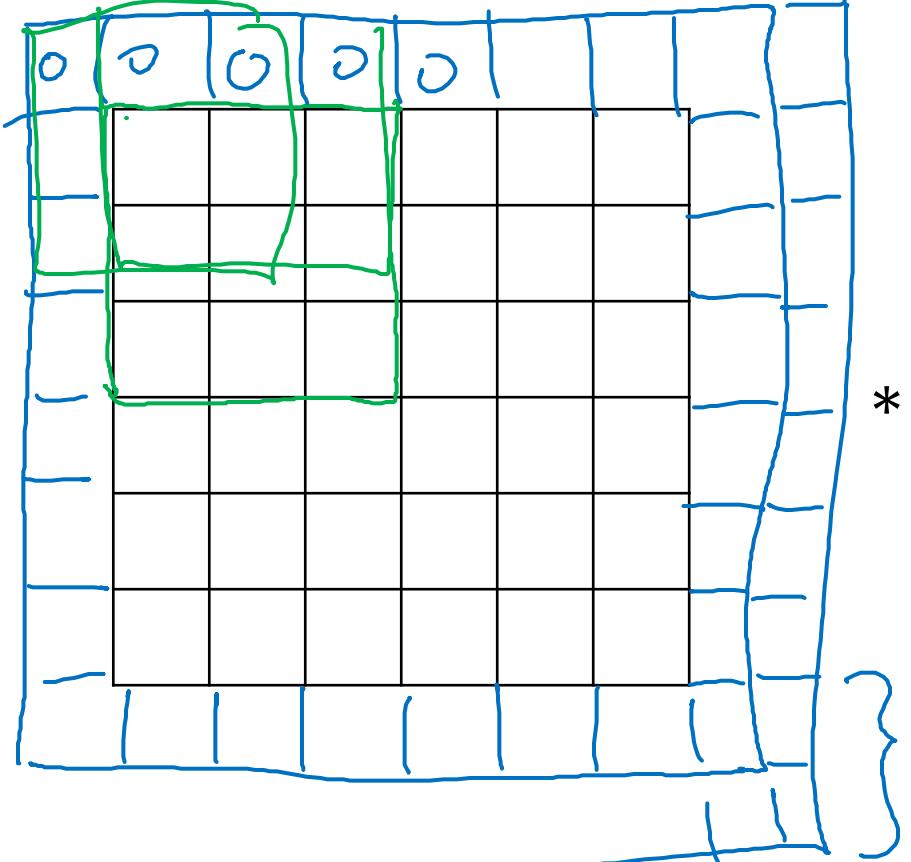


Convolutional Neural Networks

Padding

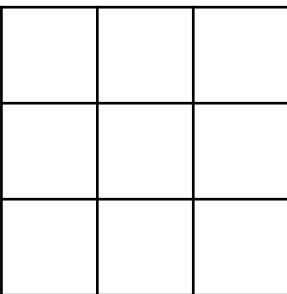
Padding



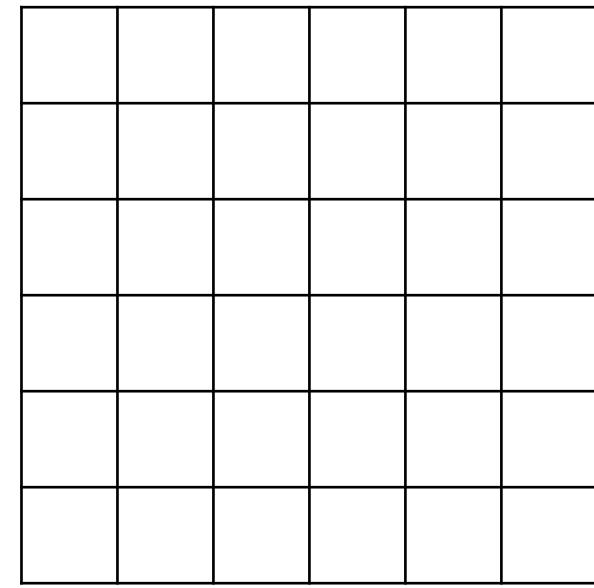
$$\frac{6 \times 6}{n \times n} \rightarrow 8 \times 8$$

$$n-f+1 \times n-f+1$$
$$6-3+1=4$$

$$P = \text{padding} = 1$$



=



$$\underline{\underline{6 \times 6}}$$

$$\xrightarrow{\quad} \underline{\underline{4 \times 4}}$$

$$n+2p-f+1 \times n+2p-f+1$$
$$6+2-3+1 \times \underline{\underline{}} = 6 \times 6$$

Valid and Same convolutions

→ $n \rightarrow$ padding

“Valid”: $n \times n \quad * \quad f \times f \quad \rightarrow \frac{n-f+1}{f} \times \frac{n-f+1}{f}$

$$\begin{array}{ccc} 6 \times 6 & * & 3 \times 3 \\ \rightarrow & & 4 \times 4 \end{array}$$

“Same”: Pad so that output size is the same as the input size.

$$n + 2p - f + 1 = n \Rightarrow p = \frac{f-1}{2}$$

$$3 \times 3 \quad p = \frac{3-1}{2} = 1 \quad \left| \begin{array}{c} 5 \times 5 \\ f=5 \end{array} \right. \quad p=2$$

$f \in$ usually odd
 1×1
 3×3
 5×5
 7×7

Convolutional Neural Networks

Strided
convolutions

Strided convolution

2	3	7	4	6	2	9
6	6	9	8	7	4	3
3	4	8	3	8	9	7
7	8	3	6	6	3	4
4	2	1	8	3	4	6
3	2	4	1	9	8	3
0	1	3	9	2	1	4

$n \times n$ * $f \times f$
padding p stride s
 $s=2$

$$\begin{matrix} & * & \\ & & \end{matrix} = \begin{matrix} 3 \times 3 \\ \text{stride } 2 \\ \lfloor z \rfloor = \text{floor}(z) \end{matrix} = \begin{matrix} 3 \times 3 \end{matrix}$$

$$\left\lfloor \frac{n+2p-f}{s} + 1 \right\rfloor \times \left\lfloor \frac{n+2p-f}{s} + 1 \right\rfloor$$
$$\frac{7+0-3}{2} + 1 = \frac{4}{2} + 1 = 3$$

Summary of convolutions

$n \times n$ image $f \times f$ filter

padding p stride s

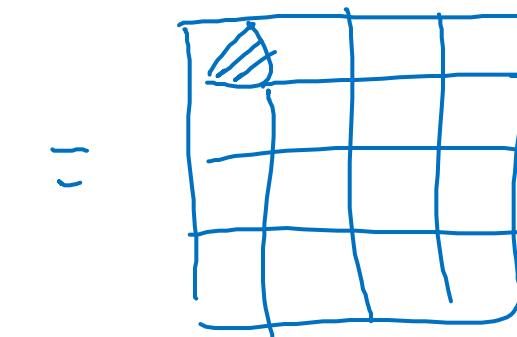
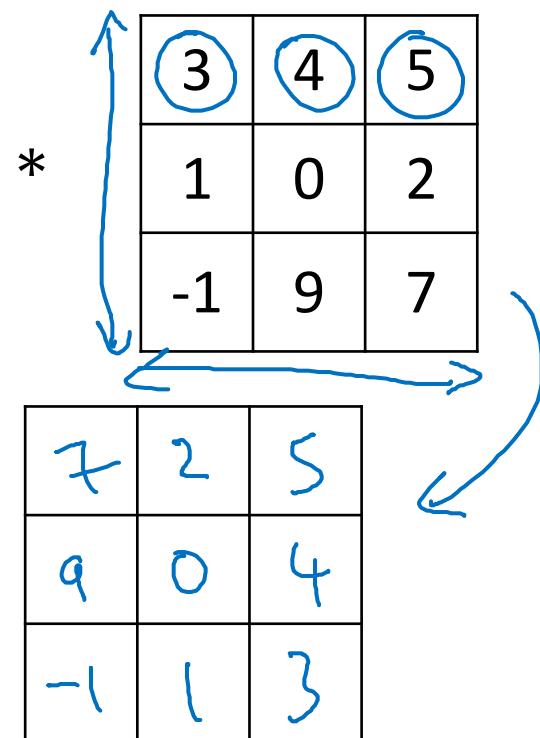
Output size:

$$\left\lfloor \frac{n+2p-f}{s} + 1 \right\rfloor \quad \times \quad \left\lfloor \frac{n+2p-f}{s} + 1 \right\rfloor$$

Technical note on cross-correlation vs. convolution

Convolution in math textbook:

2	3	7	5	4	6	2
6	6	9	4	8	7	4
3	4	8	3	3	8	9
7	8	3	6	6	6	3
4	2	1	8	3	3	4
3	2	4	1	9	8	

