

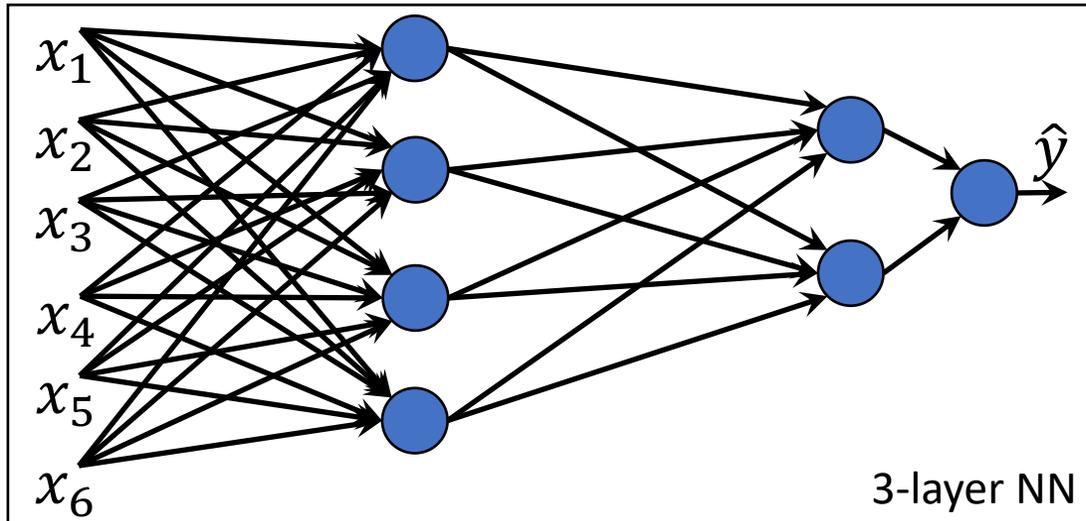
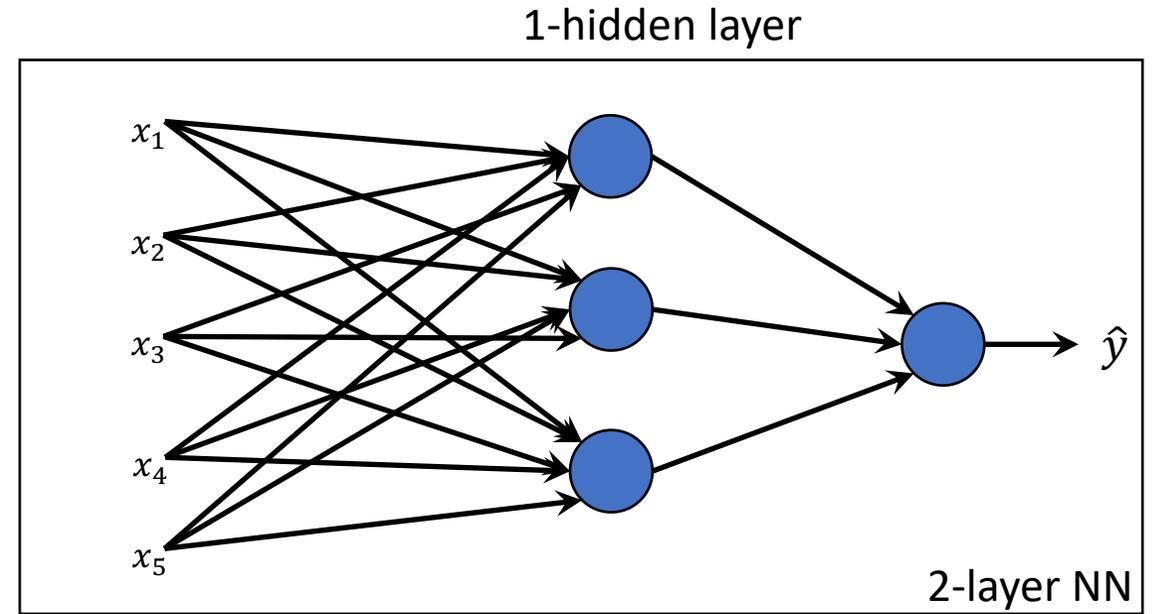
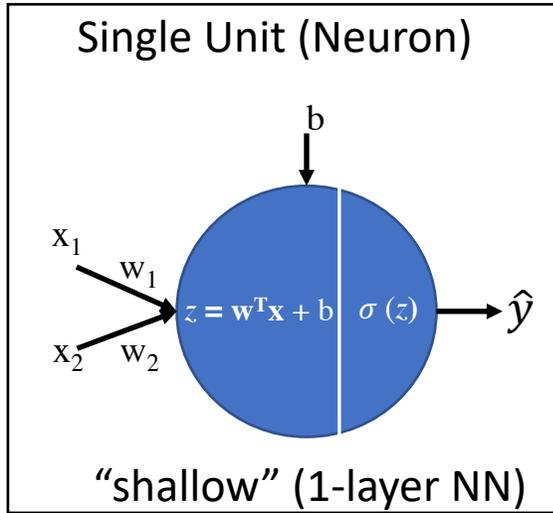
# Introduction to Computer Vision

## Intro to Convolutional Neural Networks: Part 3

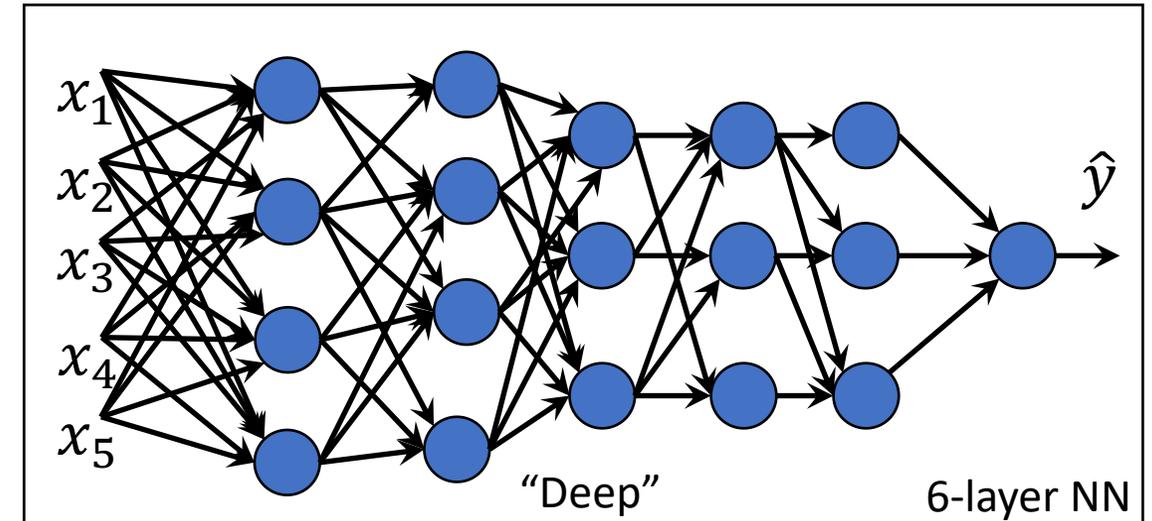


Dr. Sedat Ozer

# Deep vs. Shallow



2-hidden layers



5-hidden layers

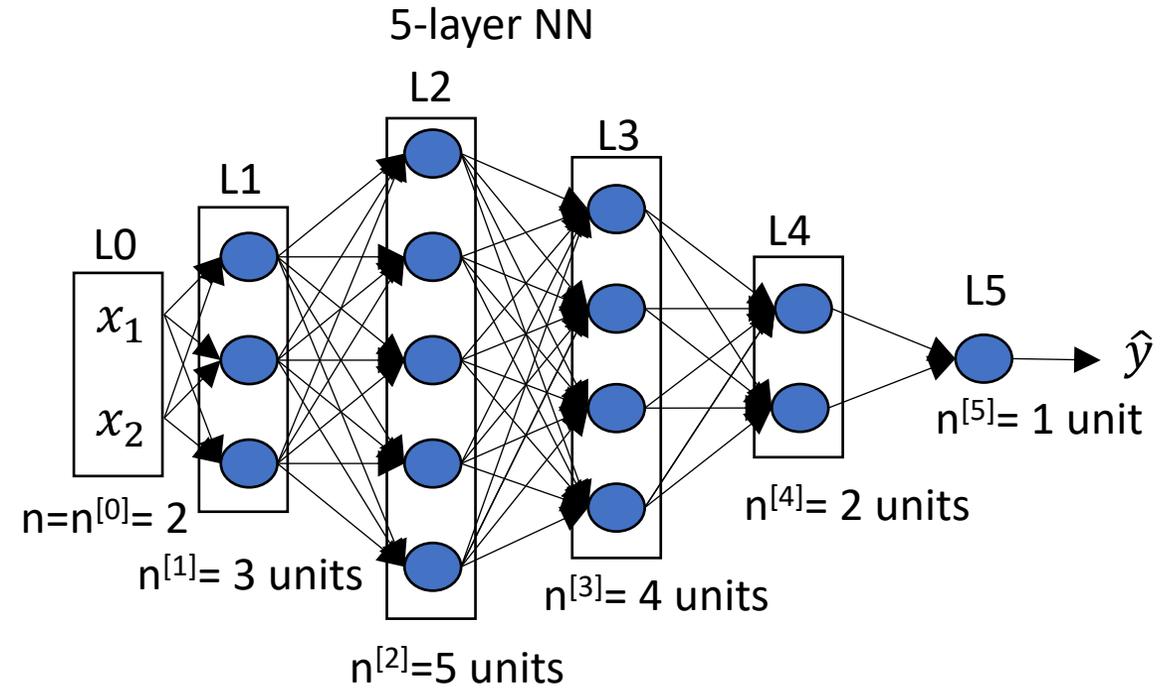
**Which one is deeper?**

# A deep NN: Notation

$$A^{[0]} = X = \begin{bmatrix} | & | & & | \\ \mathbf{x}^{(1)} & \mathbf{x}^{(2)} & \dots & \mathbf{x}^{(m)} \\ | & | & & | \end{bmatrix}_{(n^{[0]} \times m)}$$

$$A^{[l]} = \begin{bmatrix} | & | & & | \\ \mathbf{a}^{[l](1)} & \mathbf{a}^{[l](2)} & \dots & \mathbf{a}^{[l](m)} \\ | & | & & | \end{bmatrix}_{(n^{[l]} \times m)}$$

