

Convolutional Operation



Vertical Edge Detection

3	0	1	2	7	4
1	5	8	9	3	1
2	7	2	5	1	3
0	1	3	1	7	8
4	2	1	6	2	8
2	4	5	2	3	9

The diagram illustrates a convolution operation. On the left, a 3x3 kernel (highlighted in green) is shown with values 1, 0, -1. Below it, the text "Filter / kernel" is written in pink. An arrow points from the kernel to a pink speech bubble containing the text "convolution". To the left of the kernel, a blue asterisk (*) indicates the operation. To the right of the kernel, an equals sign (=) is followed by a 4x4 input image (highlighted in red) and a 4x4 output image (highlighted in blue). The input image has values -5, -4, 0, 8 in the first row, and -10, . in the second row. The output image has a value 1 in the top-left cell. The input image is labeled "3x3" in blue at the bottom.

Python : conv-forward
tensorflow : tf.nn.conv2d
Keras : Conv 2D

How it works:

it works:

vertical edge detected

doesn't matter what is in the middle

bright pixels on left dark pixels on right

negative if values flipped
* can take abs value

6x6

3x3

*

*

edge detected

Andrew Ng

Horizontal Edge Detection

10	10	10	0	0	0
10	10	10	0	0	0
10	10	10	0	0	0
0	0	0	10	10	10
0	0	0	10	10	10
0	0	0	10	10	10

6 x 6
n x n

horizontal filter / kernel

*

3 x 3
f x f

=

4 x 4
(n-f+1) x (n-f+1)

$$\begin{matrix}
 1 & 1 & 1 \\
 0 & 0 & 0 \\
 -1 & -1 & -1
 \end{matrix}
 \begin{matrix}
 0 & 0 & 0 & 0 \\
 30 & 10 & -10 & -30 \\
 30 & 10 & -10 & -30 \\
 0 & 0 & 0 & 0
 \end{matrix}$$

Other filter options:

→

1	0	-1
2	0	-2
1	0	1

Sobel filter: puts emphasis on the central row / pixel

→

3	0	-3
10	0	-10
3	0	-3

Scharr filter: puts emphasis on the central row / pixel

→

w_1	w_2	w_3
w_4	w_5	w_6
w_7	w_8	w_9

Learn these as parameters such that this gives us a good edge detector

- Learning is even more robust
- Helps learn these features it's trying to detect
- Can learn edges at angles too

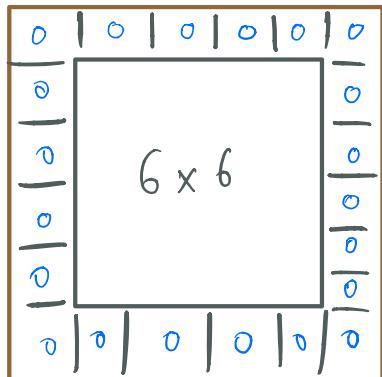
Padding

$$6 \times 6 \text{ mat} \xleftarrow{\text{convolutional}} * 3 \times 3 \text{ mat} = 4 \times 4 \text{ mat}$$

problems

→ shrinking output

→ throwing away info from edges



$$* 3 \times 3 \text{ mat} = 6 \times 6 \text{ mat}$$

$$(n+2p-f+1) \times (n+2p-f+1)$$

8x8 mat padded

- pad with zeros
- p = padding amount = 1 (in this example)