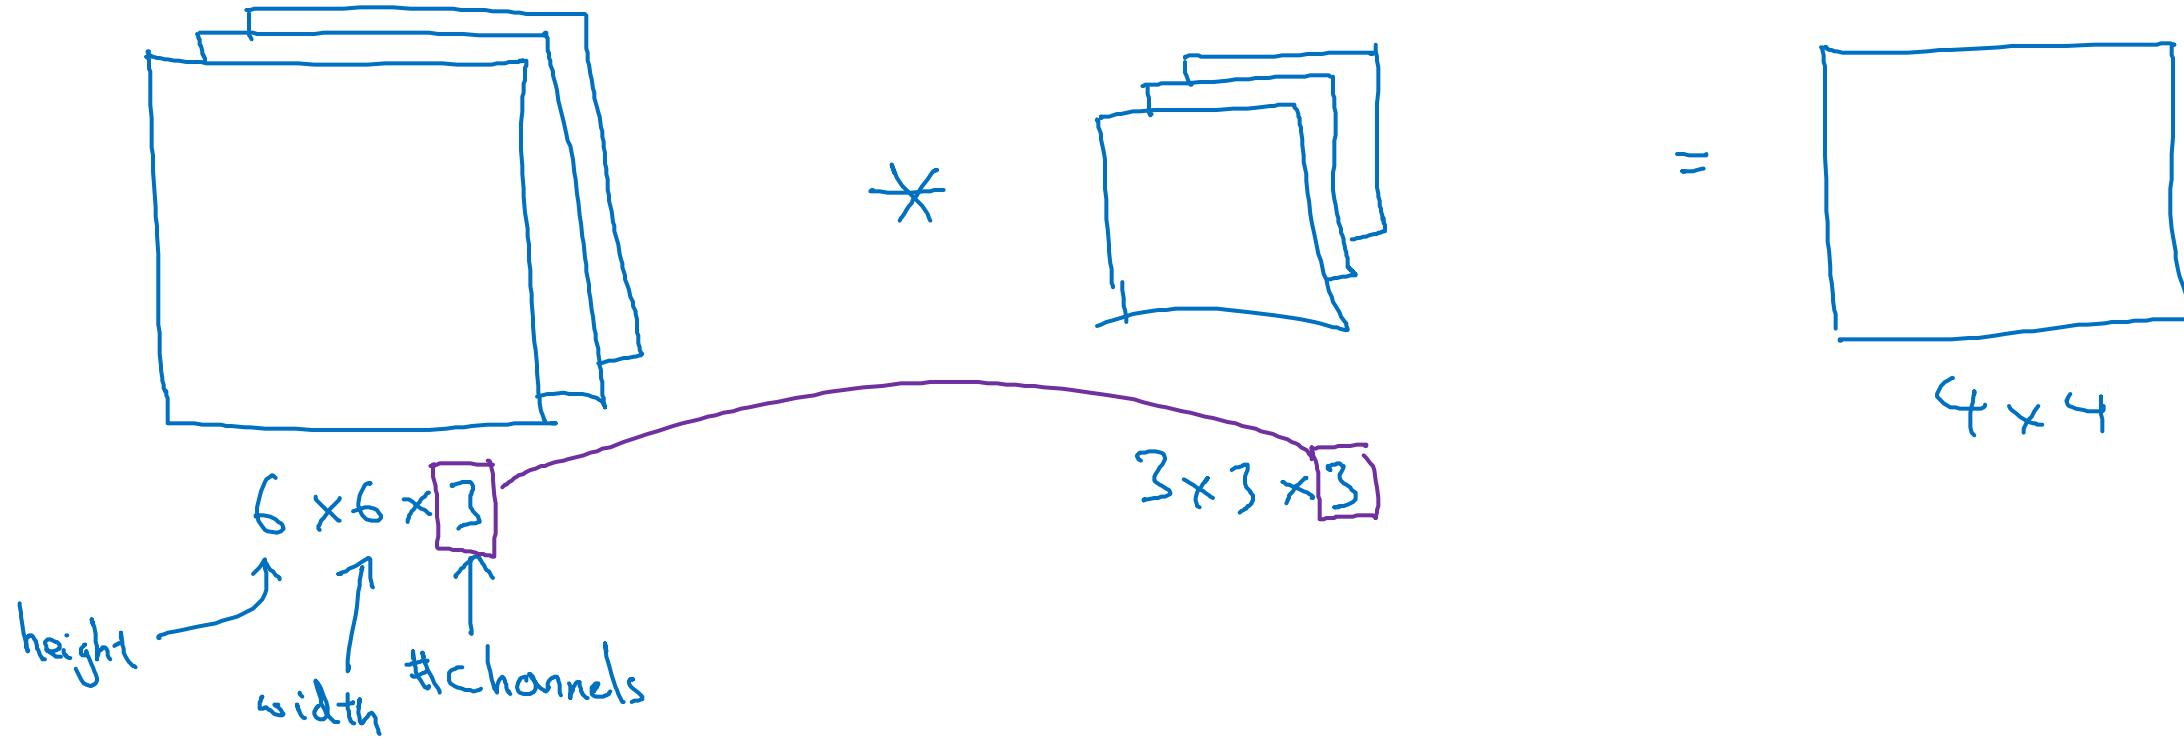


$$(A * B) * C = A * (B * C)$$

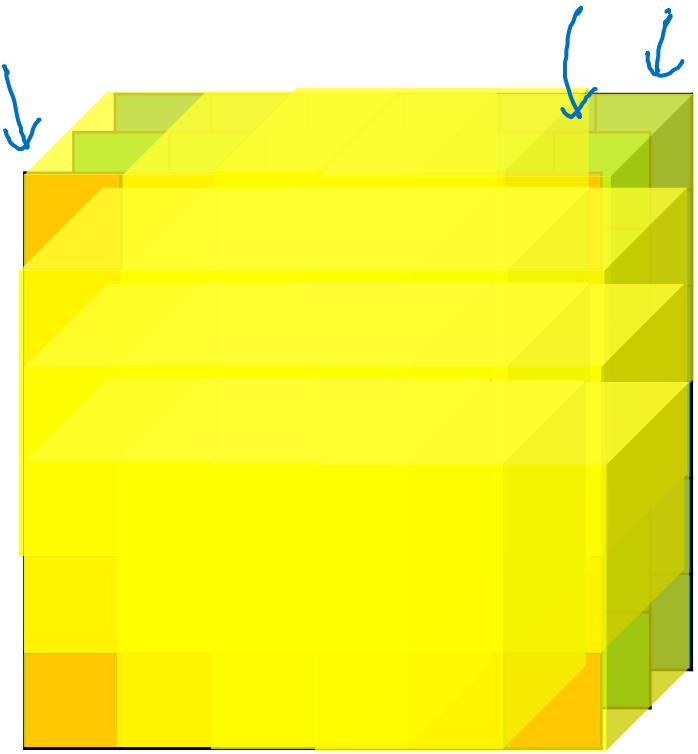
Convolutional Neural Networks

Convolutions over
volumes

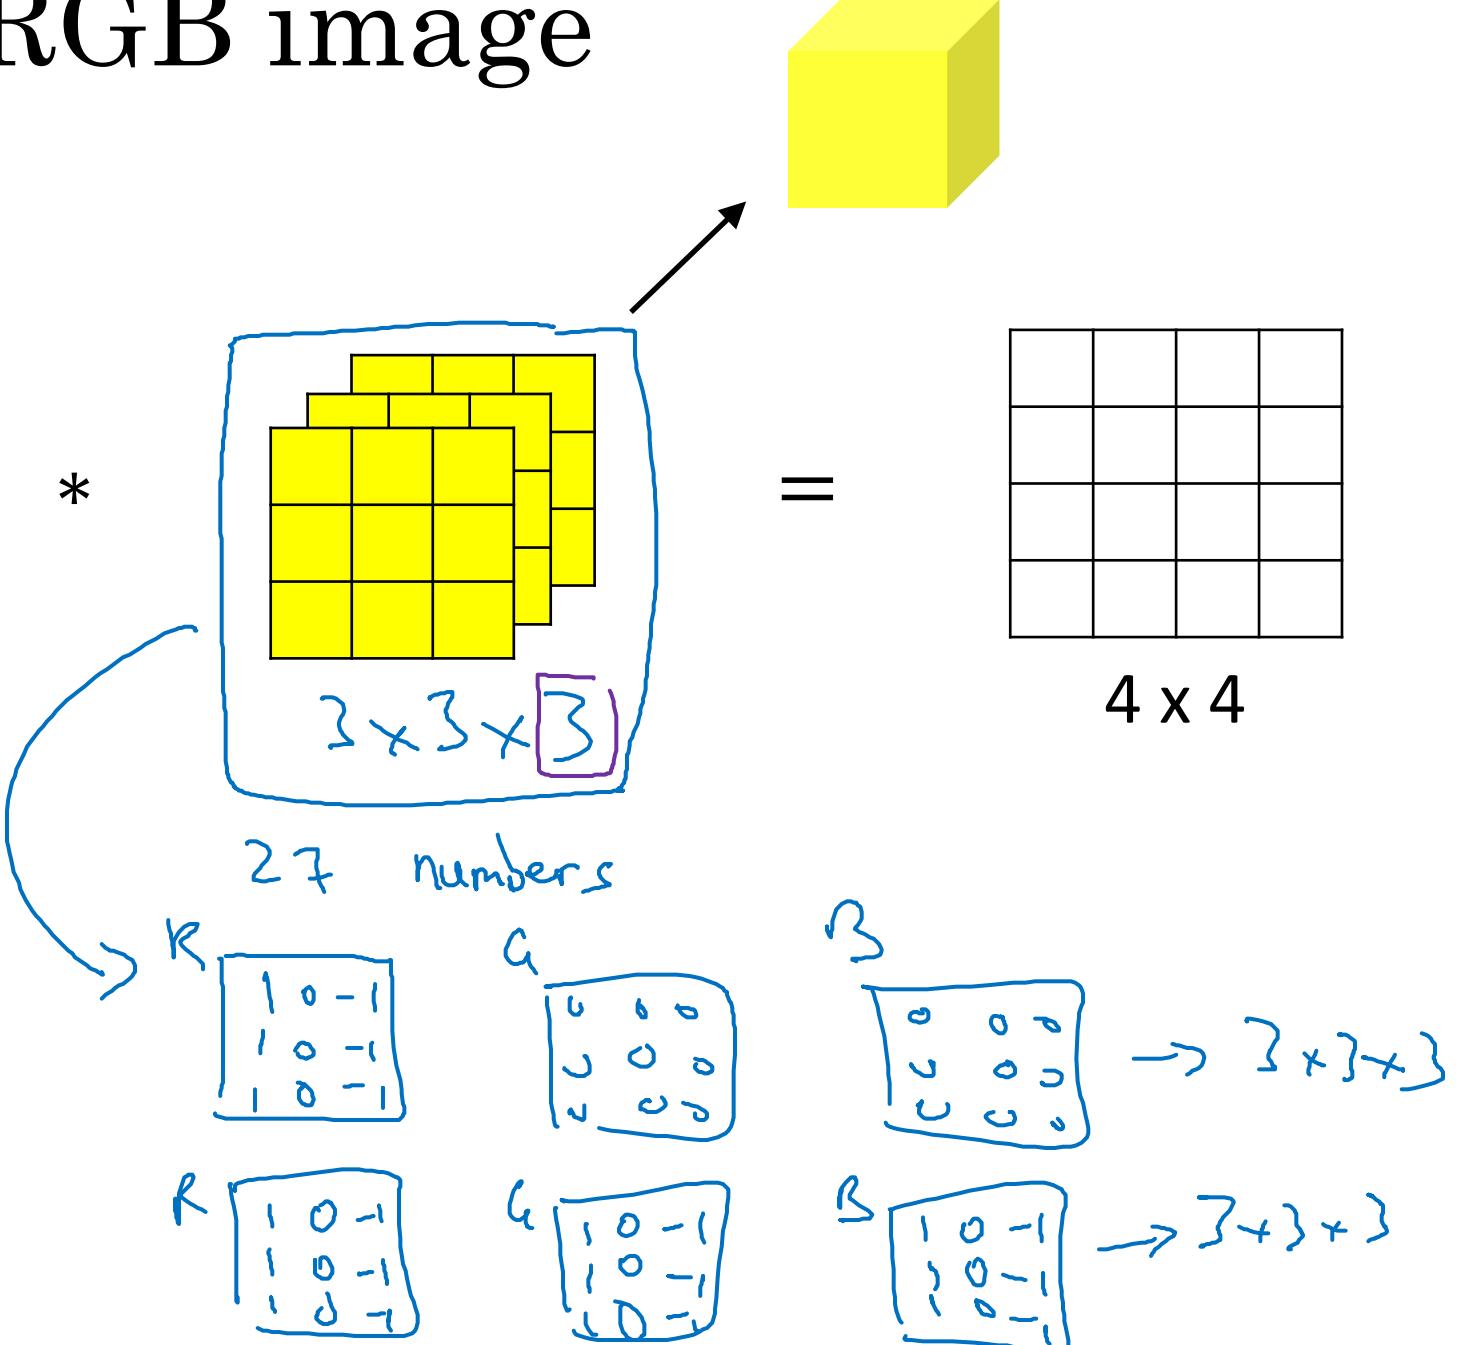
Convolutions on RGB images



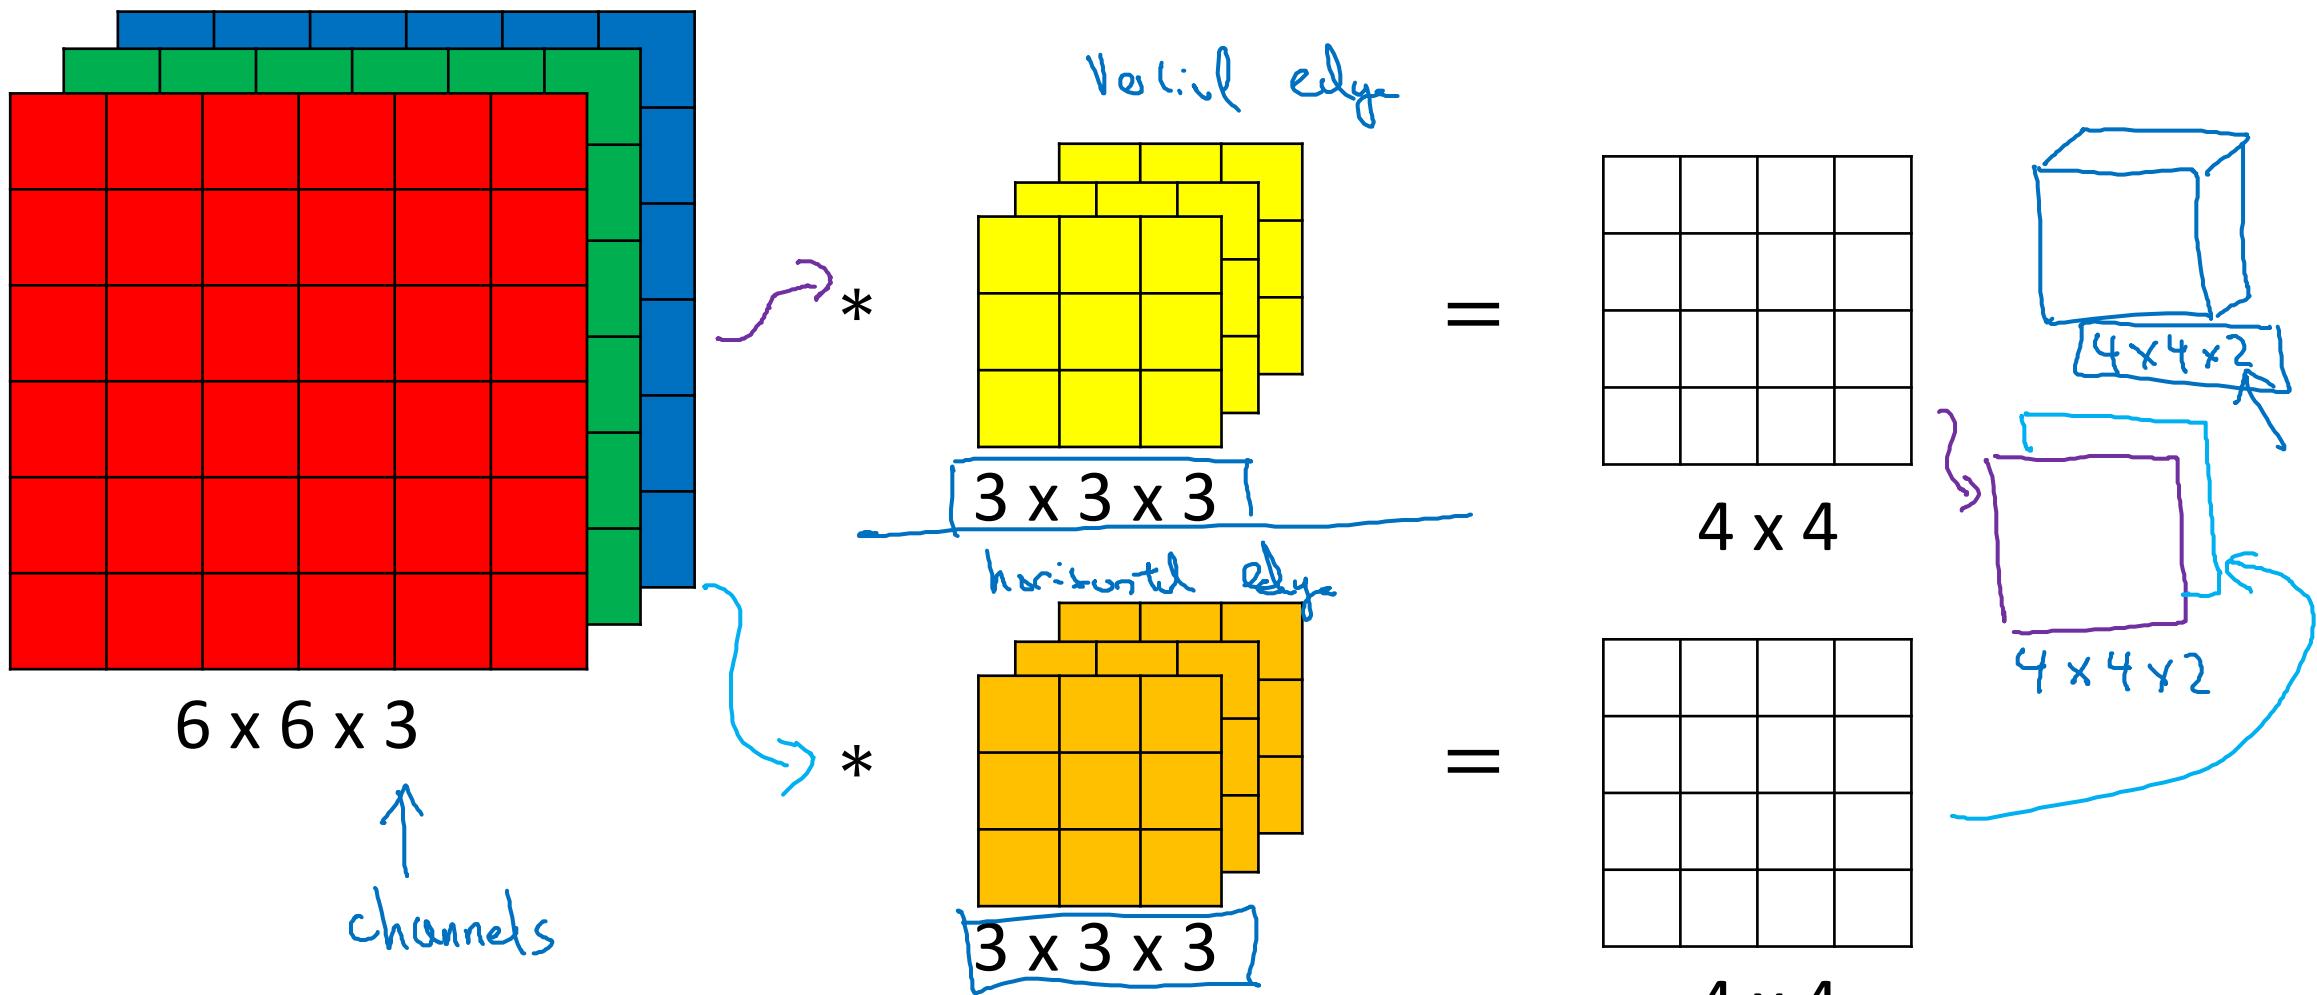