( $Z^{[l]}$  has the same shape as  $A^{[l]}$ )



$$Z^{[l]} = W^{[l]} A^{[l-1]} + \mathbf{b}^{[l]}$$

$$A^{[l]} = g(Z^{[l]})$$

Dims

$$\mathbf{z}^{[l]}: (n^{[l]} \times 1)$$

$$\mathbf{a}^{[l]}: (n^{[l]} \times 1)$$

$$\mathbf{W}^{[l]}: (n^{[l]} \times n^{[l-1]})$$

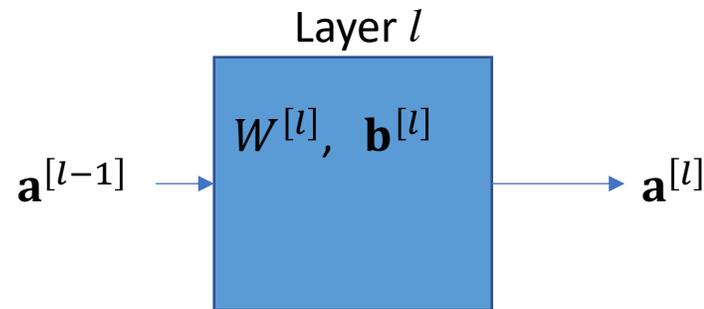
# Hyperparameters

- We have parameters:  $W^{[l]}$  ,  $\mathbf{b}^{[l]}$  for each layer (actual parameters)
- We also have additional parameters affecting the performance of our algorithm: Hyperparameters. I.e., the hyperparameters are the parameters that affect our actual parameters. Examples:
  - Learning rate,
  - number of iterations,
  - total number of hidden layers,
  - hidden unit numbers in each layer,
  - Choice of activation function,
  - Momentum,
  - mini batch size,
  - regularization parameters,
  - choice of cost function, choice of the optimization technique, ...

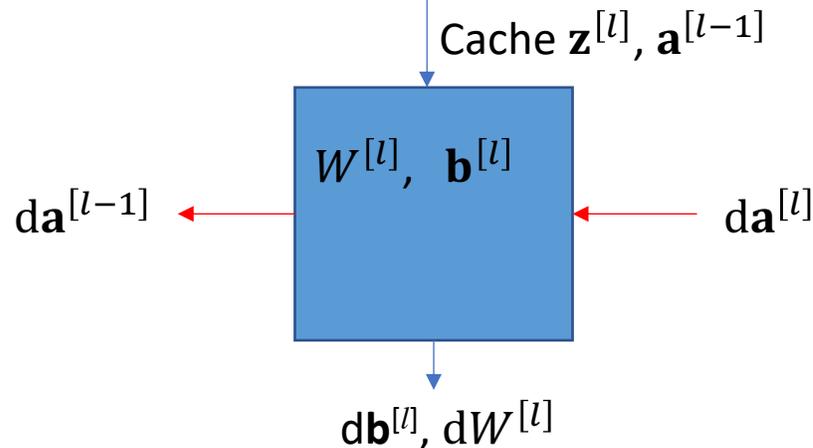
# Forward vs. Backward computation

- Forward direction computes the output and the loss.
- Backward direction computes the derivatives to find the update directions for the parameters: weights & biases.

Forward direction:



Backward direction:



Element-wise product

$$dz^{[l]} = da^{[l]} * g^{[l]'}(z^{[l]})$$

$$dW^{[l]} = dz^{[l]} a^{[l-1]}$$

$$db^{[l]} = dz^{[l]}$$

$$da^{[l-1]} = W^{[l]T} dz^{[l]}$$

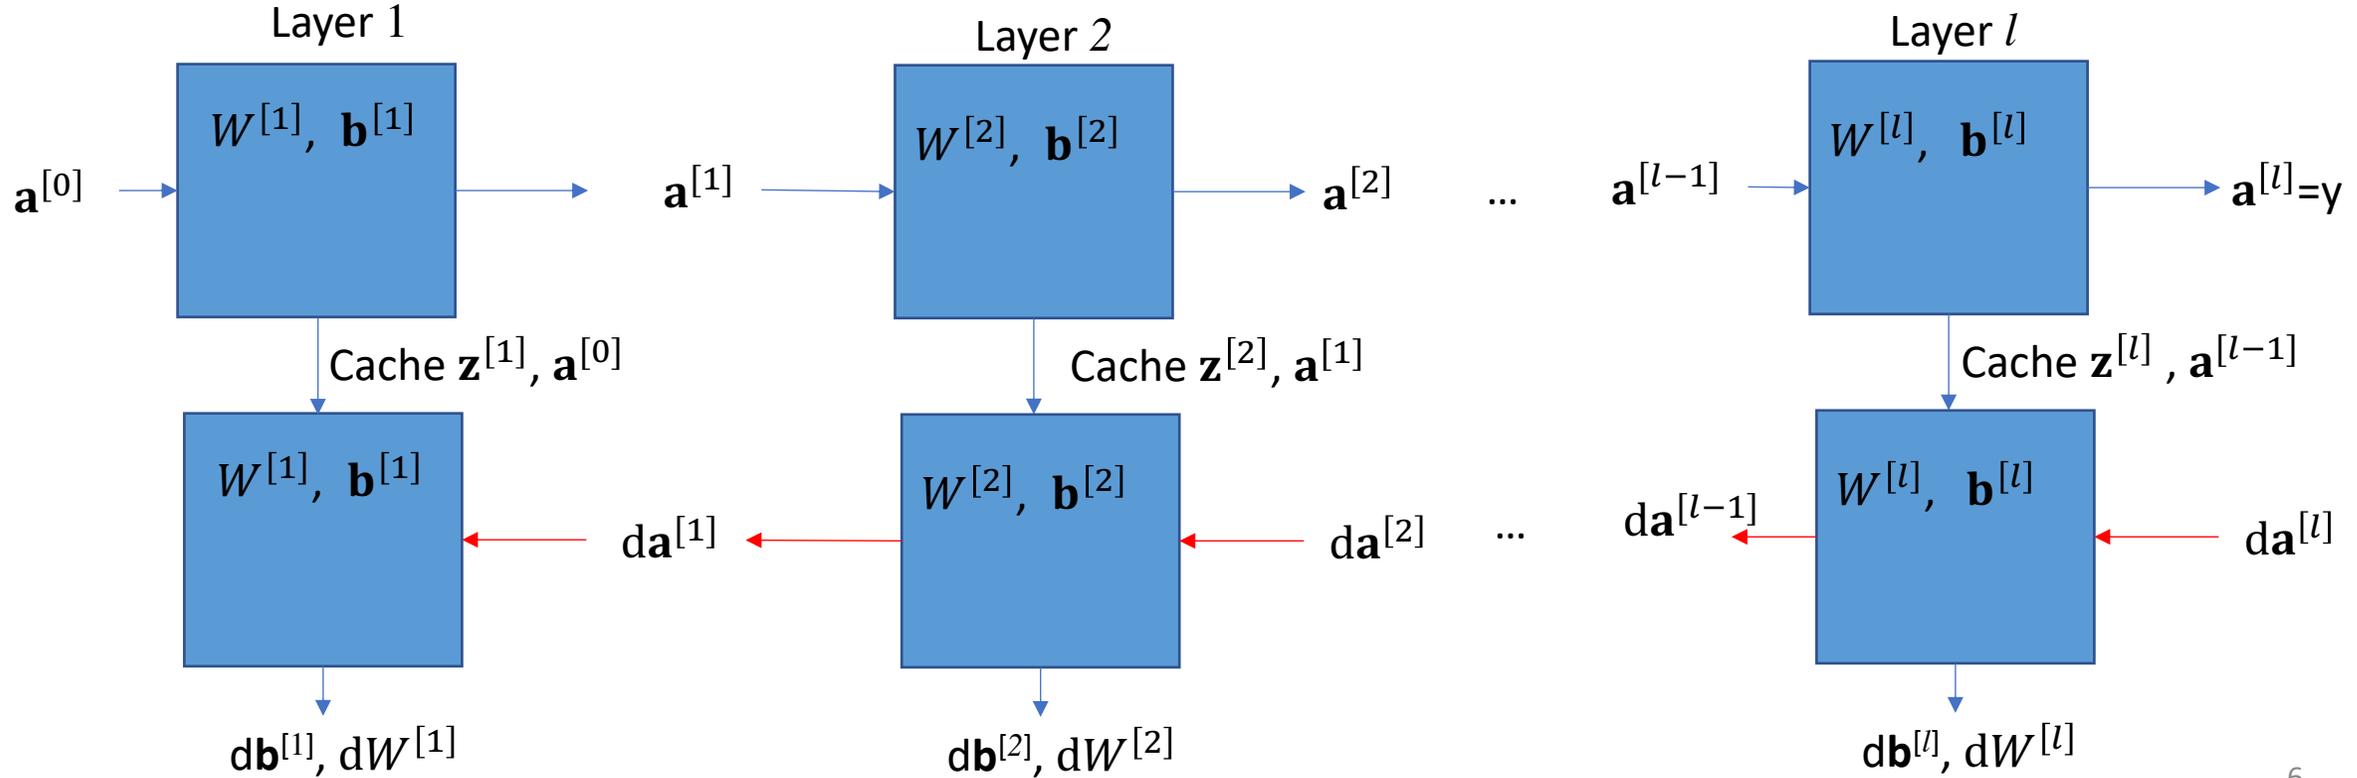
# Forward vs. Backward computation

$$dz^{[l]} = da^{[l]} * g^{[l]'}(z^{[l]})$$

$$dW^{[l]} = dz^{[l]} a^{[l-1]}$$

$$db^{[l]} = dz^{[l]}$$

$$da^{[l-1]} = W^{[l]T} dz^{[l]}$$



# Backward Propagation

Element-wise multiplication

$$d\mathbf{z}^{[l]} = d\mathbf{a}^{[l]} * g^{[l]'}(\mathbf{z}^{[l]})$$

$$dW^{[l]} = d\mathbf{z}^{[l]} \mathbf{a}^{[l-1]}$$

$$d\mathbf{b}^{[l]} = d\mathbf{z}^{[l]}$$

$$d\mathbf{a}^{[l-1]} = W^{[l]T} d\mathbf{z}^{[l]}$$

For a single sample

Element-wise multiplication

$$d\mathbf{Z}^{[l]} = d\mathbf{A}^{[l]} * g^{[l]'}(\mathbf{Z}^{[l]})$$

$$dW^{[l]} = \frac{1}{m} d\mathbf{Z}^{[l]} \mathbf{A}^{[l-1]T}$$

$$d\mathbf{b}^{[l]} = \frac{1}{m} \text{np.sum}(d\mathbf{Z}^{[l]}, \text{axis} = 1, \text{keepdims} = \text{True})$$

$$d\mathbf{A}^{[l-1]} = W^{[l]T} d\mathbf{Z}^{[l]}$$

For  $m$  data samples (“Vectorized” versions)

