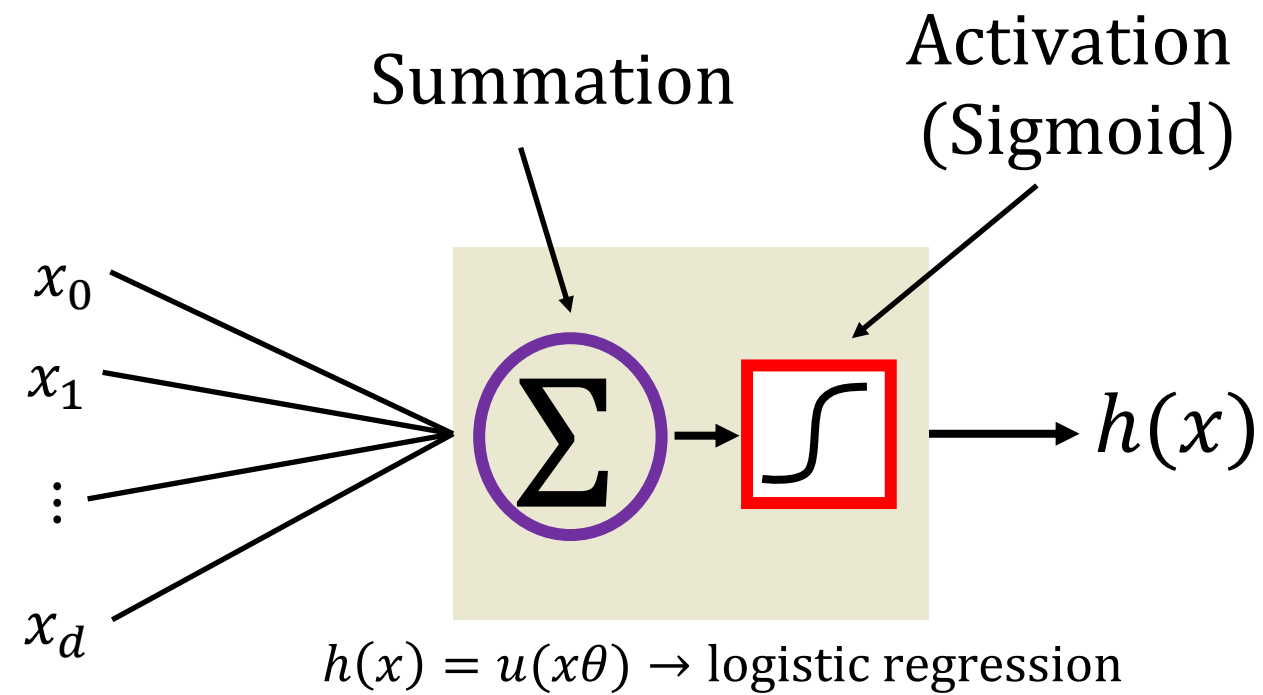


CONVOLUTIONAL NEURAL NETWORK

Nakul Gopalan
Georgia Tech



Slides are based on Ming Li and Mahdi Roozbahani



$$\text{output} = \text{activation}(x\theta + b)$$


Name of the neuron	Activation function: $\text{activation}(z)$
Linear unit	$x\theta$
Threshold/sign unit	$\text{sign}(x\theta)$
Sigmoid unit	$\frac{1}{1 + \exp(-x\theta)}$
Rectified linear unit (ReLU)	$\max(0, x\theta)$
Tanh unit	$\tanh(x\theta)$

Put an image in



Smaller Network: CNN


- We know it is good to learn a small model.
- From this fully connected model, do we really need all the edges?
- Can some of these be ^{red}shared?



Consider learning an image:



- Some patterns are much smaller than the whole image

Can represent a small region with fewer parameters




Same pattern appears in different places:
They can be compressed!

What about training a lot of such “small” detectors
and each detector must “move around”.



A convolutional layer

A CNN is a neural network with some convolutional layers (and some other layers). A convolutional layer has a number of filters that does convolutional operation. Neocognitron by Kunihiko Fukushima (1980).



Convolution

1	0	0	0	0	1
0	1	0	0	1	0
0	0	1	1	0	0
1	0	0	0	1	0
0	1	0	0	1	0
0	0	1	0	1	0

6 x 6 image

These are the network parameters to be learned.

weights

1	-1	-1
-1	1	-1
-1	-1	1

Filter 1

-1	1	-1
-1	1	-1
-1	1	-1

Filter 2

⋮ ⋮

*Filtered to 10
→ 1000
detected*

Each filter detects a small pattern (3 x 3).