Valid & Same convolutions

Valid → no padding applied

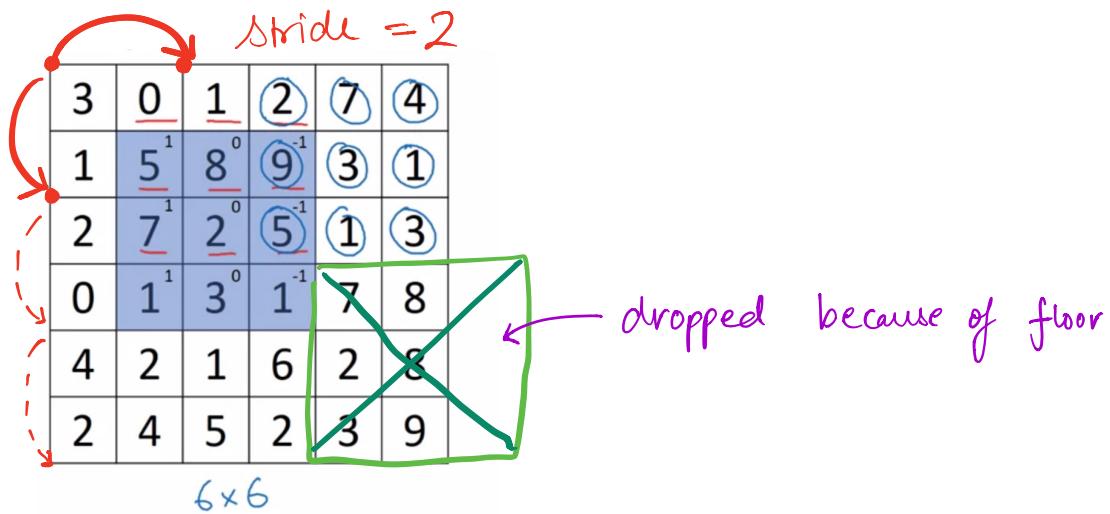
$$n \times n * f \times f = (n-f+1) \times (n-f+1)$$

Same → pad so the output is same as the input size

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

* f is usually odd

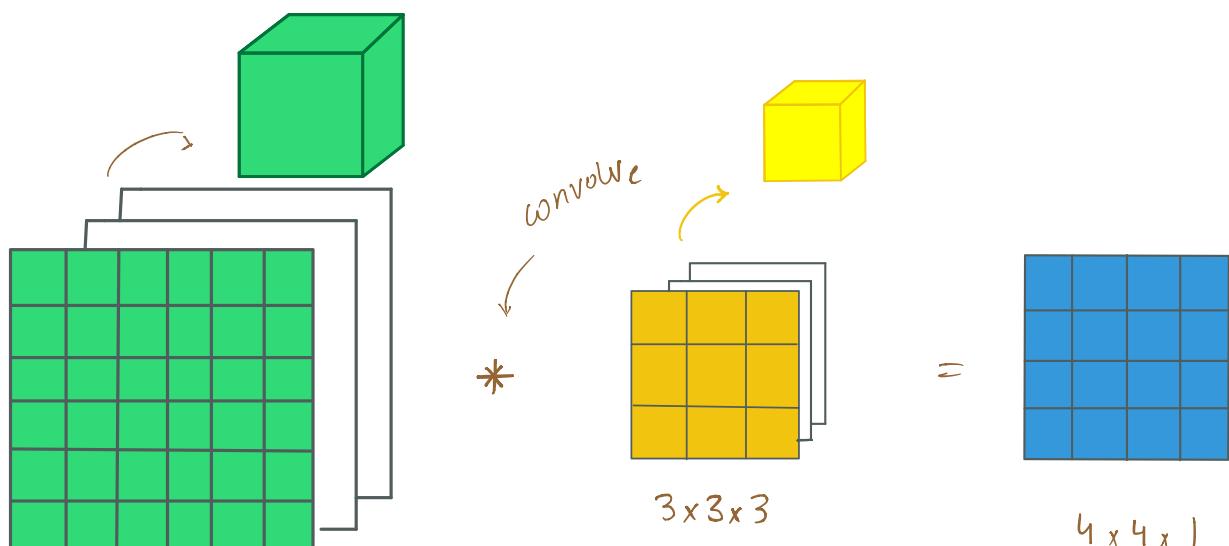
Strided Convolutions



dimensions -

$$n \times n \text{ * } f \times f \quad \xrightarrow{\text{padding } p \quad \text{stride } s} \quad \left\lfloor \frac{n + 2p - f}{s} + 1 \right\rfloor \times \left\lfloor \frac{n + 2p - f}{s} + 1 \right\rfloor$$