# Convolutional Neural Networks

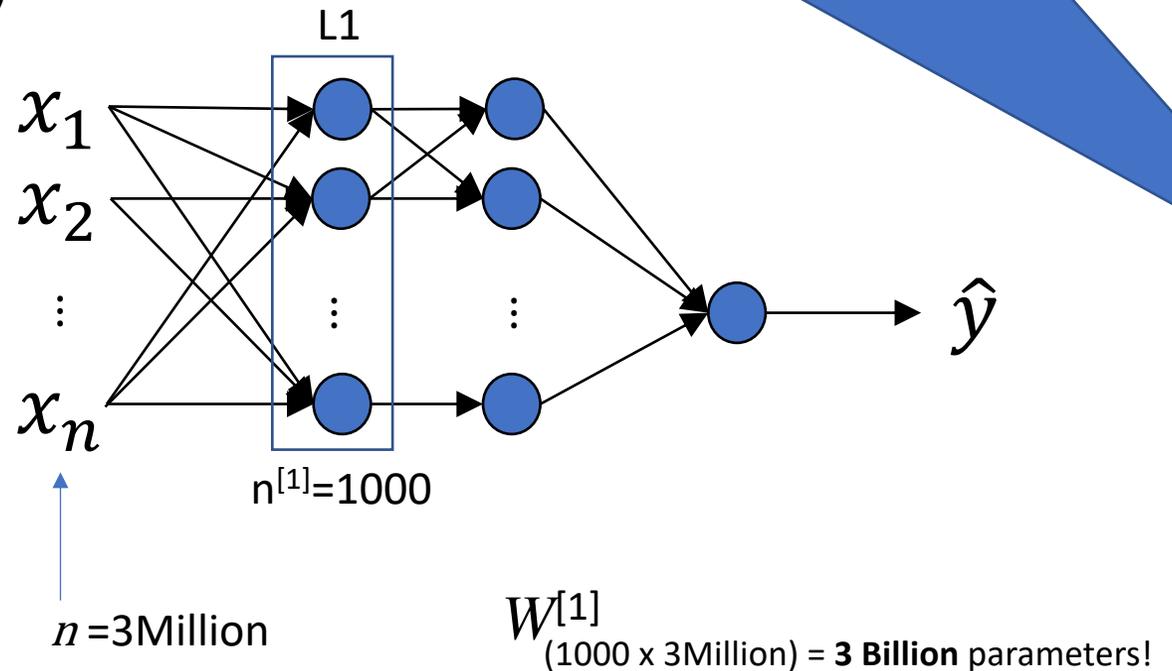
# Story so far...

- Up until now, we learned how to
- Lets consider the case where the

As an alternative approach:  
Use a convolutional NN !!



Problem: Is there a person in this picture?  
(answer: yes/no)

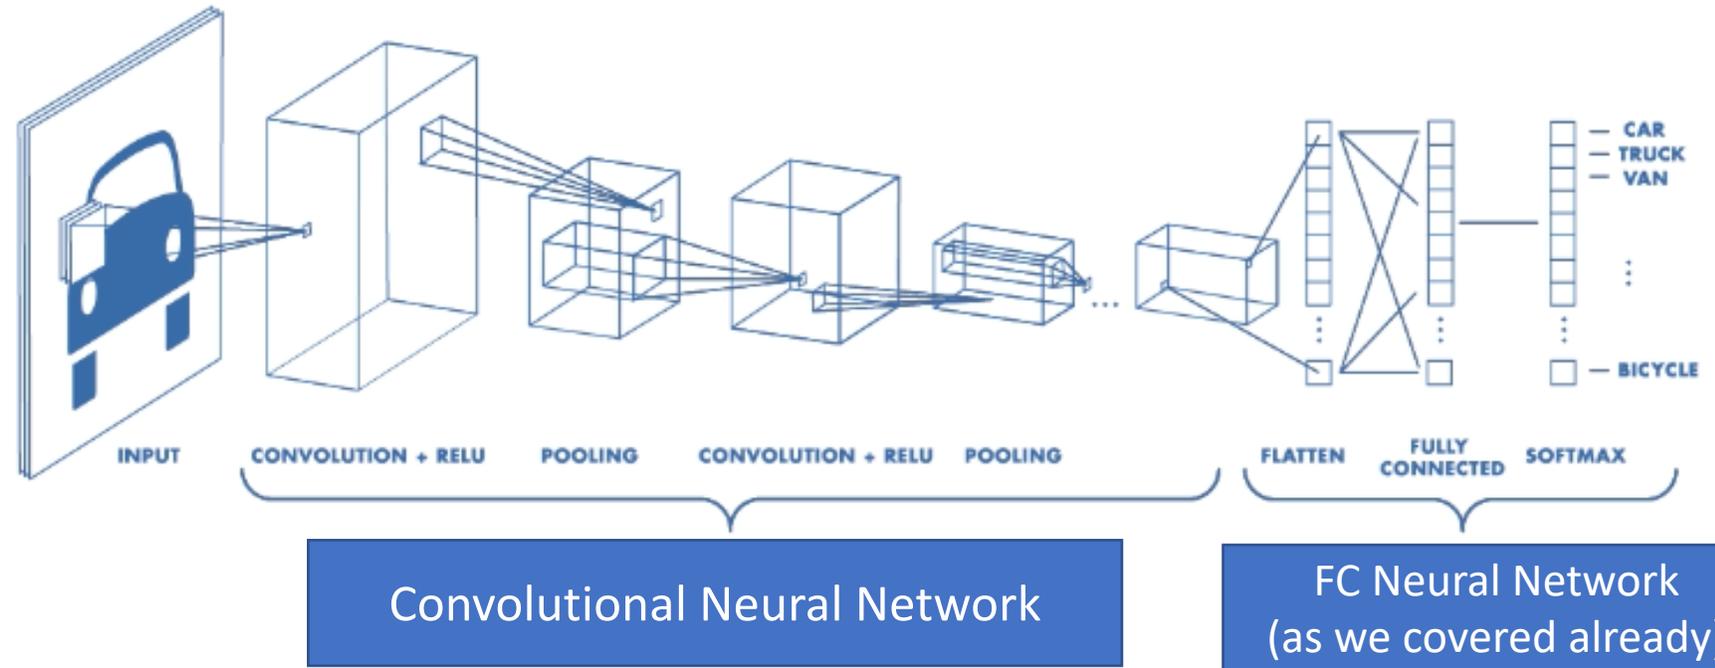


A high-resolution image:  $1000 \times 1000 \times 3$   
= (3Million x 1)

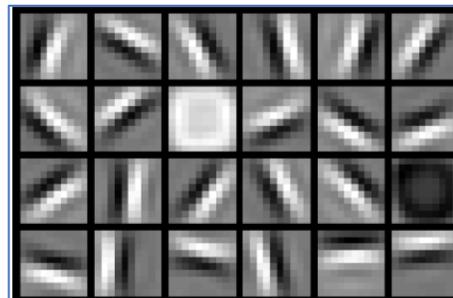
# Convolutional Neural Networks (CNN)

Includes many additional components - operations. Most common ones are:

- Filtering,
- Padding,
- Stride,
- Pooling.



# Features ?



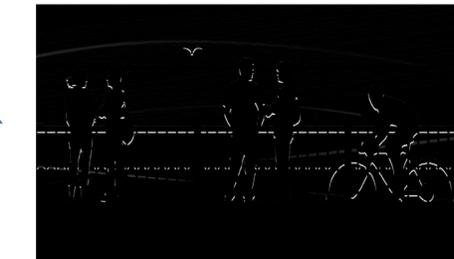
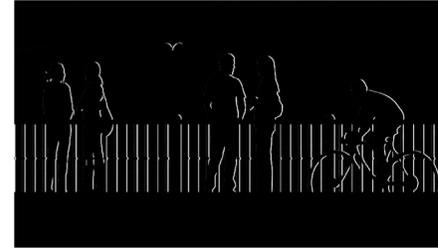
# Edge Detection

- To explain the Convolutional NNs, we will look at the edge detection example first.