Convolution


1	-1	-1
-1	1	-1
-1	-1	1

Filter 1

stride=1

1	0	0	0	0	1
0	1	0	0	1	0
0	0	1	1	0	0
1	0	0	0	1	0
0	1	0	0	1	0
0	0	1	0	1	0

Dot product



6 x 6 image

Convolution


If stride=2

1	0	0	0	0	1
0	1	0	0	1	0
0	0	1	1	0	0
1	0	0	0	1	0
0	1	0	0	1	0
0	0	1	0	1	0

6 x 6 image

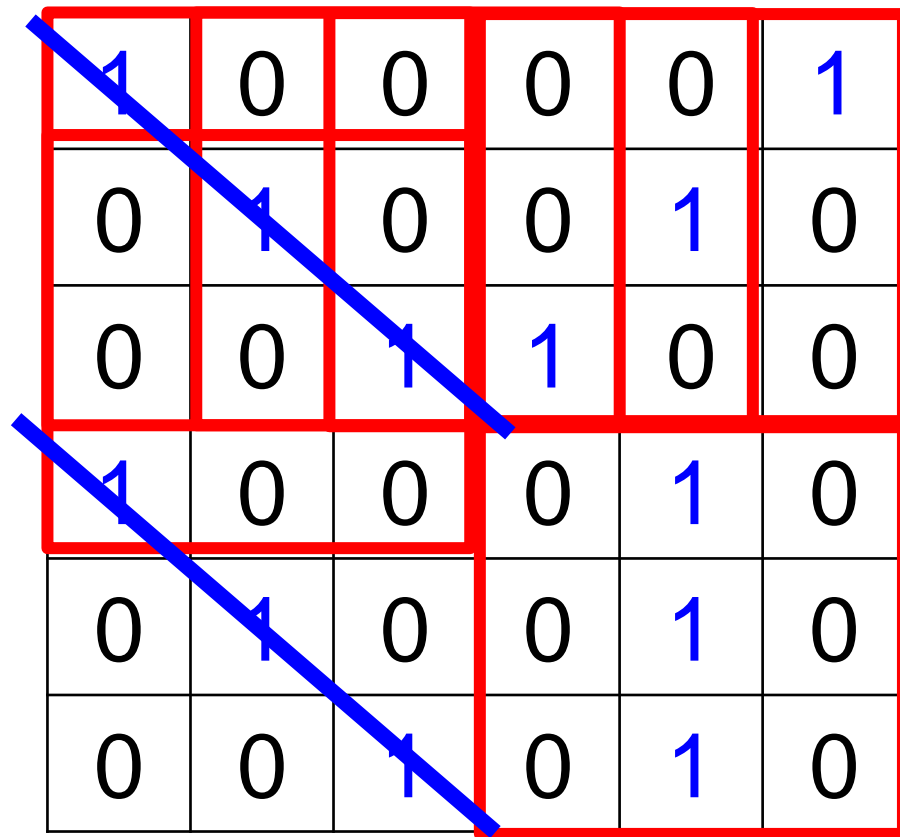
1	-1	-1
-1	1	-1
-1	-1	1

Filter 1



Convolution – diagonal edges?

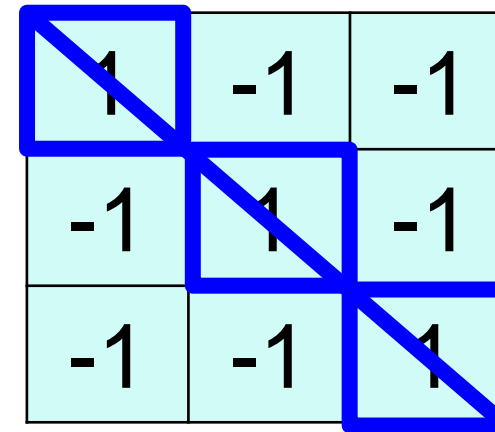
stride=1



A 6x6 grid representing an image. The grid is outlined with a thick red border. A blue diagonal line runs from the top-left corner to the bottom-right corner. The values in the grid are as follows:

1	0	0	0	0	1
0	1	0	0	1	0
0	0	1	1	0	0
1	0	0	0	1	0
0	1	0	0	1	0
0	0	1	0	1	0

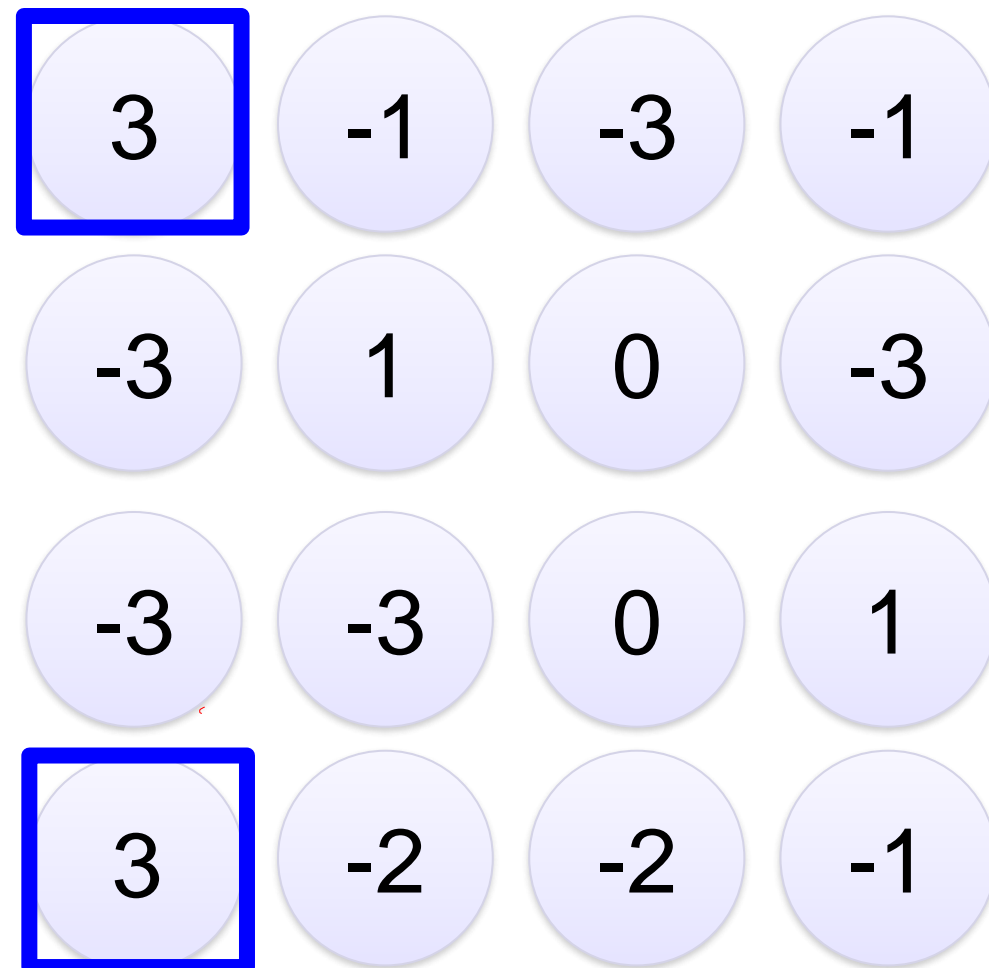
6 x 6 image



A 3x3 grid representing a filter. The grid is outlined with a thick blue border. A blue diagonal line runs from the top-left corner to the bottom-right corner. The values in the grid are as follows:

1	-1	-1
-1	1	-1
-1	-1	1

Filter 1



A 4x4 grid of circles representing the convolution result. The top-left circle (0,0) and the bottom-left circle (3,0) are highlighted with a thick blue border. The values in the circles are as follows:

3	-1	-3	-1
-3	1	0	-3
-3	-3	0	1
3	-2	-2	-1

Convolution - Vertical edges?