Convolutions Over Volume



height $6 \times 6 \times 3$
width width
channels

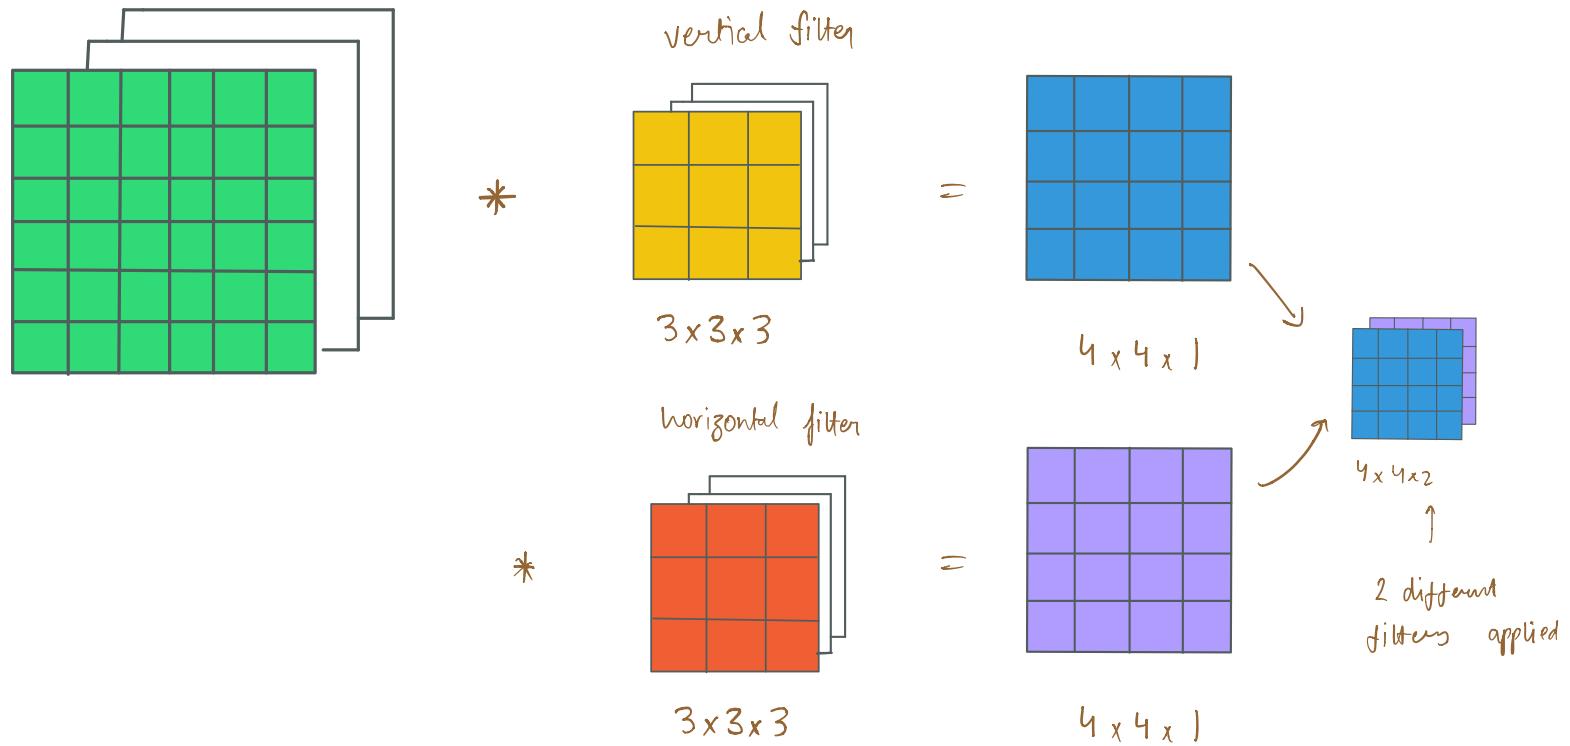
edges in red channel

R
1 0 -1
1 0 -1
1 0 -1

G
0 0 0
0 0 0
0 0 0

B
0 0 0
0 0 0
0 0 0

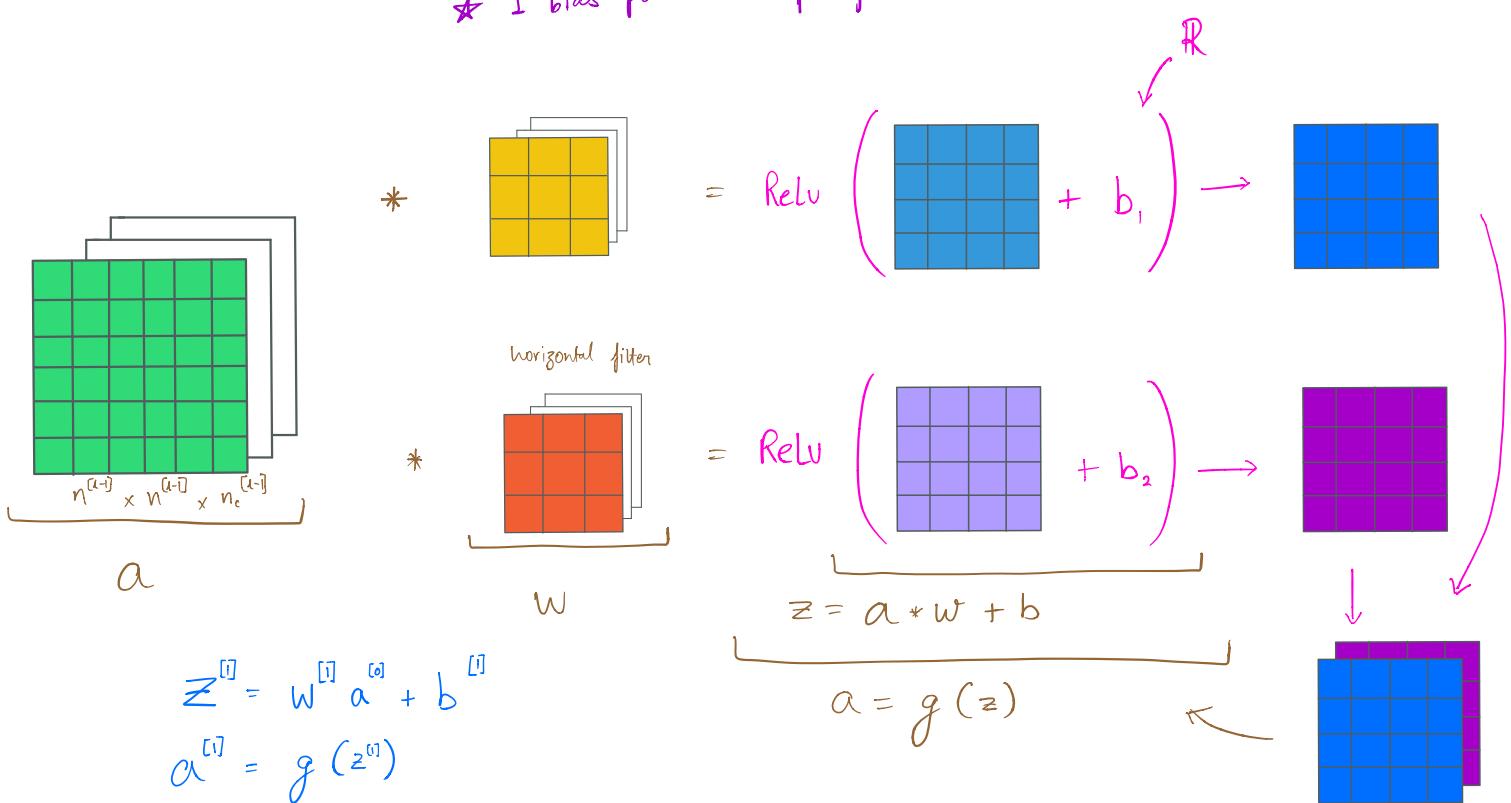
Multiple Filters



dimension:

$$n \times n \times n_c * f \times f \times n_c \rightarrow \frac{n-f+1}{4} \times \frac{n-f+1}{4} \times n_c \quad \begin{matrix} \uparrow \\ \# \text{ filters} \end{matrix} \quad \begin{matrix} \uparrow \\ \text{depth} \end{matrix}$$