# Edges



vertical edges

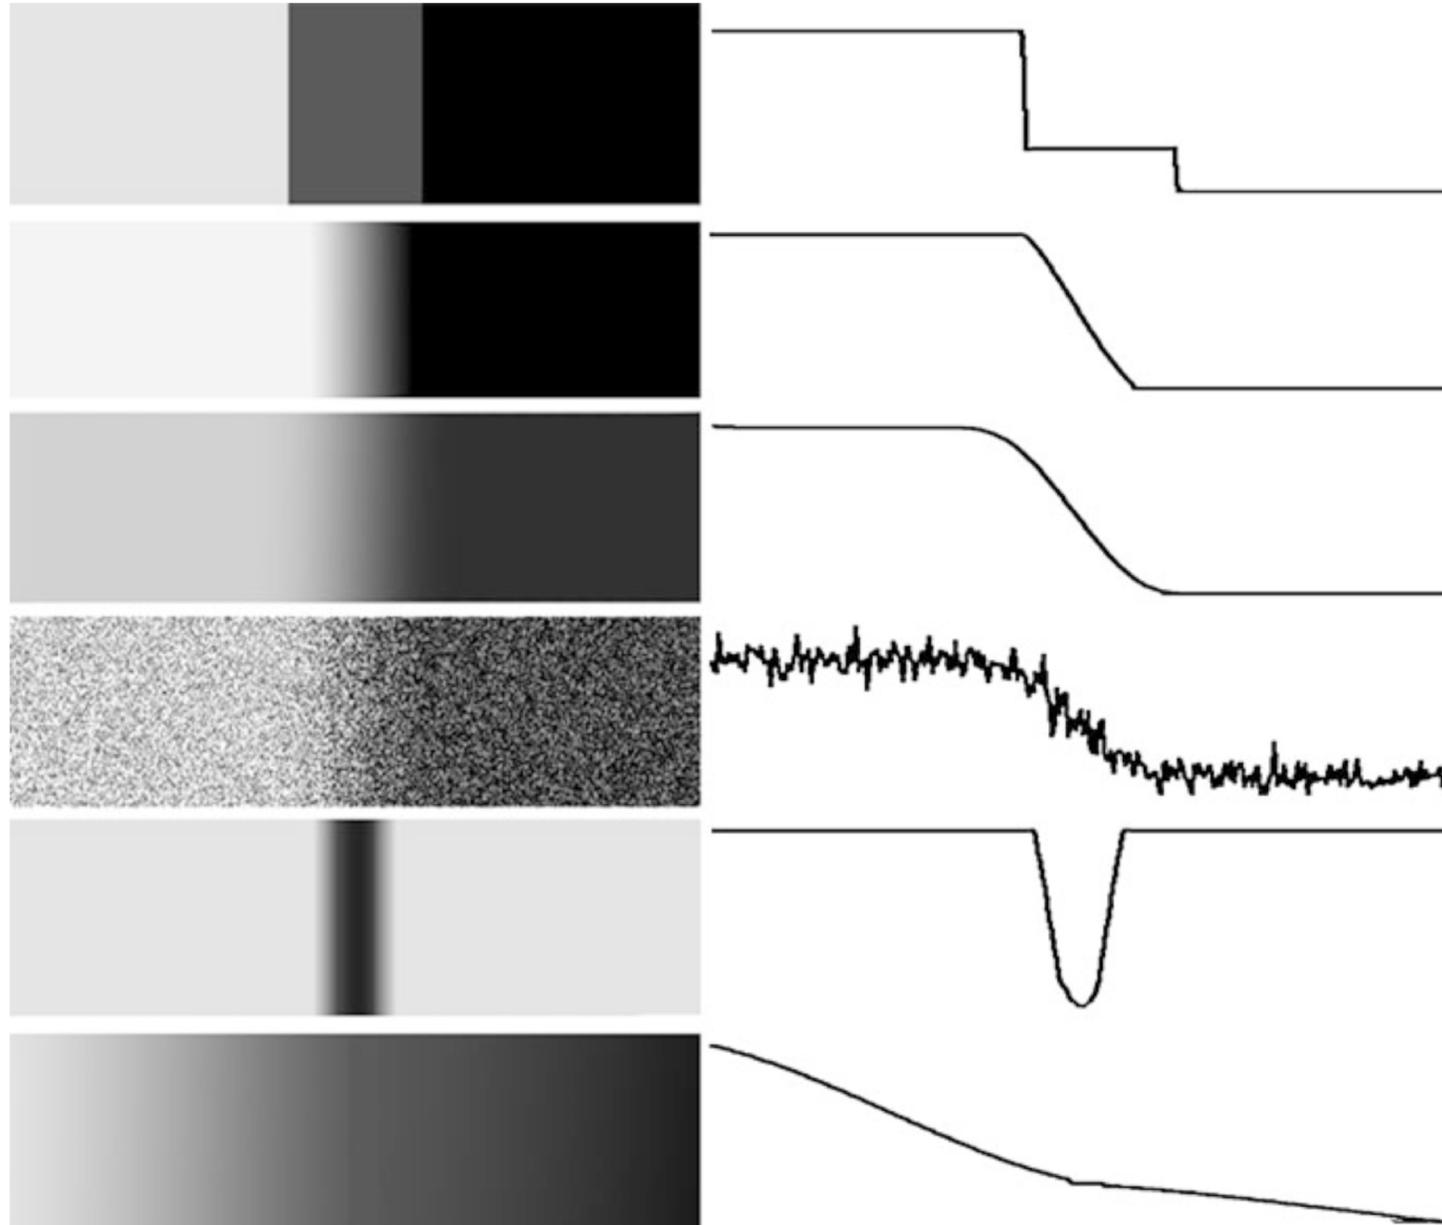


horizontal edges

# EDGES

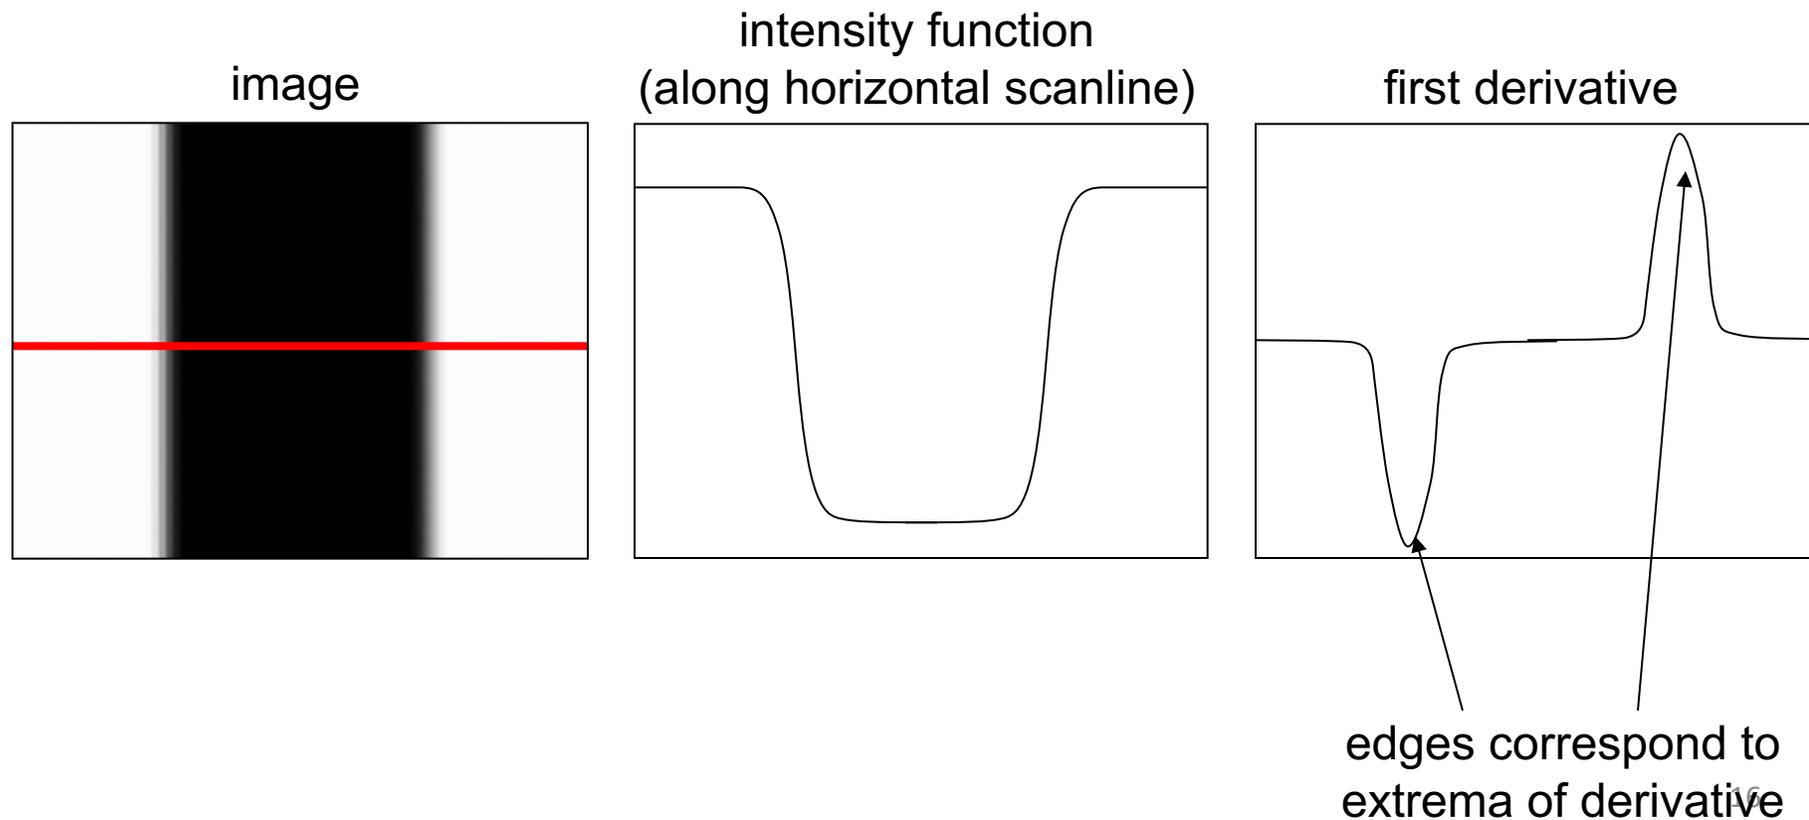
- **Discontinuities** in images are features that are often useful for initializing an image analysis procedure.
- Edges are important information for understanding an image; by removing “non-edge” data we also **simplify** the data.

# Edge Models



# Characterizing Edges

- An edge is a place of rapid change in the image intensity function



# Derivatives and Average

- **Derivative:** rate of change
- Examples:
  - Speed is a rate of change of a distance,  $X=V.t$
  - Acceleration is a rate of change of speed,  $V=a.t$

# Derivative

$$\frac{df}{dx} = \lim_{\Delta x \rightarrow 0} \frac{f(x) - f(x - \Delta x)}{\Delta x} = f'(x) = f_x$$

Example:  $y = x^2 + x^4$

$$\frac{dy}{dx} = 2x + 4x^3$$

# Discrete Derivative

$$\frac{df}{dx} = \lim_{\Delta x \rightarrow 0} \frac{f(x) - f(x - \Delta x)}{\Delta x} = f'(x)$$

$$\frac{df}{dx} = \frac{f(x) - f(x-1)}{1} = f'(x)$$

$$\frac{df}{dx} = f(x) - f(x-1) = f'(x)$$

# Discrete Derivative / Finite Difference

$$\frac{df}{dx} = \frac{f(x) - f(x-1)}{1} = f'(x) \quad \text{Backward difference}$$

$$\frac{df}{dx} = \frac{f(x) - f(x+1)}{-1} = f'(x) \quad \text{Forward difference}$$

$$\frac{df}{dx} = \frac{f(x+1) - f(x-1)}{2} = f'(x) \quad \text{Central difference}$$

# Example: Finite Difference

$$\begin{array}{cccccccc} f(x) = & 10 & 15 & 10 & 10 & 25 & 20 & 20 & 20 \\ & & [-1 & 1] & [-1 & 1] & [-1 & 1] & [-1 & 1] & [-1 & 1] & [-1 & 1] \\ f'(x) = & 0 & 5 & -5 & 0 & 15 & -5 & 0 & 0 \end{array}$$

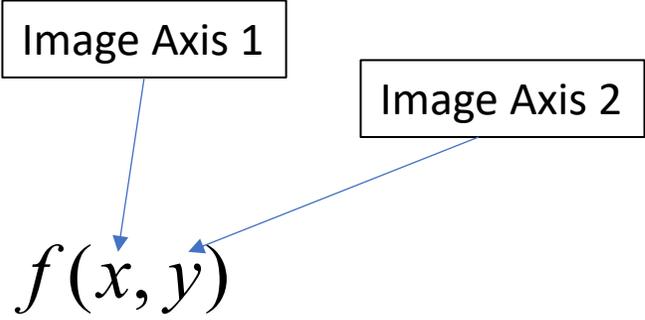
## Derivative Masks

Backward difference      [-1 1]

Forward difference      [1 -1]

Central difference      [-1 0 1]