Convolutions on RGB image



$$\begin{matrix} & 6 \times 6 \times 3 \\ \uparrow & \uparrow & \uparrow \\ \boxed{\text{ }} \quad * \quad \boxed{\text{ }} & = & \boxed{\text{ }} \end{matrix}$$



Multiple filters



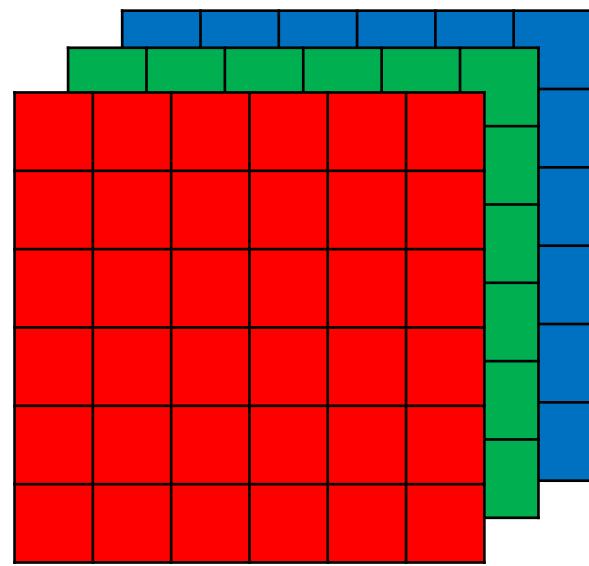
Summary: $n \times n \times n_c \times f \times f \times n_c \rightarrow \frac{n-f+1}{4} \times \frac{n-f+1}{4} \times \frac{n_c}{2} \# \text{filters}$

$$6 \times 6 \times 3 \quad 3 \times 3 \times 3 \quad \frac{6-3+1}{4} \times \frac{6-3+1}{4} \times \frac{3}{2} \# \text{filters}$$