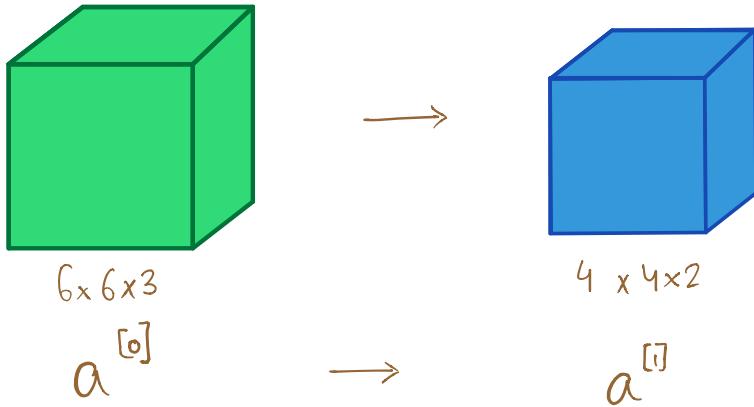
* 1 bias parameter per filter



$$\text{derivatives} - \delta A = \sum_{h=0}^{n_h} \sum_{w=0}^{n_w} w_c \times \delta Z_{hw}$$

$$\delta w_c = \sum_{h=0}^{n_h} \sum_{w=0}^{n_w} a_{\text{slice}} \times \delta Z_{hw}$$

$$\delta b = \sum \sum \delta Z_{hw}$$



Notation

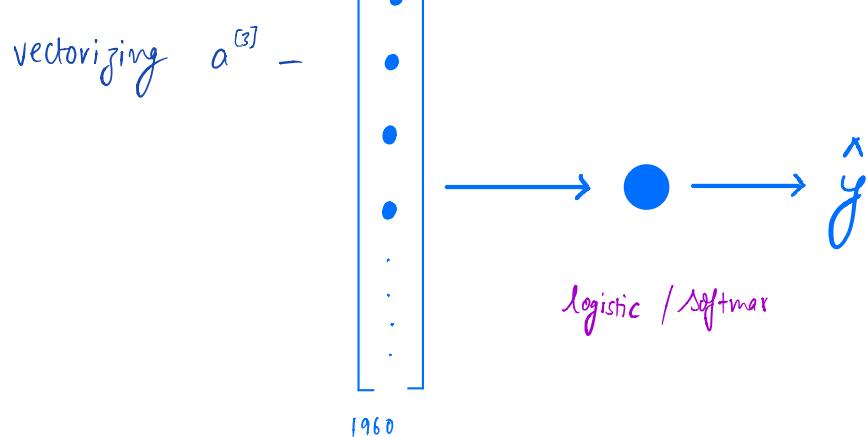
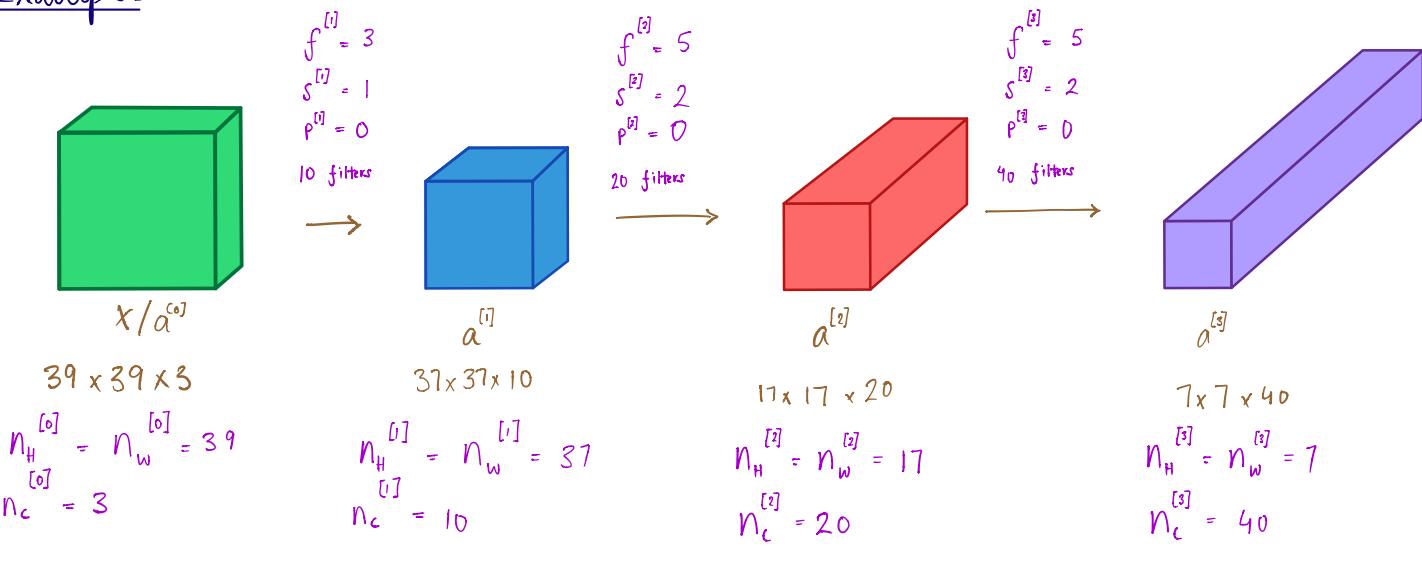
If layer 1 is a convolution layer -