# Derivative in 2-D

|                    |  |
|--------------------|--|
|                    |   |
| Given function     | $f(x, y)$  |
| Gradient vector    | $\nabla f(x, y) = \begin{bmatrix} \frac{\partial f(x, y)}{\partial x} \\ \frac{\partial f(x, y)}{\partial y} \end{bmatrix} = \begin{bmatrix} f_x \\ f_y \end{bmatrix}$ |
| Gradient magnitude | $ \nabla f(x, y)  = \sqrt{f_x^2 + f_y^2}$  |
| Gradient direction | $\theta = \tan^{-1} \frac{f_x}{f_y}$   |

# Derivative of Images

Derivative masks (filters):

$$f_x \Rightarrow \frac{1}{3} \begin{bmatrix} -1 & 0 & 1 \\ -1 & 0 & 1 \\ -1 & 0 & 1 \end{bmatrix} \quad f_y \Rightarrow \frac{1}{3} \begin{bmatrix} 1 & 1 & 1 \\ 0 & 0 & 0 \\ -1 & -1 & -1 \end{bmatrix}$$

$$I = \begin{bmatrix} 10 & 10 & 20 & 20 & 20 \\ 10 & 10 & 20 & 20 & 20 \\ 10 & 10 & 20 & 20 & 20 \\ 10 & 10 & 20 & 20 & 20 \\ 10 & 10 & 20 & 20 & 20 \end{bmatrix}$$

$$I_x = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & 10 & 10 & 0 & 0 \\ 0 & 10 & 10 & 0 & 0 \\ 0 & 10 & 10 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

# (Vertical) Edge Detection

|                |                |                 |                 |                 |                 |
|----------------|----------------|-----------------|-----------------|-----------------|-----------------|
| 4 <sup>1</sup> | 1 <sup>0</sup> | 4 <sup>-0</sup> | 0 <sup>-0</sup> | 1 <sup>-0</sup> | 2 <sup>-1</sup> |
| 2 <sup>1</sup> | 7 <sup>0</sup> | 4 <sup>-0</sup> | 3 <sup>-0</sup> | 0 <sup>-0</sup> | 1 <sup>-1</sup> |
| 0 <sup>1</sup> | 6 <sup>0</sup> | 5 <sup>-0</sup> | 4 <sup>-0</sup> | 2 <sup>-0</sup> | 2 <sup>-1</sup> |
| 5 <sup>1</sup> | 2 <sup>0</sup> | 1 <sup>-0</sup> | 0 <sup>-0</sup> | 4 <sup>-0</sup> | 2 <sup>-1</sup> |
| 1              | 5              | 3               | 5               | 6               | 6               |
| 3              | 2              | 2               | 1               | 7               | 5               |

Convolution



\*

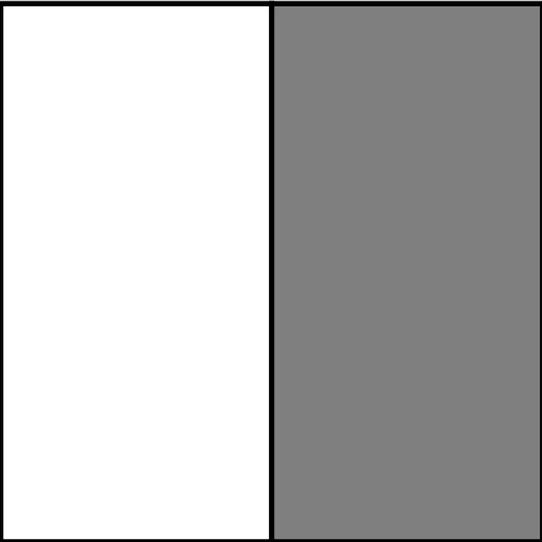
|   |   |    |
|---|---|----|
| 1 | 0 | -1 |
| 1 | 0 | -1 |
| 1 | 0 | -1 |

=

|    |    |     |    |
|----|----|-----|----|
| -7 | 7  | 10  | 2  |
| -3 | 8  | 4   | 2  |
| -3 | -4 | -3  | -1 |
| 3  | 3  | -11 | -7 |

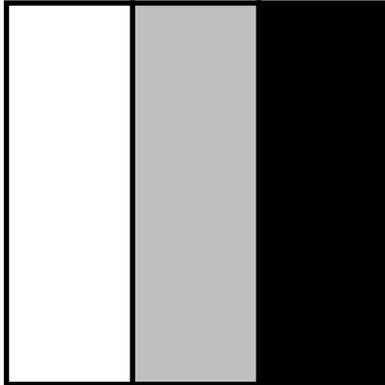
# Edge Detection (1)

|    |    |    |   |   |   |
|----|----|----|---|---|---|
| 20 | 20 | 20 | 0 | 0 | 0 |
| 20 | 20 | 20 | 0 | 0 | 0 |
| 20 | 20 | 20 | 0 | 0 | 0 |
| 20 | 20 | 20 | 0 | 0 | 0 |
| 20 | 20 | 20 | 0 | 0 | 0 |
| 20 | 20 | 20 | 0 | 0 | 0 |



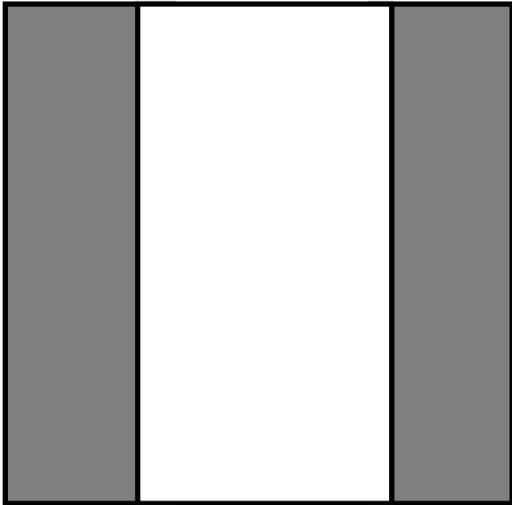
\*

|   |   |    |
|---|---|----|
| 1 | 0 | -1 |
| 1 | 0 | -1 |
| 1 | 0 | -1 |



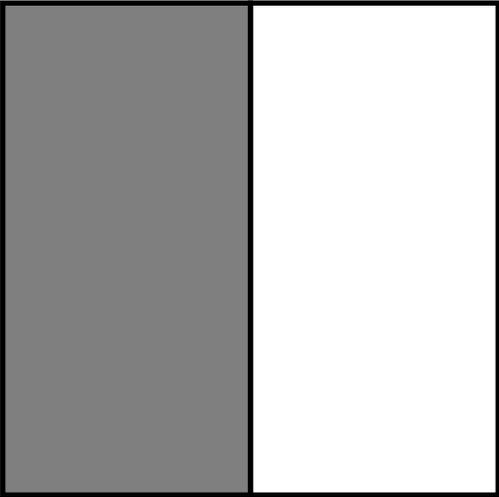
=

|   |    |    |   |
|---|----|----|---|
| 0 | 60 | 60 | 0 |
| 0 | 60 | 60 | 0 |
| 0 | 60 | 60 | 0 |
| 0 | 60 | 60 | 0 |



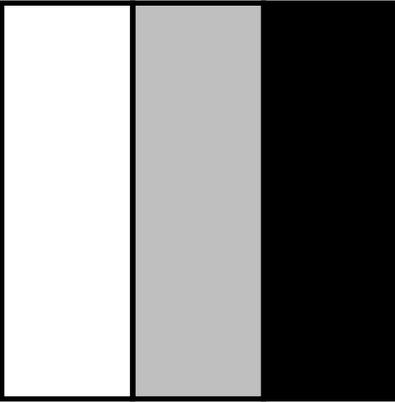
# Edge Detection (2)

|   |   |   |    |    |    |
|---|---|---|----|----|----|
| 0 | 0 | 0 | 20 | 20 | 20 |
| 0 | 0 | 0 | 20 | 20 | 20 |
| 0 | 0 | 0 | 20 | 20 | 20 |
| 0 | 0 | 0 | 20 | 20 | 20 |
| 0 | 0 | 0 | 20 | 20 | 20 |
| 0 | 0 | 0 | 20 | 20 | 20 |