Convolutional Neural Networks

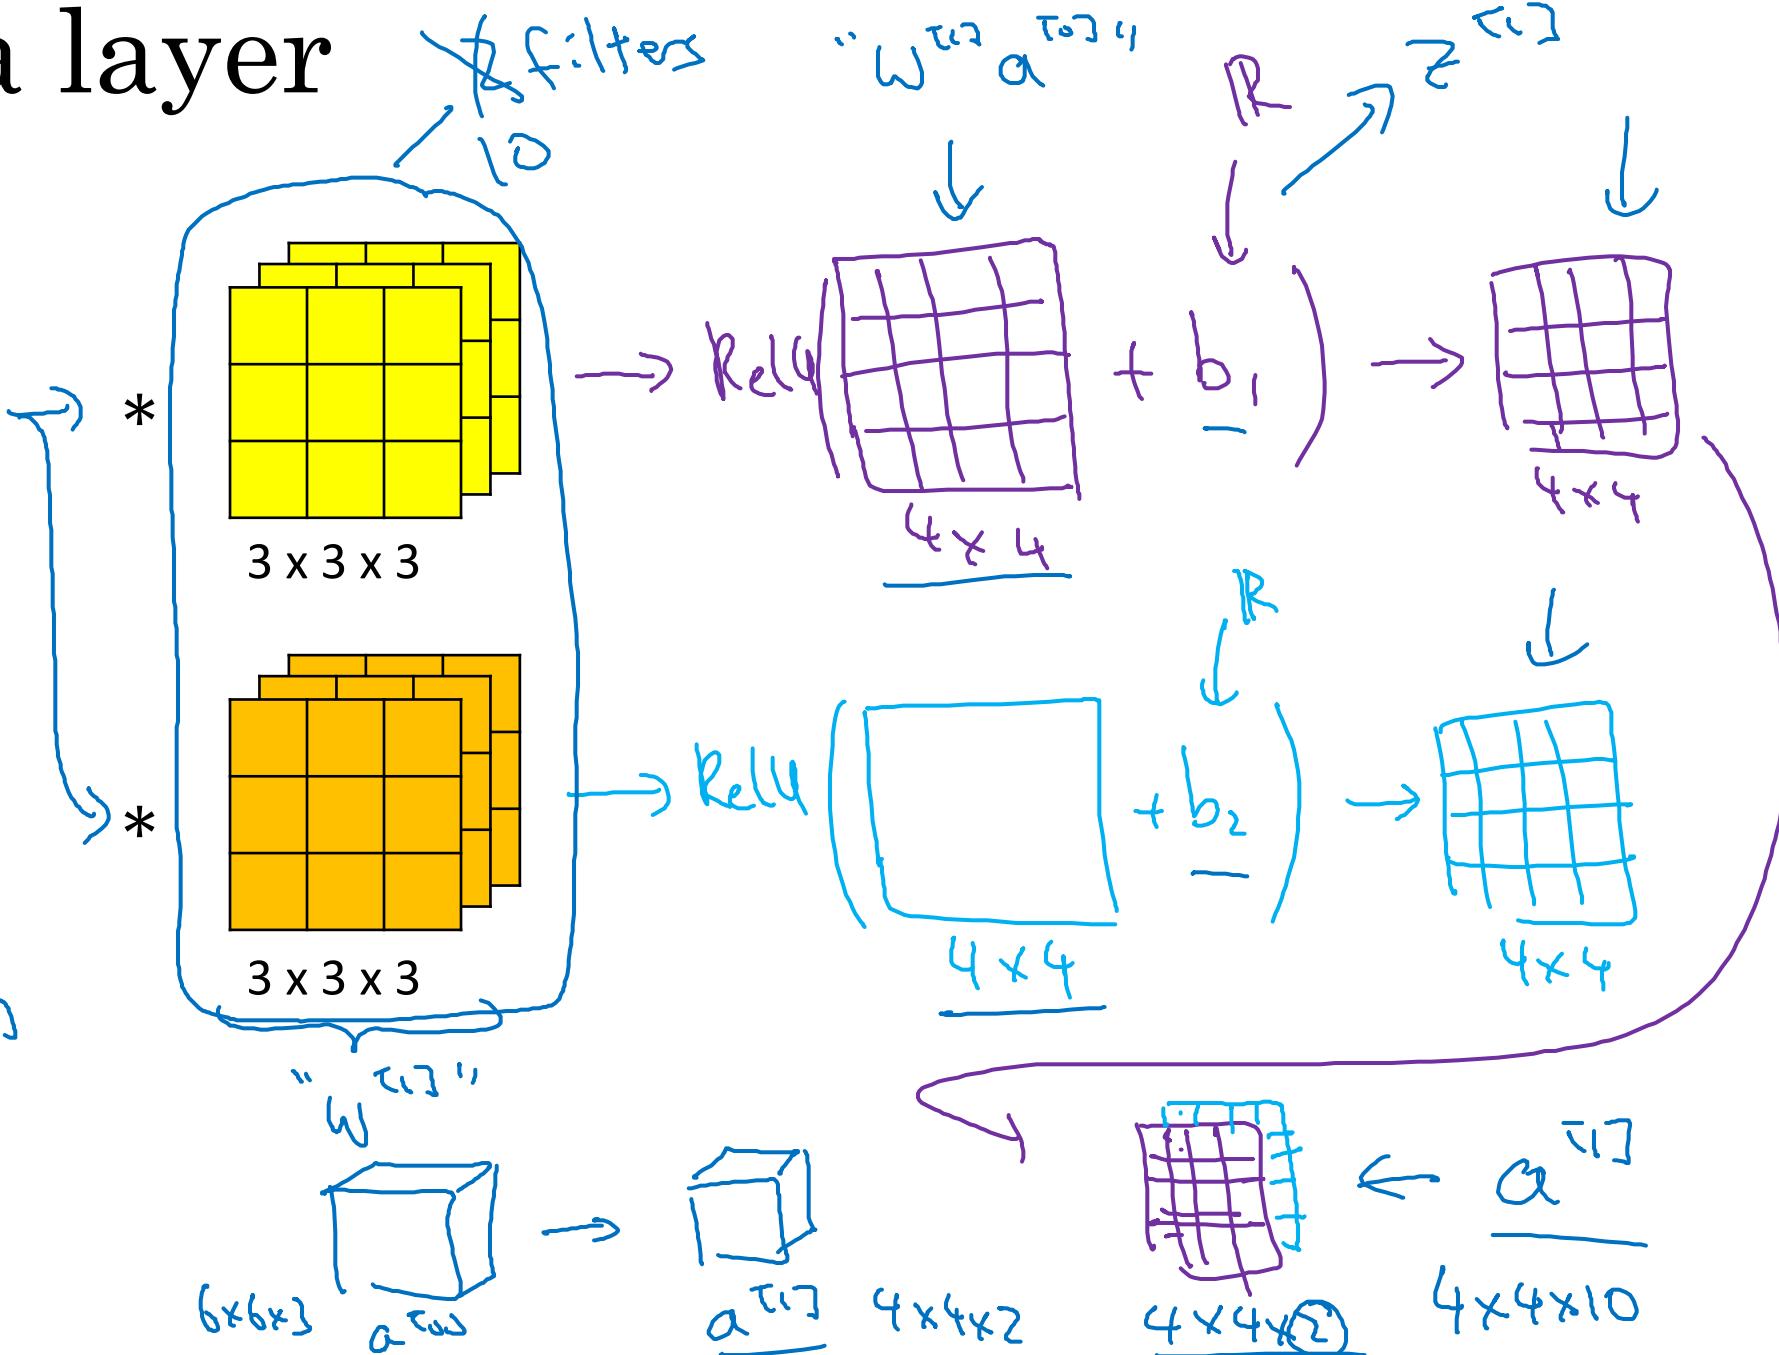
One layer of a
convolutional
network

Example of a layer



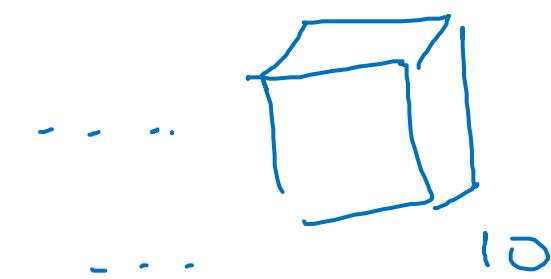
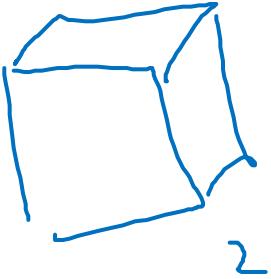
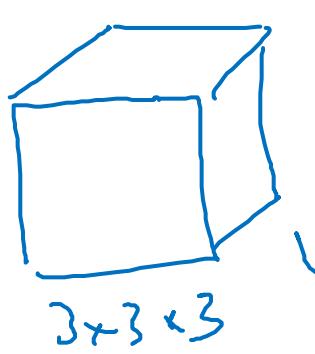
$$z^{(1)} = w^{(1)} a^{(0)} + b^{(1)}$$

$$a^{(1)} = g(z^{(1)})$$



Number of parameters in one layer

If you have 10 filters that are $3 \times 3 \times 3$ in one layer of a neural network, how many parameters does that layer have?



27 parameters.

+ bias

→ 28 parameters.

280 parameters.

Summary of notation

If layer l is a convolution layer:

$f^{[l]}$ = filter size

$p^{[l]}$ = padding

$s^{[l]}$ = stride

$n_c^{[l]}$ = number of filters

→ Each filter is: $f^{[l]} \times f^{[l]} \times n_c^{[l]}$

Activations: $a^{[l]} \rightarrow n_H^{[l]} \times n_W^{[l]} \times n_C^{[l]}$.

Weights: $f^{[l]} \times f^{[l]} \times n_c^{[l-1]} \times n_c^{[l]}$

bias: $n_c^{[l]}$

Input: $n_H^{[l-1]} \times n_W^{[l-1]} \times n_c^{[l-1]}$

Output: $n_H^{[l]} \times n_W^{[l]} \times n_c^{[l]}$

$$n_H^{[l]} = \left\lfloor \frac{n_H^{[l-1]} + 2p^{[l]} - f^{[l]}}{s^{[l]}} + 1 \right\rfloor$$