stride=1


1	0	0	0	0	1
0	1	0	0	1	0
0	0	1	1	0	0
1	0	0	0	1	0
0	1	0	0	1	0
0	0	1	0	1	0

6 x 6 image

-1	1	-1
-1	1	-1
-1	1	-1

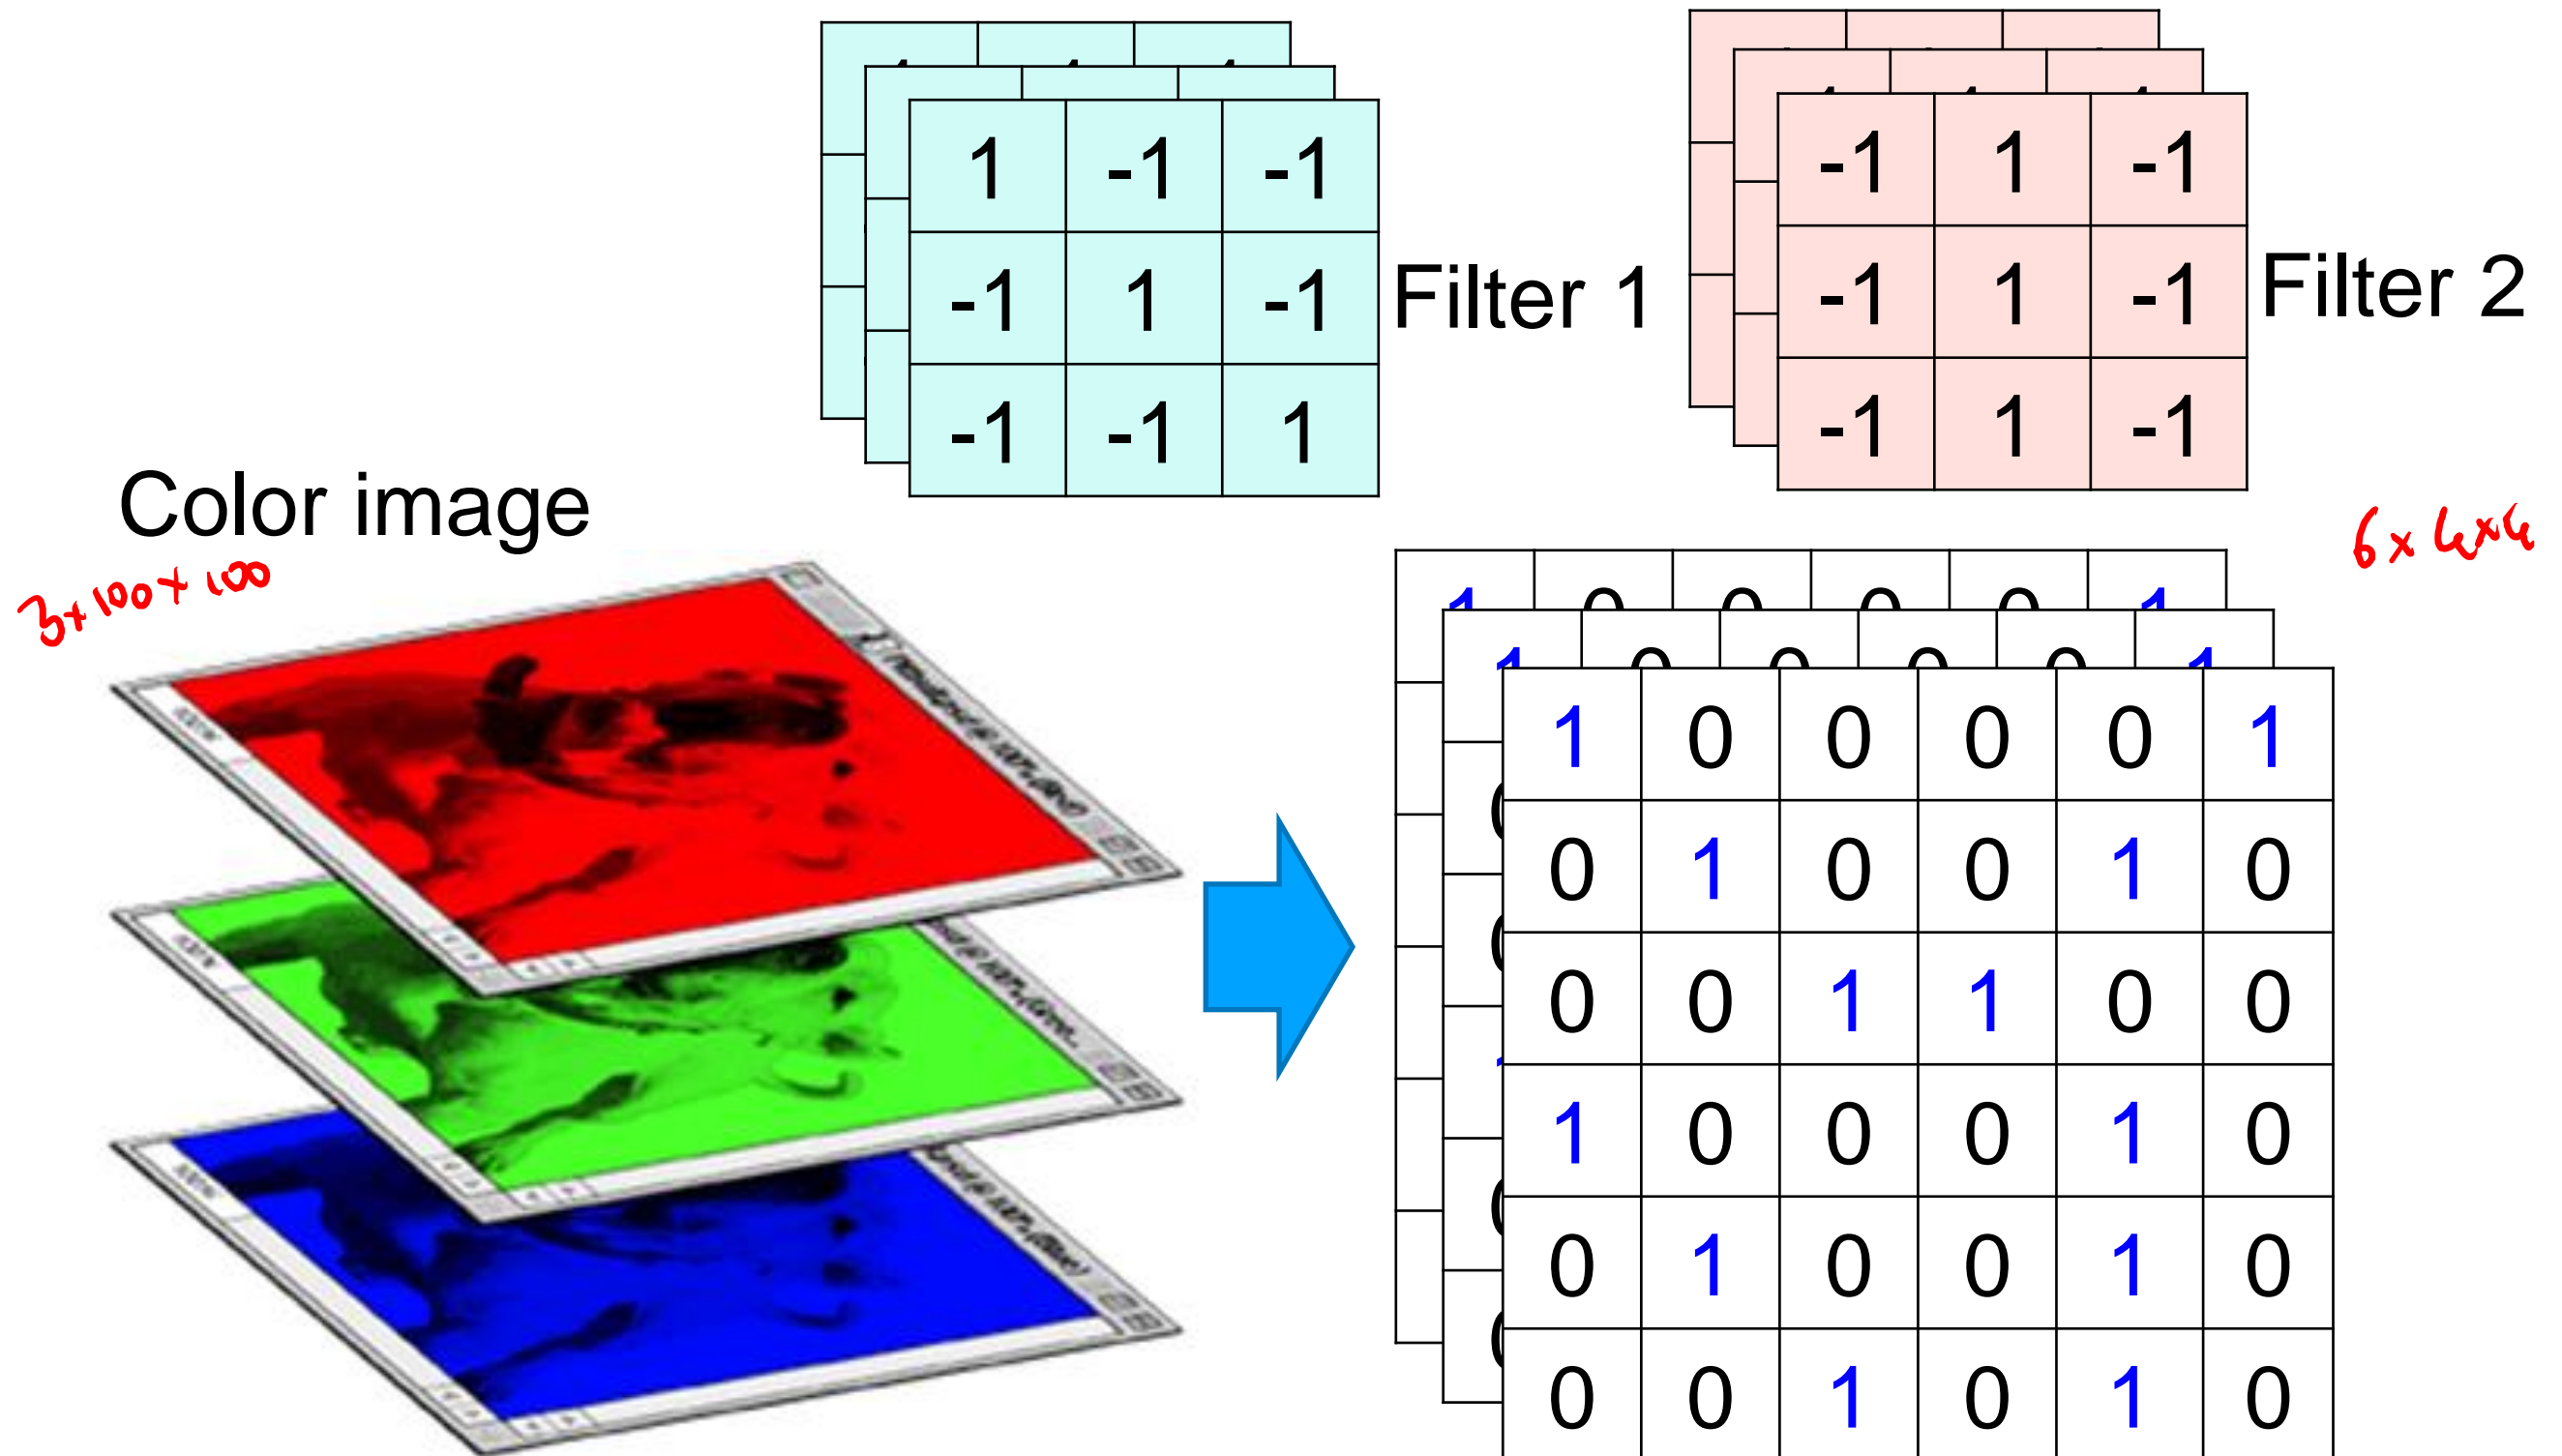
Filter 2

Repeat this for each filter



Two 4 x 4 images
Forming 2 x 4 x 4 matrix

Color image: RGB 3 channels



Convolution v.s. Fully Connected

1	0	0	0	0	1
0	1	0	0	1	0
0	0	1	1	0	0
1	0	0	0	1	0
0	1	0	0	1	0
0	0	1	0	1	0

image

1	-1	-1
-1	1	-1
-1	-1	1

-1	1	-1
-1	1	-1
-1	1	-1




convolution

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

Fully-
connected