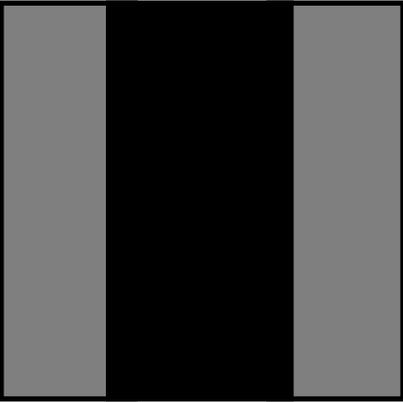
\*

|   |   |    |
|---|---|----|
| 1 | 0 | -1 |
| 1 | 0 | -1 |
| 1 | 0 | -1 |

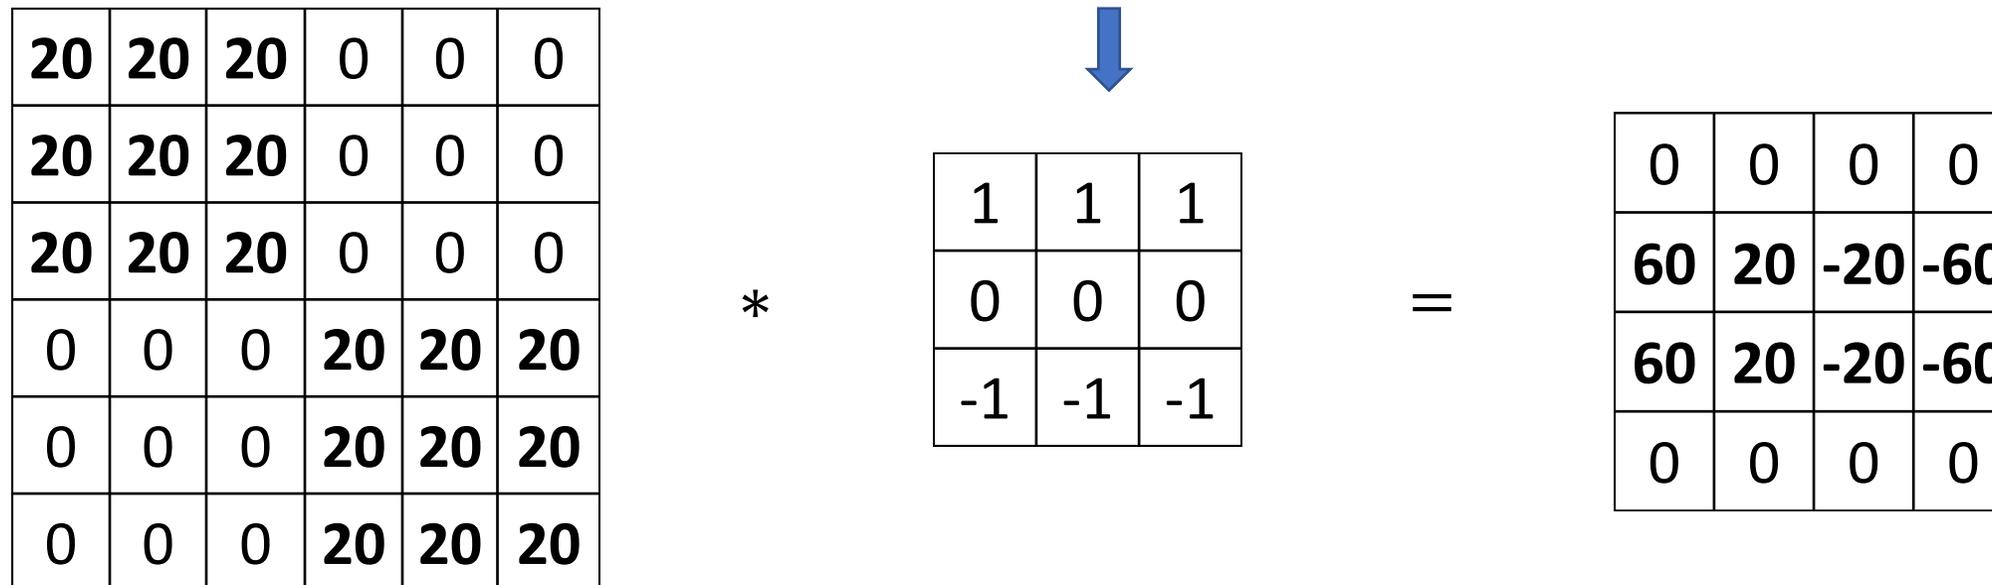
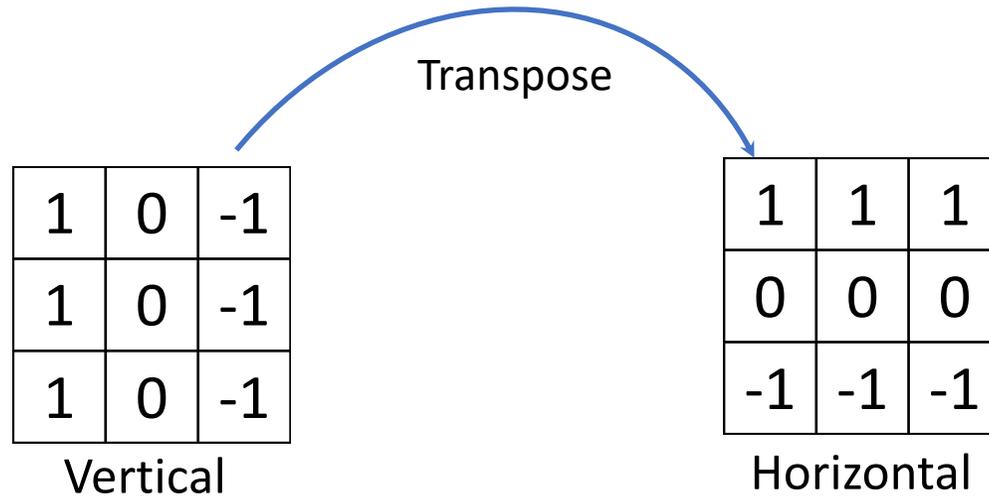


=

|   |     |     |   |
|---|-----|-----|---|
| 0 | -60 | -60 | 0 |
| 0 | -60 | -60 | 0 |
| 0 | -60 | -60 | 0 |
| 0 | -60 | -60 | 0 |



# Vertical vs. Horizontal Edge Detection



# More Filters

|   |   |    |
|---|---|----|
| 1 | 0 | -1 |
| 1 | 0 | -1 |
| 1 | 0 | -1 |

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

Sobel filter

|    |   |     |
|----|---|-----|
| 3  | 0 | -3  |
| 10 | 0 | -10 |
| 3  | 0 | -2  |

Scharr filter

# Which filter to pick? (Why to pick! Learn it)

|   |   |   |   |   |   |
|---|---|---|---|---|---|
| 4 | 1 | 4 | 0 | 1 | 2 |
| 2 | 7 | 4 | 3 | 0 | 1 |
| 0 | 6 | 5 | 4 | 2 | 2 |
| 5 | 2 | 1 | 0 | 4 | 2 |
| 1 | 5 | 3 | 5 | 6 | 6 |
| 3 | 2 | 2 | 1 | 7 | 5 |

6x6

\*

|       |       |       |
|-------|-------|-------|
| $w_1$ | $w_2$ | $w_3$ |
| $w_4$ | $w_5$ | $w_6$ |
| $w_7$ | $w_8$ | $w_9$ |

3x3

=

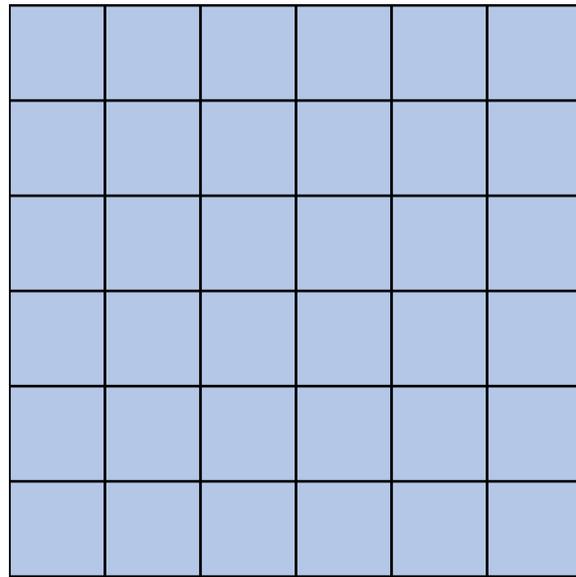
|  |  |  |  |
|--|--|--|--|
|  |  |  |  |
|  |  |  |  |
|  |  |  |  |
|  |  |  |  |

4x4

Convolution

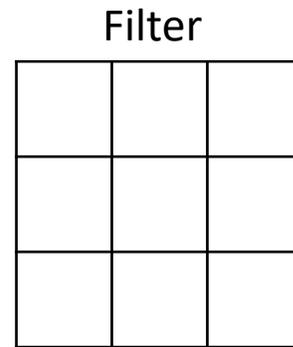
# Padding

Remember the original example:



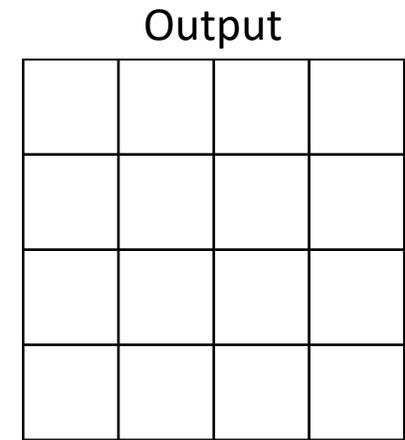
6x6

\*



3x3

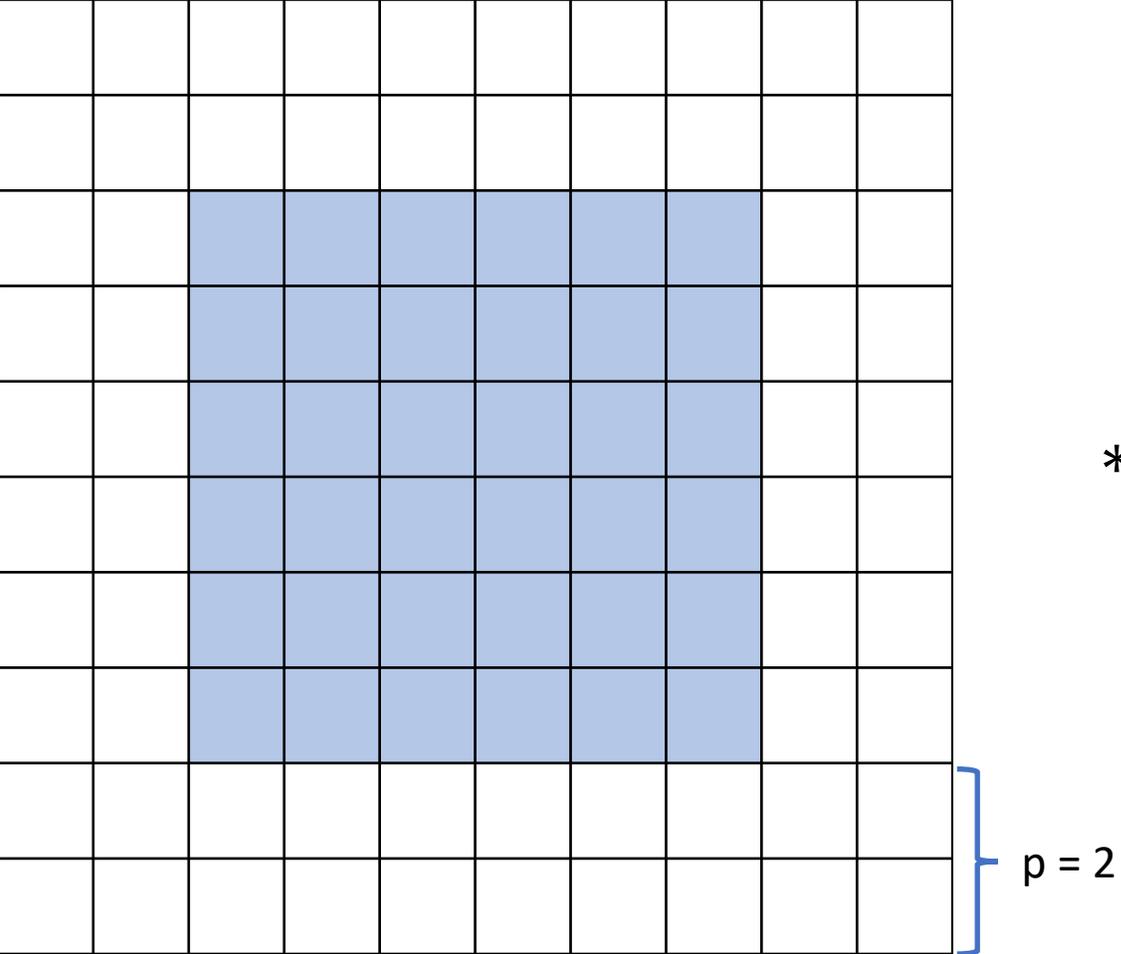
=



4x4

# Padding

Original image:  $6 \times 6 = n \times n$

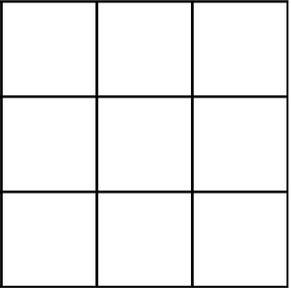


Padded image:  $10 \times 10$

$p = 2$   
 $p = \text{padding}$

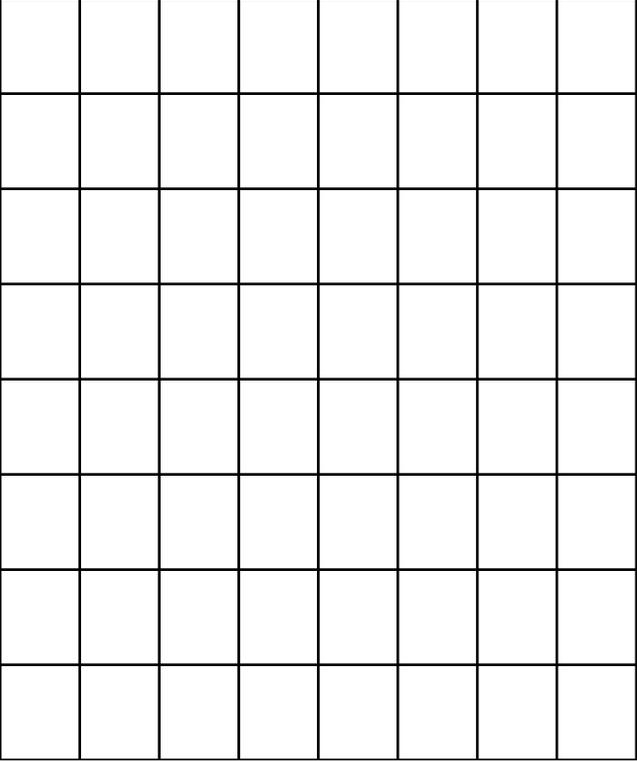
\*

Filter



$3 \times 3$   
 $f \times f$

=



Output:  $8 \times 8$

Final output dims =  $(n-f+2p+1) \times (n-f+2p+1)$

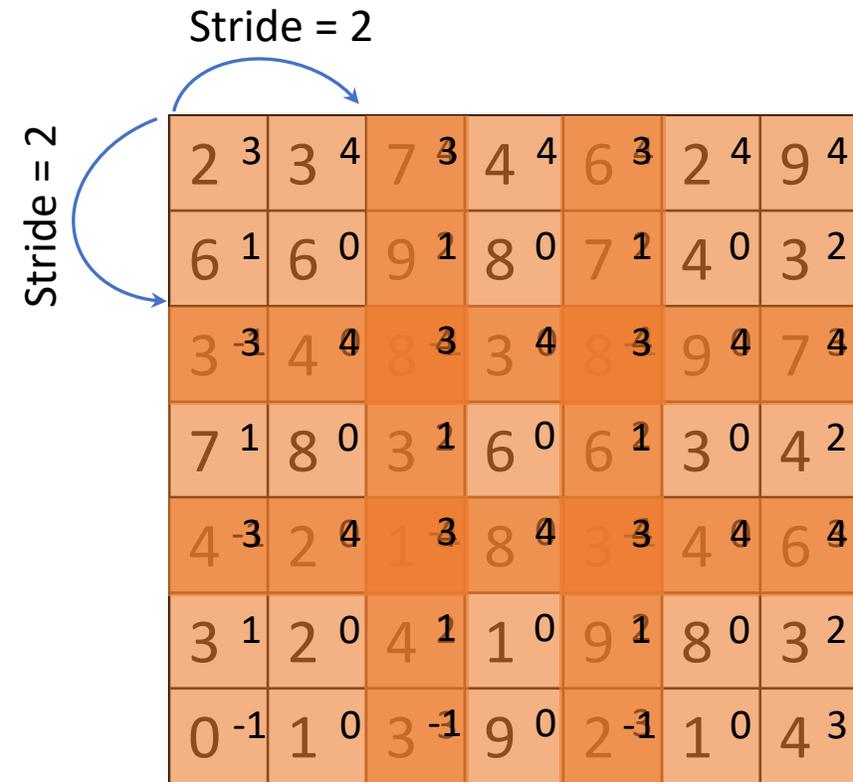
# “Valid” & “Same” Convolutions

“Valid”: if no padding!

“Same”: Pad such that the output size matches the input size.

$$P = (f-1)/2$$

# Strided convolution



\*

|    |   |   |
|----|---|---|
| 3  | 4 | 4 |
| 1  | 0 | 2 |
| -1 | 0 | 3 |

=

|    |     |     |
|----|-----|-----|
| 91 | 100 | 83  |
| 69 | 91  | 127 |
| 44 | 72  | 74  |

# Output Size (Output “Shape”)

Image dims:  $n \times n$

Filter dims:  $f \times f$

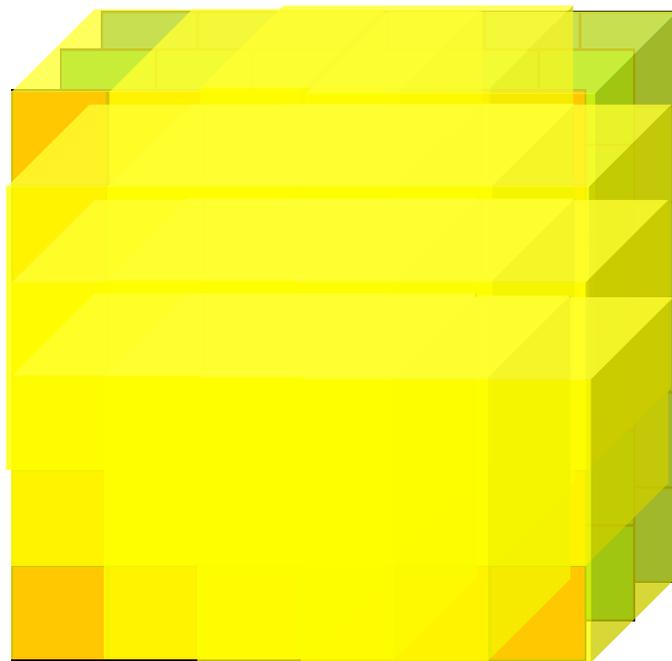
Padding:  $p$

Stride:  $s$

Floor (round to the lowest integer)

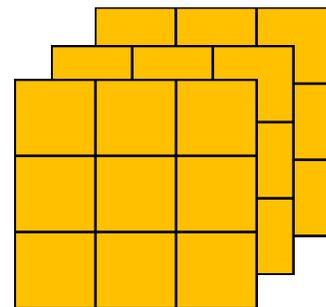
$$\text{Output dims: } \left\lfloor \frac{n+2p-f}{s} + 1 \right\rfloor \times \left\lfloor \frac{n+2p-f}{s} + 1 \right\rfloor$$