$$A^{[l]} \rightarrow m \times n_H^{[l]} \times n_W^{[l]} \times n_c^{[l]}$$

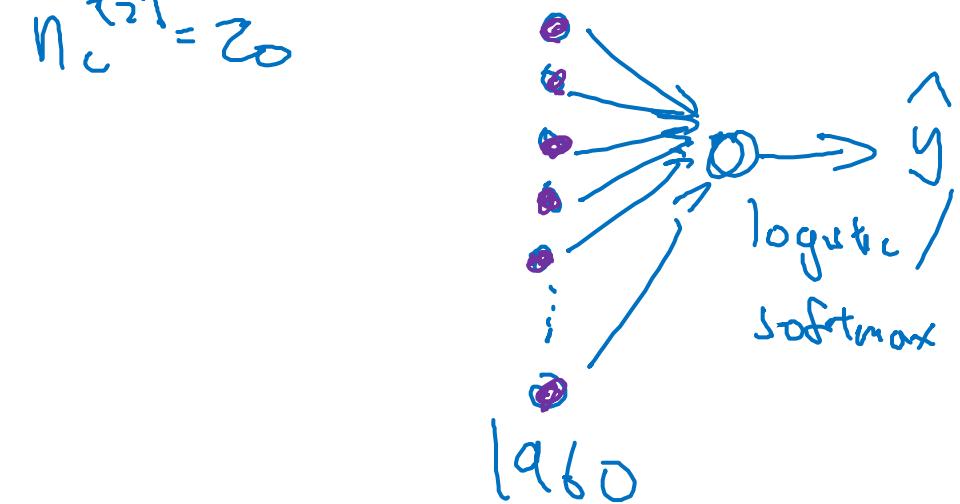
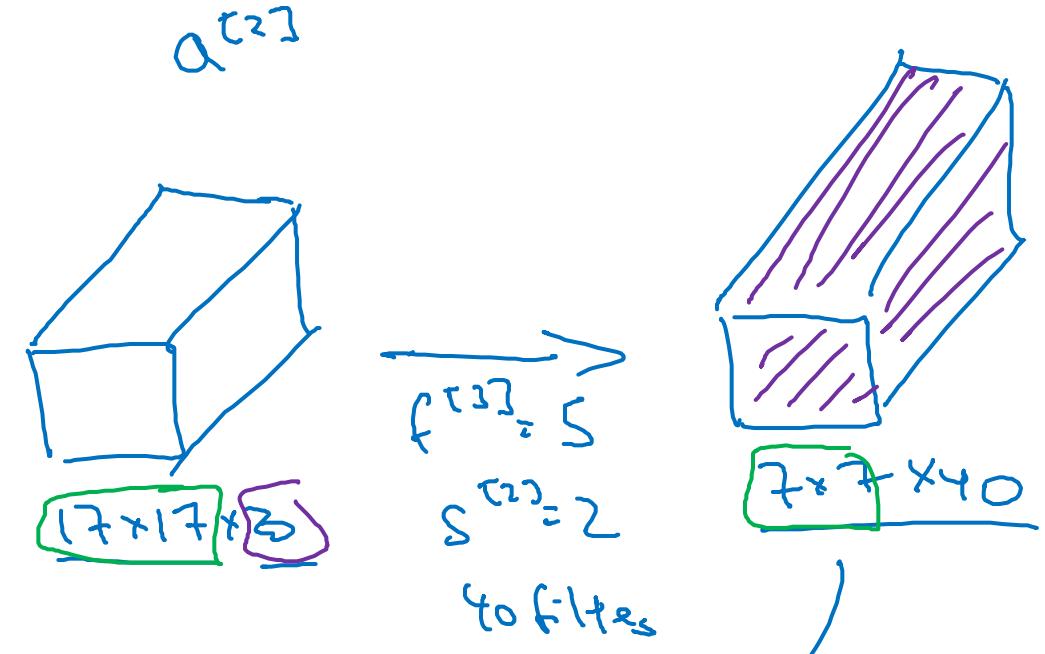
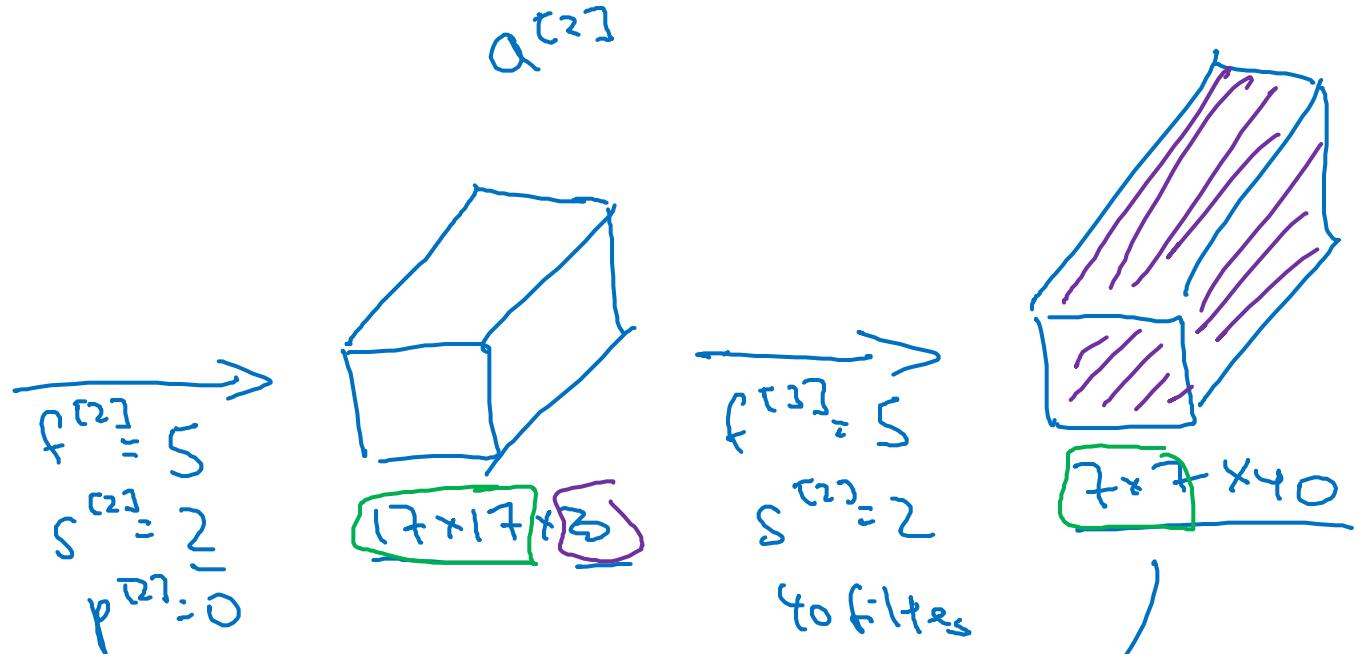
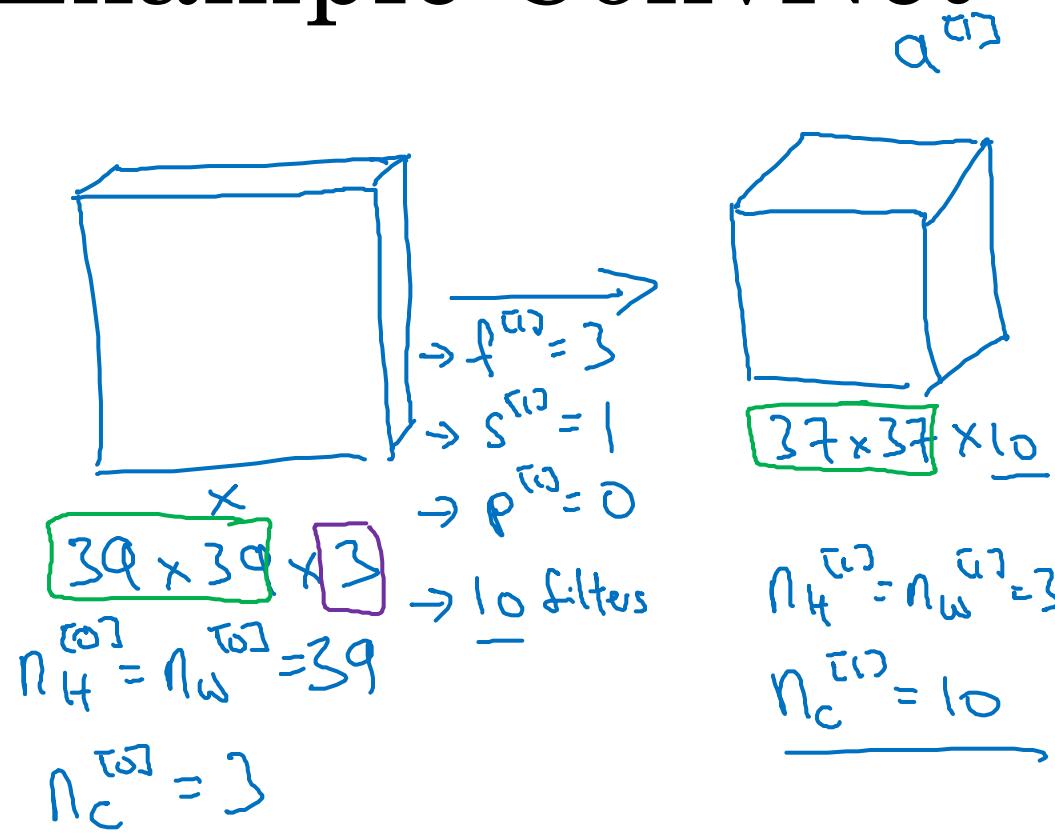
#f: #filters in layer l.

$$n_c^{[l]} \times n_H^{[l]} \times n_W^{[l]}$$

Convolutional Neural Networks

A simple convolution
network example

Example ConvNet



Types of layer in a convolutional network:

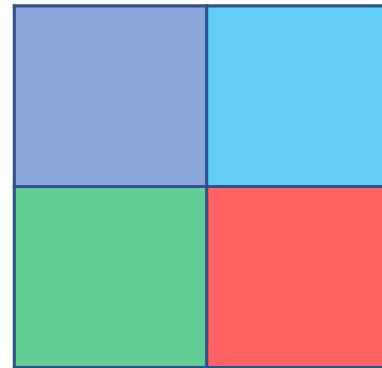
- Convolution
- Pooling
- Fully connected

Convolutional Neural Networks

Pooling layers

Pooling layer: Max pooling

1	3	2	1
2	9	1	1
1	3	2	3
5	6	1	2



Pooling layer: Max pooling

1	3	2	1	3
2	9	1	1	5
1	3	2	3	2
8	3	5	1	0
5	6	1	2	9

$5 \times 5 \times 2 n_c$



9	9	5
9	9	5
8	6	9

$$\frac{f=3}{s=1} \quad 3 \times 3 \times 2 \quad n_c$$

$$\left(\frac{n+p-f}{s} + 1 \right)$$

Pooling layer: Average pooling

1	3	2	1
2	9	1	1
1	4	2	3
5	6	1	2



3.75	1.25
4	2

$$f=2$$

$$s=2$$

$$\underbrace{7 \times 7 \times 1000}_{\longrightarrow} \rightarrow 1 \times 1 \times 1000$$