- $f^{[l]}$ = filter size
- $p^{[l]}$ = padding
- $s^{[l]}$ = stride
- input - $n_h^{[l-1]} \times n_w^{[l-1]} \times n_c^{[l-1]}$
- output - $n_h^{[l]} \times n_w^{[l]} \times n_c^{[l]}$
- $n^{[l]} = \left\lfloor \frac{n^{[l-1]} + 2p^{[l]} - f^{[l]}}{s^{[l]}} + 1 \right\rfloor$

$$\text{filter} - f^{[l]} \times f^{[l]} \times n_c^{[l-1]}$$

- activations - $a^{[0]} \rightarrow n_H^{[1]} \times n_W^{[1]} \times n_c^{[1]}$
- weights - $f^{[1]} \times f^{[1]} \times n_c^{[1]} \times n_c^{[1]}$
depth from last layer \uparrow \nwarrow # of filters in layer 1
- bias - $n_c^{[1]} \rightarrow 1, 1, 1, n_c^{[1]}$

Example

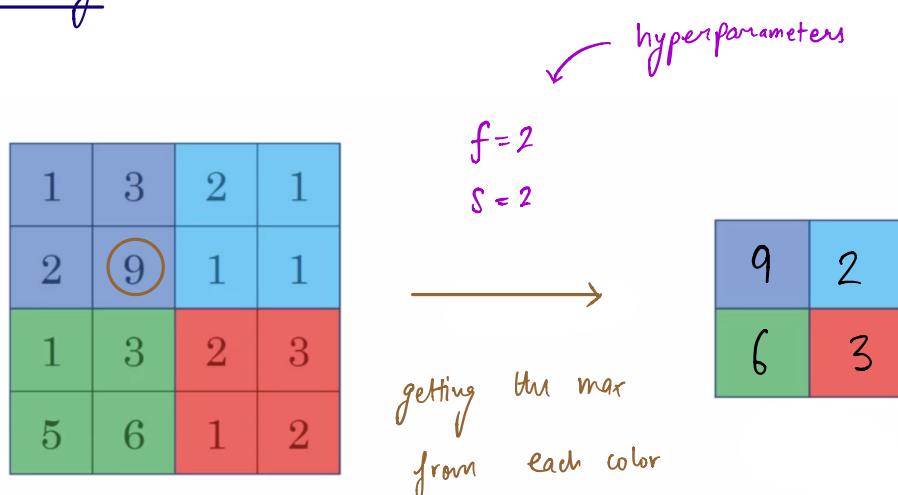


Types of layers in a convolutional network -

- convolution (conv)
- pooling (pool)
- fully connected (FC)

