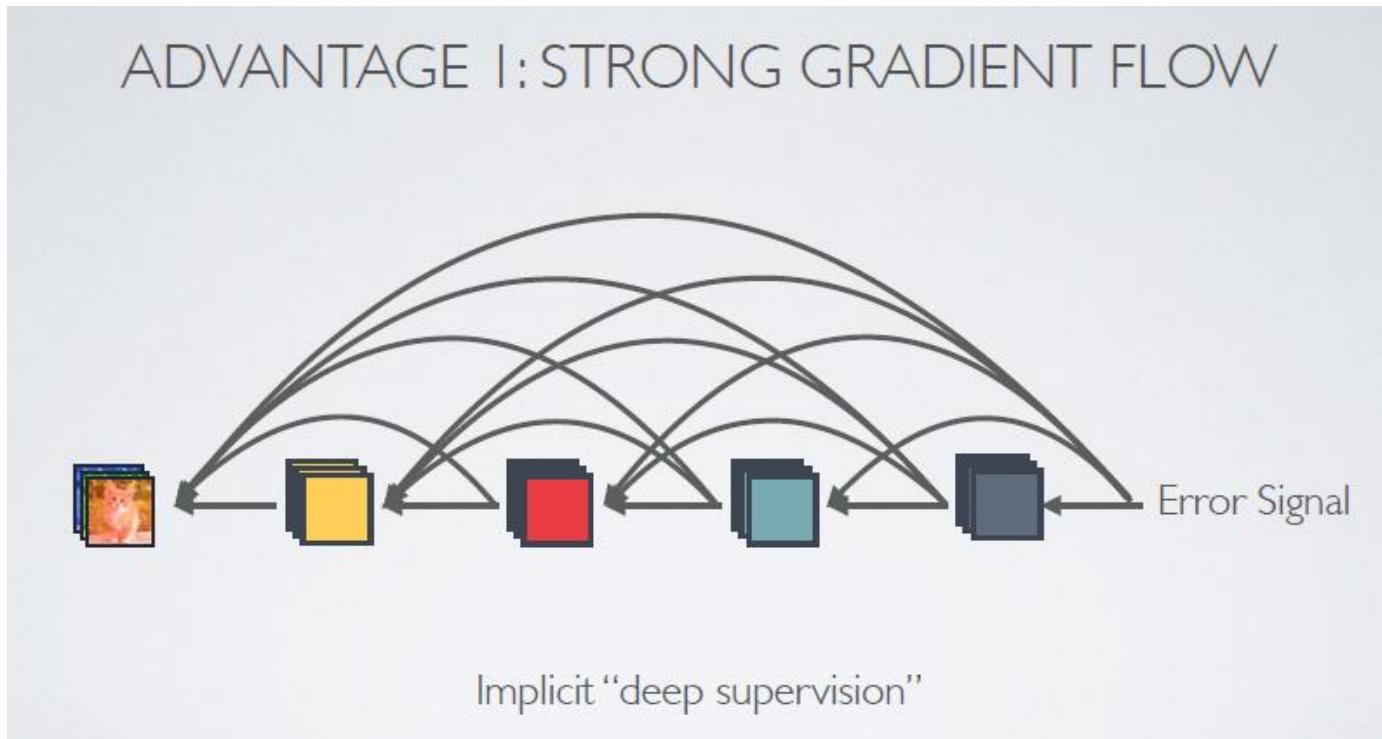


**Figure 3:** Comparison of the DenseNets and ResNets top-1 error rates (single-crop testing) on the ImageNet validation dataset as a function of learned parameters (*left*) and FLOPs during test-time (*right*).