Summary of pooling

Hyperparameters:

f : filter size

s : stride

Max or average pooling

$\rightarrow p: \text{padding}.$

No parameters to learn!

$$n_H \times n_w \times \underline{n_c}$$



$$\left\lfloor \frac{n_H - f + 1}{s} \right\rfloor \times \left\lfloor \frac{n_w - f}{s} + 1 \right\rfloor$$

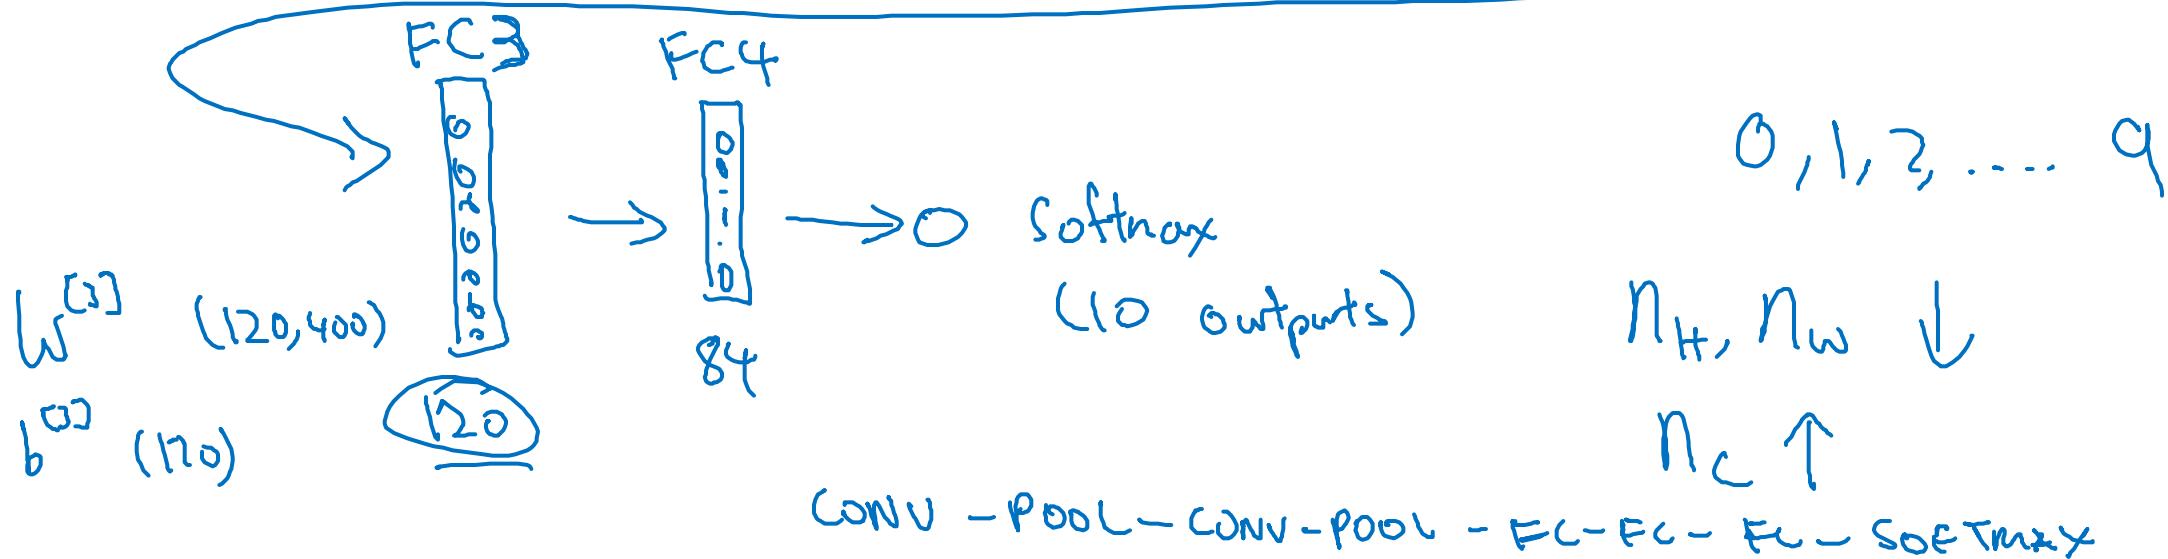
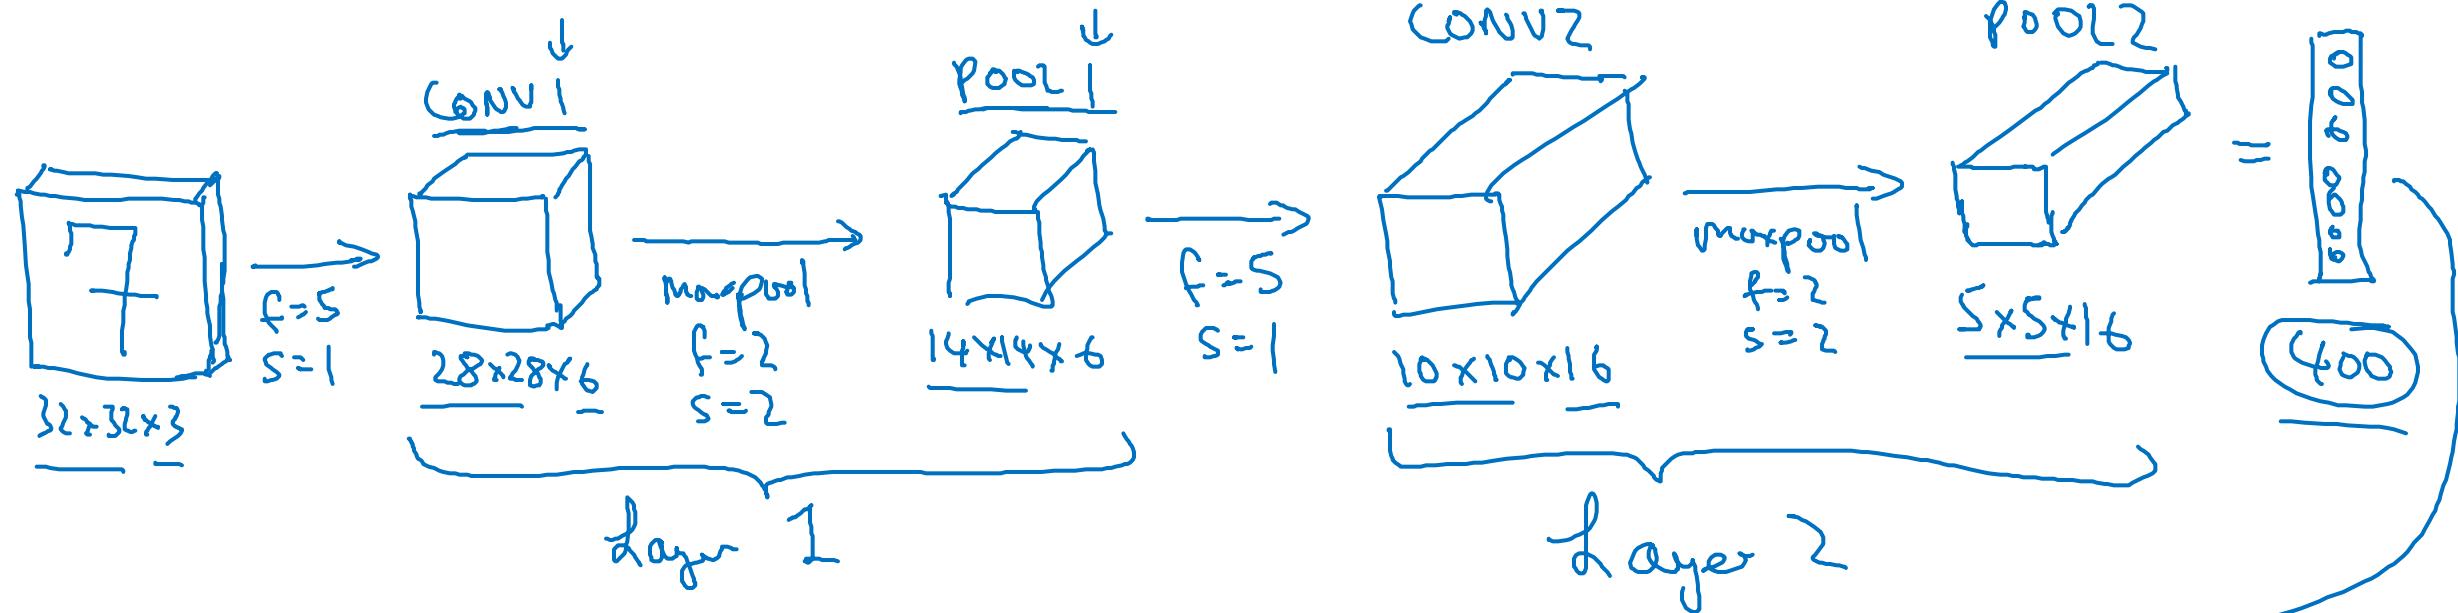
$$\times \underline{n_c}$$