Pooling Layers

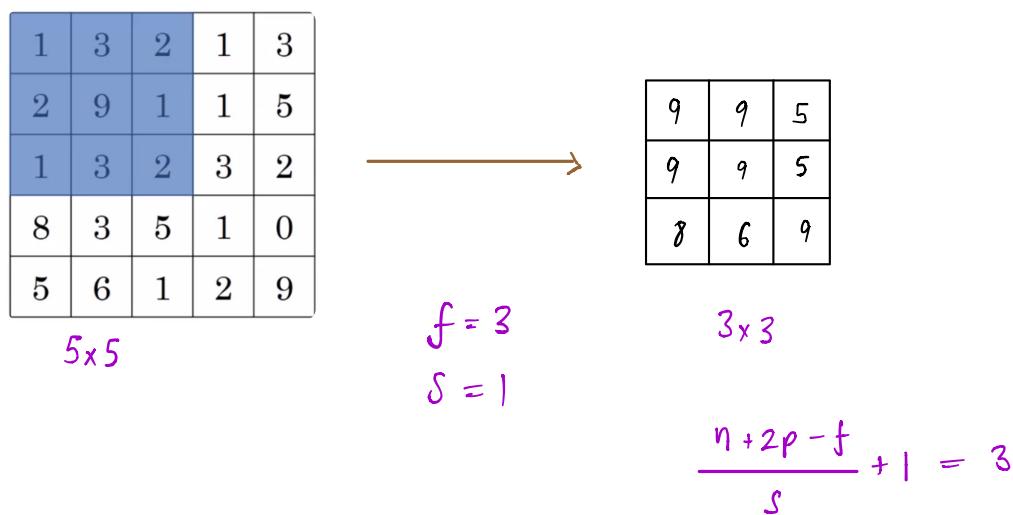
Max Pooling



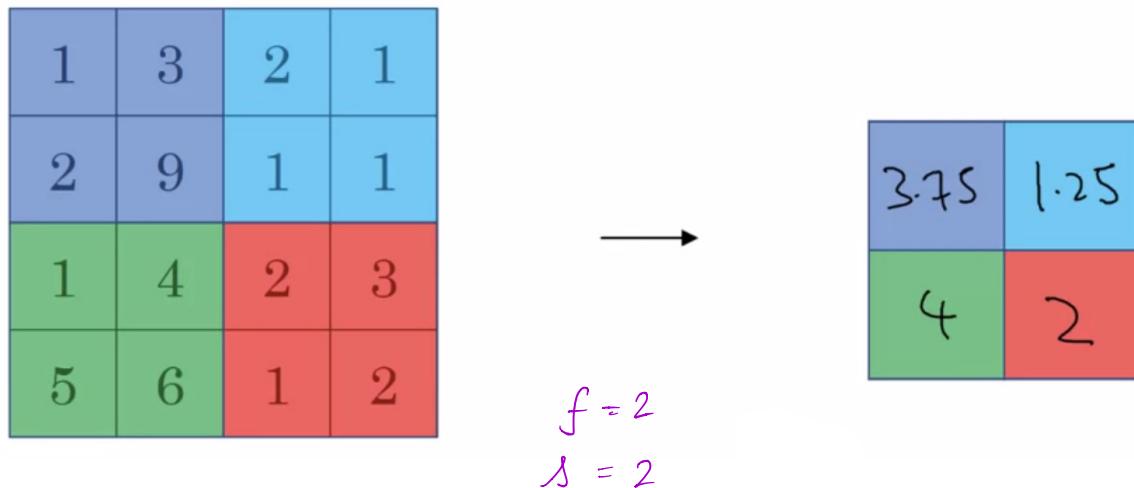
- Some set of features or activations from previous layer, then max pooling detects a particular feature and preserves it due to the high number

→ no parameters, only hyperparameters \Rightarrow fixed computation

example:



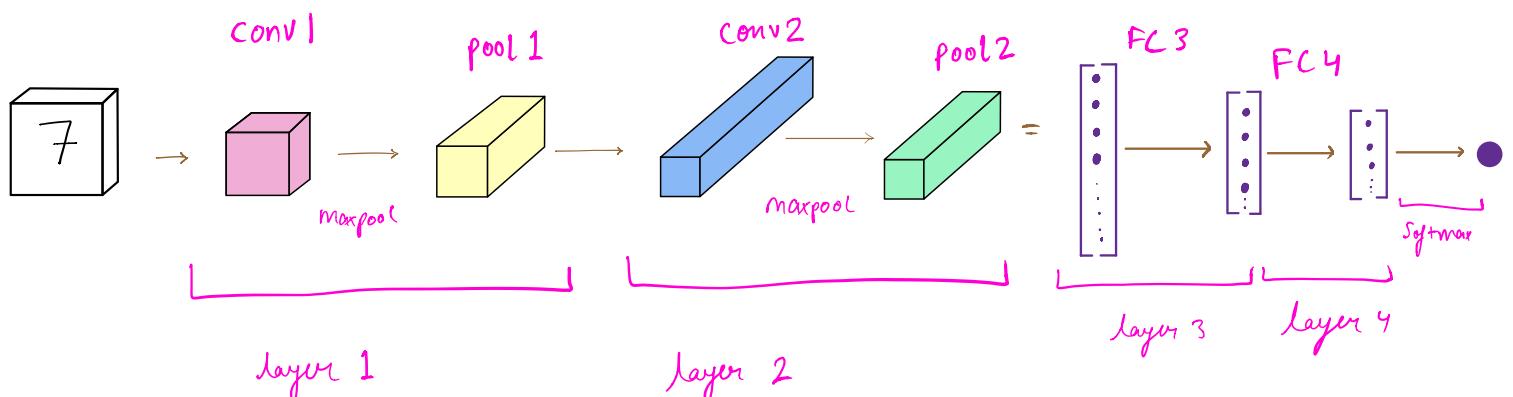
Average Pooling



$$7 \times 7 \times 1000 \rightarrow 1 \times 1 \times 1000$$

→ helps make the size of the matrix smaller by taking the averages

Example of a full Convolutional Network



	Activation shape	Activation Size	# parameters
Input:	(32,32,3)	3,072	0
CONV1 (f=5, s=1)	(28,28,8)	6,272	608
POOL1	(14,14,8)	1,568	0
CONV2 (f=5, s=1)	(10,10,16)	1,600	3216
POOL2	(5,5,16)	400	0
FC3	(120,1)	120	48120
FC4	(84,1)	84	10164
Softmax	(10,1)	10	850

Why Convolutions

CONV nets have comparatively way less parameters

↪ Parameter Sharing - A feature detector like vertical / horizontal edge detectors are useful in all parts of the image

↪ Sparsity of connections - in each layer, each output value depends only on a small number of inputs

As we take a picture suppose RGB, and we apply a filter/kernel to it, we are able to make that picture go through a filter like edge detections and leave out other features of the image. The image we get is particular to that filter. As we apply more filters/kernels to the image, we can stack these output images and we get images that have certain filters applied to them and these images tend to be smaller in width and height as compared to the actual images.

We can view this as a neural network if we try to stack all the pixels of the original picture and connect it to the weights which are the values of the pixels in the filter. The concept of convolutional neural networks is that we connect a lot of these pixels values / input units to the weights / hidden units but not all of them to each one of the hidden units. The missing connections or the less connections help us to end up with an easier network with less parameters than a traditional neural network where all input units are connected to every one of the hidden units.

As use this filter to get another image that we feed again to another convolutional network, this helps us separate or individually feed these features to further layers which eventually help us identify what the image is. We train the parameters, i.e the pixel values of the filter, to make new filters that help us detect features important in identifying the image. As we apply the filter we get a part or a feature of the image important for our classification.



What pooling does is it helps us to kind of average out and make the size of the height and width smaller to prevent having a large number of parameters. It is like a control check on the network and helps us keep the number of parameters in control.