# Convolutions on RGB image



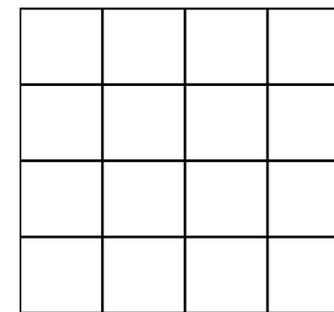
6x6x3

\*

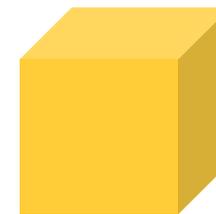


3x3x3

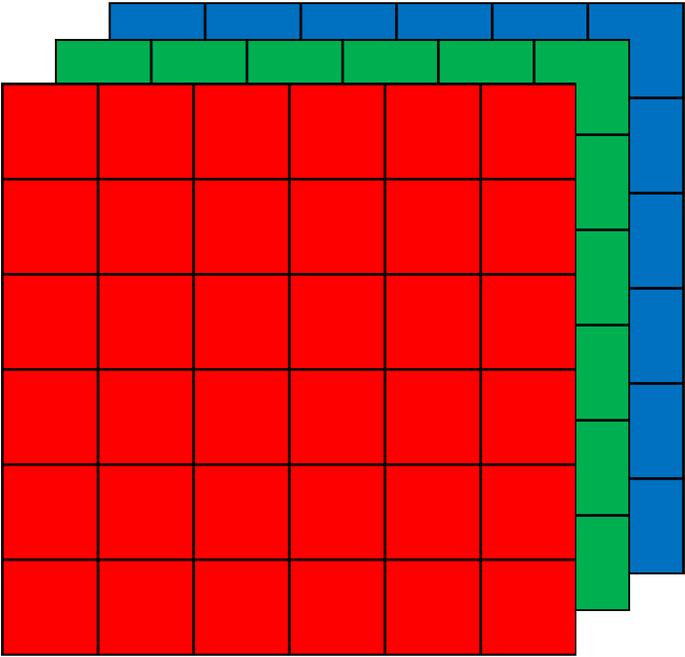
=



4 x 4



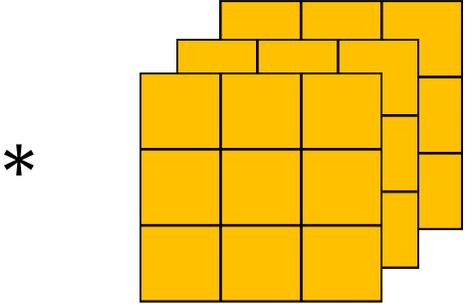
# Multiple filters



6 x 6 x 3

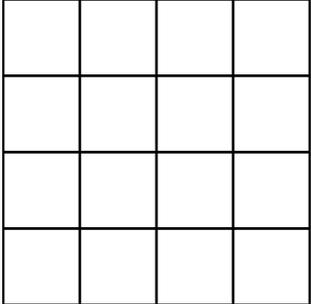
Number of channels

$n \times n \times n_c$

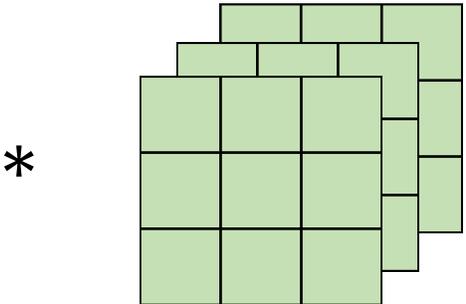


3 x 3 x 3

=

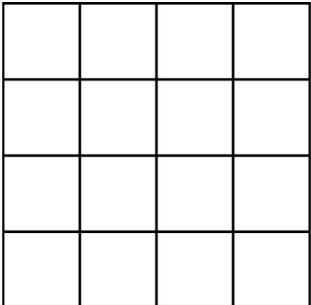


4 x 4



3 x 3 x 3

=



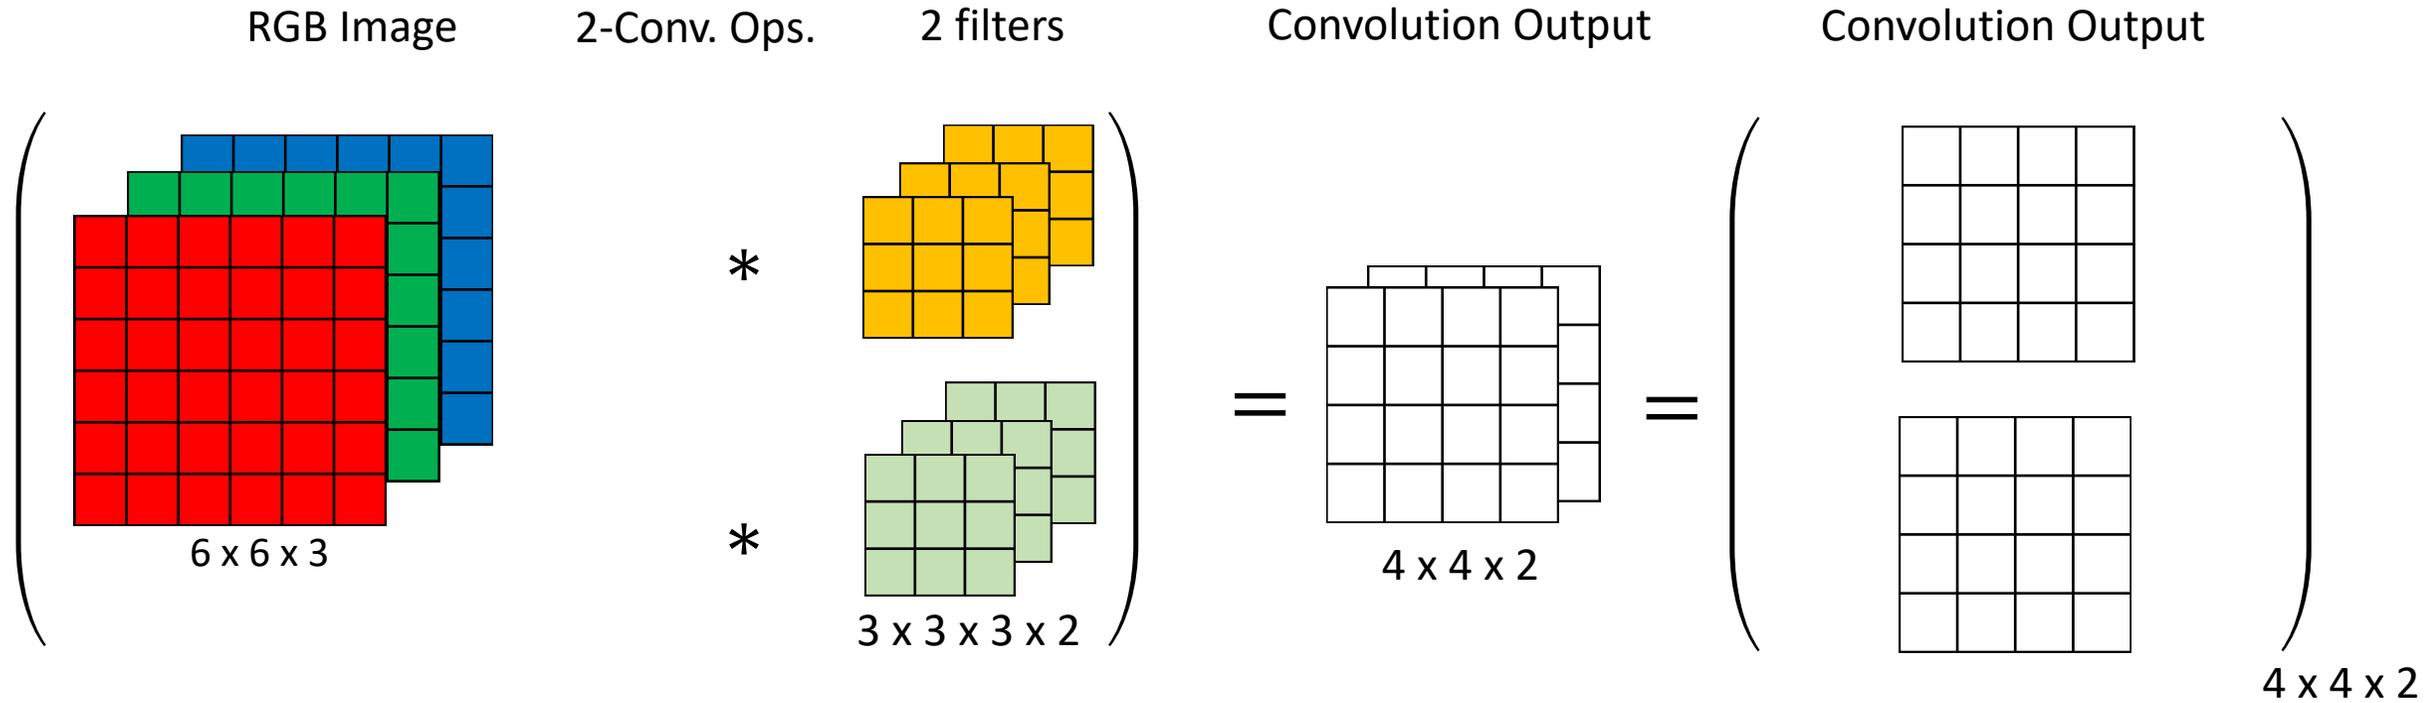
4 x 4

Number of filters

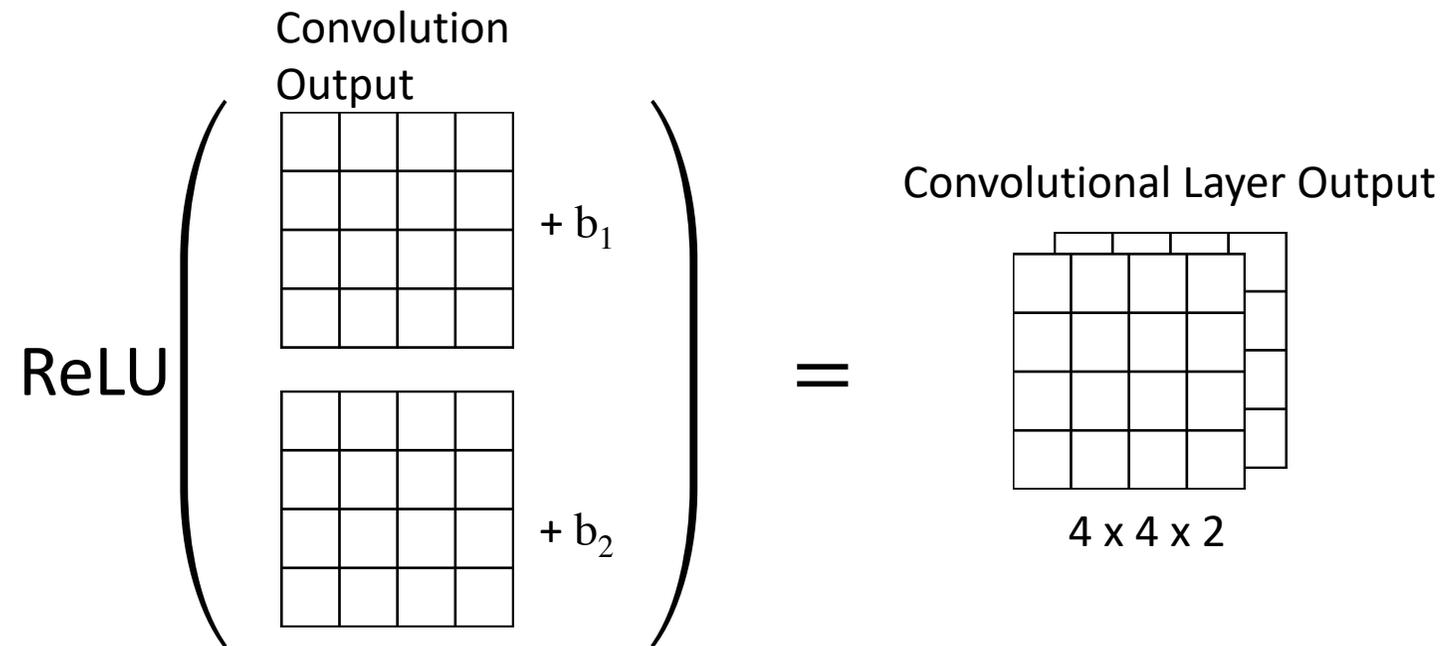
Number of filters

\*  $f \times f \times n_c \times 2 = (n - f + 1) \times (n - f + 1) \times 2$

# Multiple filters



# A Convolutional Layer of CNN (ConvNet)



Convolutional Layer Output =  $\text{ReLU} ( (RGBImage * filter) + b )$  } Similar notation!  
Output of a neuron in a FC NN:  $a = \text{ReLU} ( \mathbf{w}^T \mathbf{x} + b )$  in FC NN } (but not equal)

# Input – Output Dimensions for $l^{th}$ Convolutional Layer

Input (One image):  $n_H^{[l-1]} \times n_W^{[l-1]} \times n_C^{[l-1]}$   $\longrightarrow$  Output:  $n_H^{[l]} \times n_W^{[l]} \times n_C^{[l]}$

Padding size:  $p^{[l]} \times p^{[l]}$

Stride size:  $s^{[l]} \times s^{[l]}$

Filter size:  $f^{[l]} \times f^{[l]} \times n_C^{[l-1]} \times n_C^{[l]} = \text{One\_Filter}_{Dims} \times n_C^{[l]}$

(Number of Weights = Filter size), Bias size:  $n_C^{[l]}$

$$\text{Output dims (for one image)} : \underbrace{\left[ \frac{n_H^{[l-1]} + 2p^{[l]} - f^{[l]}}{s^{[l]}} + 1 \right]}_{n_H^{[l]}} \times \underbrace{\left[ \frac{n_W^{[l-1]} + 2p^{[l]} - f^{[l]}}{s^{[l]}} + 1 \right]}_{n_W^{[l]}} \times n_C^{[l]}$$

The output dimensions for  $m$  input samples  $A^{[l]}$ :  $m \times n_H^{[l]} \times n_W^{[l]} \times n_C^{[l]}$

This lecture contains material from Andrew Ng and Ulas Bagci