Convolutional Neural Networks

Convolutional neural
network example

Neural network example

(LeNet-5)



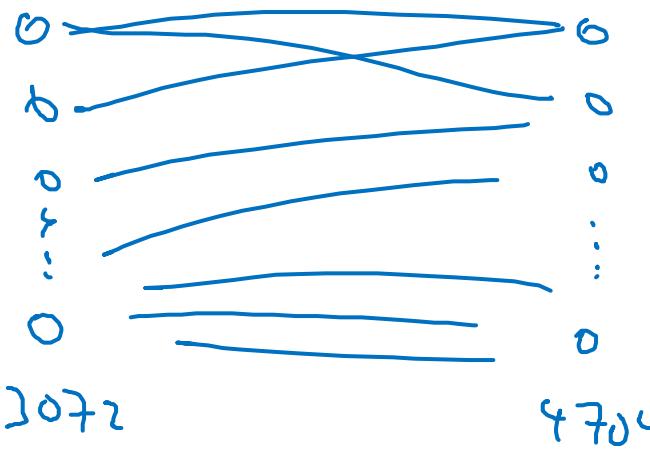
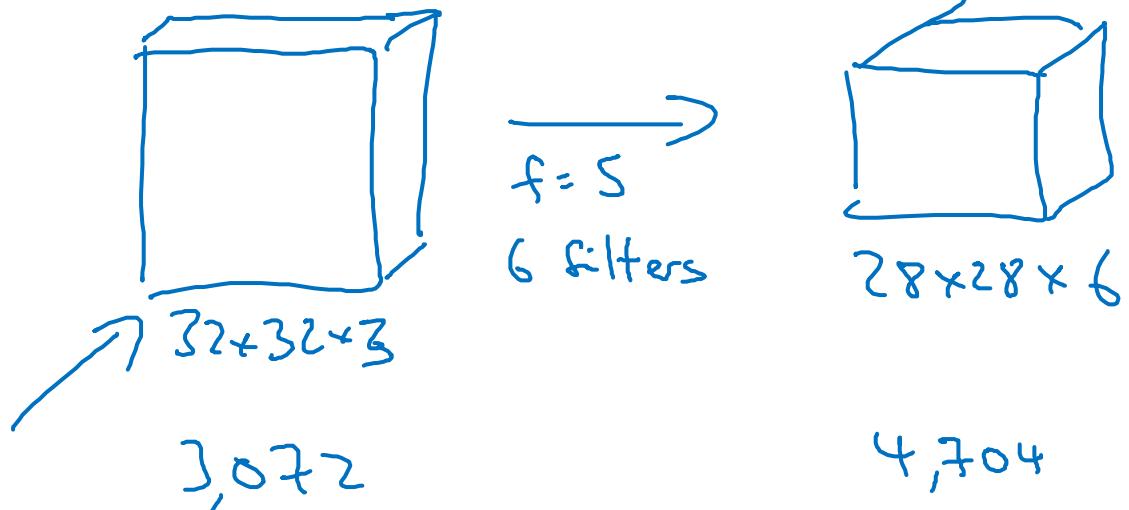
Neural network example

	Activation shape	Activation Size	# parameters
Input:	(32,32,3)	3,072 $a^{[1]}$	0
CONV1 (f=5, s=1)	(28,28,8)	6,272	456 ←
POOL1	(14,14,8)	1,568	0 ←
CONV2 (f=5, s=1)	(10,10,16)	1,600	1516 ←
POOL2	(5,5,16)	400	0 ←
FC3	(120,1)	120	48120 {
FC4	(84,1)	84	10164 }
Softmax	(10,1)	10	850

Convolutional Neural Networks

Why convolutions?

Why convolutions



$$\begin{aligned} & 5 \times 5 - 25 \\ & 26 \\ & 6 \times 26 = 156 \text{ Parameters} \\ & 3,072 \times 4,704 \approx 14M \end{aligned}$$

Why convolutions

$$\begin{array}{|c|c|c|c|c|c|} \hline 10 & 10 & 10 & 0 & 0 & 0 \\ \hline 10 & 10 & 10 & 0 & 0 & 0 \\ \hline 10 & 10 & 10 & 0 & 0 & 0 \\ \hline 10 & 10 & 10 & 0 & 0 & 0 \\ \hline 10 & 10 & 10 & 0 & 0 & 0 \\ \hline 10 & 10 & 10 & 0 & 0 & 0 \\ \hline \end{array} * \begin{array}{|c|c|c|} \hline 1 & 0 & -1 \\ \hline 1 & 0 & -1 \\ \hline 1 & 0 & -1 \\ \hline \end{array} = \begin{array}{|c|c|c|c|} \hline 0 & 30 & 30 & 0 \\ \hline 0 & 30 & 30 & 0 \\ \hline 0 & 30 & 30 & 0 \\ \hline 0 & 30 & 30 & 0 \\ \hline \end{array}$$

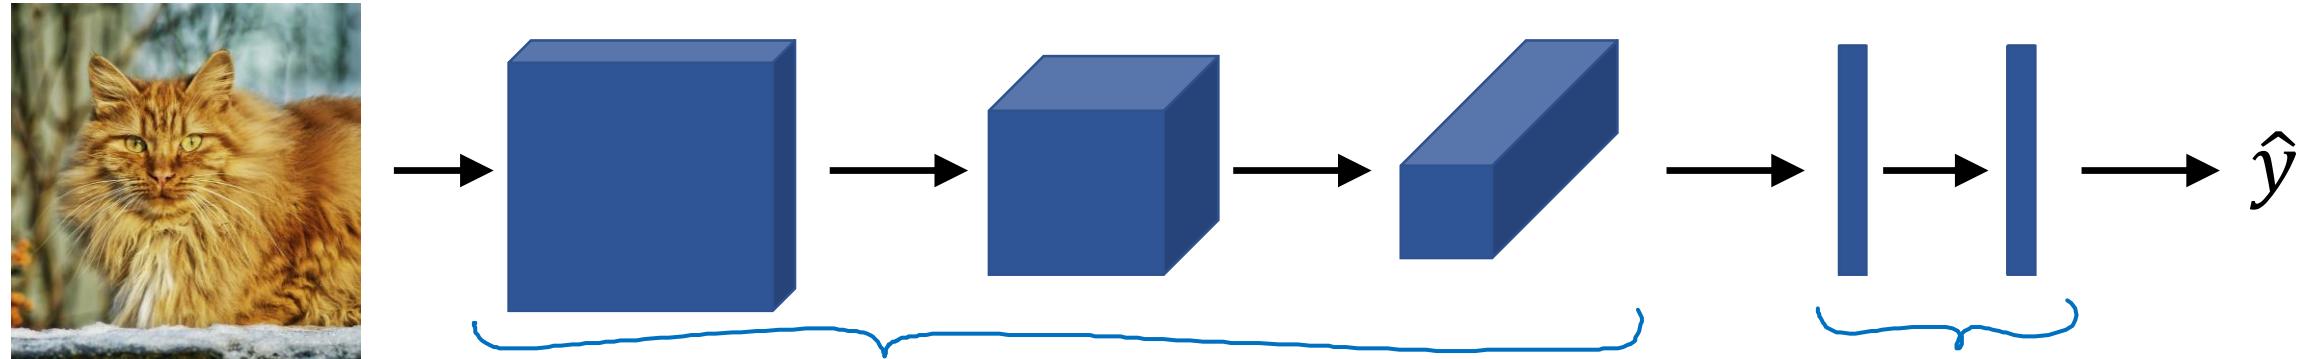
The diagram illustrates a convolution operation. On the left is a 6x6 input matrix with values 10 in the first five rows and 0 in the last row. In the center is a 3x3 kernel matrix with values 1, 0, -1 in each row. Below the kernel is a blue bracket labeled "3x3". An asterisk (*) indicates the multiplication operation, followed by an equals sign (=). On the right is the resulting 4x4 output matrix, which has values 30 in the first three rows and 0 in the last row.

Parameter sharing: A feature detector (such as a vertical edge detector) that's useful in one part of the image is probably useful in another part of the image.

Sparsity of connections: In each layer, each output value depends only on a small number of inputs.

Putting it together

Training set $(x^{(1)}, y^{(1)}) \dots (x^{(m)}, y^{(m)})$.



$$\text{Cost } J = \frac{1}{m} \sum_{i=1}^m \mathcal{L}(\hat{y}^{(i)}, y^{(i)})$$

Use gradient descent to optimize parameters to reduce J