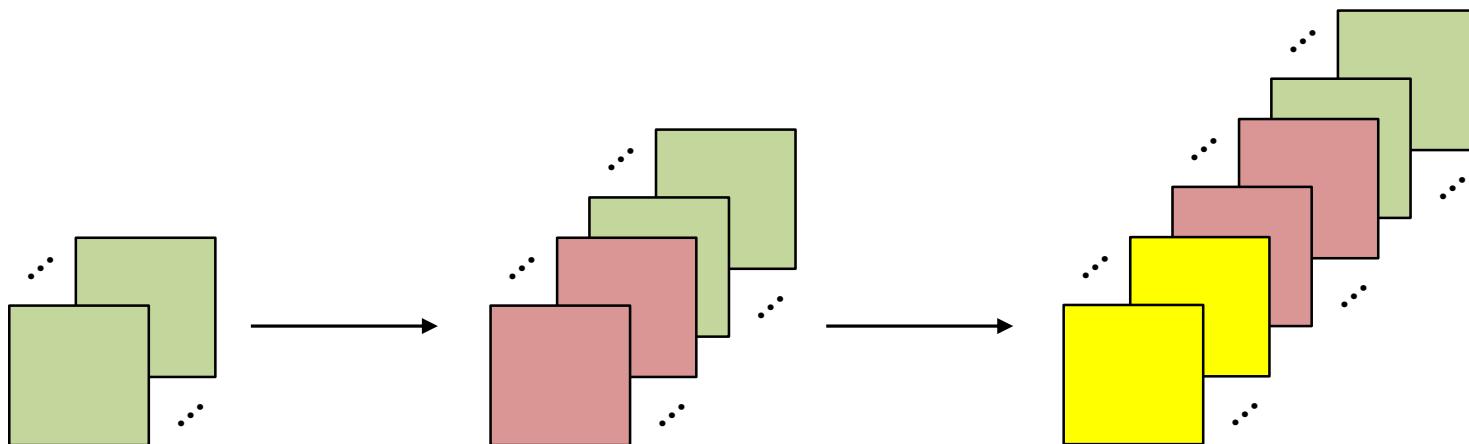
\* FLOPs(floating-point operations per second)  
컴퓨터의 연산 속도를 나타내는 단위

- 1. They alleviate the vanishing-gradient problem**
- 2. Strengthen feature propagation**
- 3. Encourage feature reuse**
- 4. Substantially reduce the number of parameters**

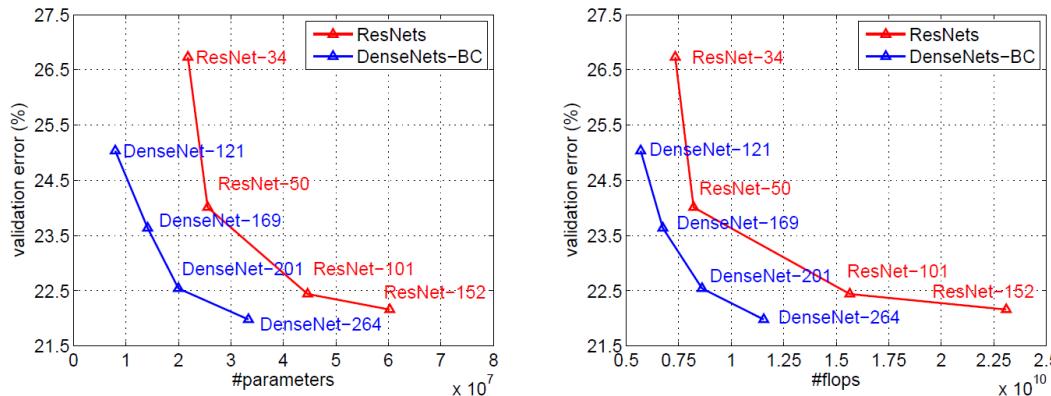
1. They alleviate the vanishing-gradient problem
2. Strengthen feature propagation



1. They alleviate the vanishing-gradient problem
2. Strengthen feature propagation
3. Encourage feature reuse



1. They alleviate the vanishing-gradient problem
2. Strengthen feature propagation
3. Encourage feature reuse
4. Substantially reduce the number of parameters



**Figure 3:** Comparison of the DenseNets and ResNets top-1 error rates (single-crop testing) on the ImageNet validation dataset as a function of learned parameters (left) and FLOPs during test-time (right).

**Q & A**