1	0	0	0	0	1
0	1	0	0	1	0
0	0	1	1	0	0
1	0	0	0	1	0
0	1	0	0	1	0
0	0	1	0	1	0




Conventional
Fully Connected
layers
(FC layers)

1	0	0	0	0	1
0	1	0	0	1	0
0	0	1	1	0	0
1	0	0	0	1	0
0	1	0	0	1	0
0	0	1	0	1	0


6 x 6 image


1	1	
2	0	
3	0	
4	0	
5	0	
6	1	
7	0	
8	1	
⋮		⋮
31	0	
32	0	
33	1	
34	0	
35	1	
36	0	

features 1st hidden layer




fewer parameters!







Fewer parameters

Even fewer parameters



An example classifier using CNNs



Can repeat many times

Flattened



Max Pooling

1	-1	-1
-1	1	-1
-1	-1	1

Filter 1

-1	1	-1
-1	1	-1
-1	1	-1

Filter 2



Why Pooling

- Subsampling pixels will not change the object
bird




Subsampling



bird

We can subsample the pixels to make image smaller



fewer parameters to characterize the image


A CNN compresses a fully connected network in two ways:

- Reducing number of connections
- Shared weights on the edges
- Moreover, Max pooling further reduces the complexity

Max Pooling

1	0	0	0	0	1
0	1	0	0	1	0
0	0	1	1	0	0
1	0	0	0	1	0
0	1	0	0	1	0
0	0	1	0	1	0


6 x 6 image




Conv



Max
Pooling




New image
but smaller



2 x 2 image

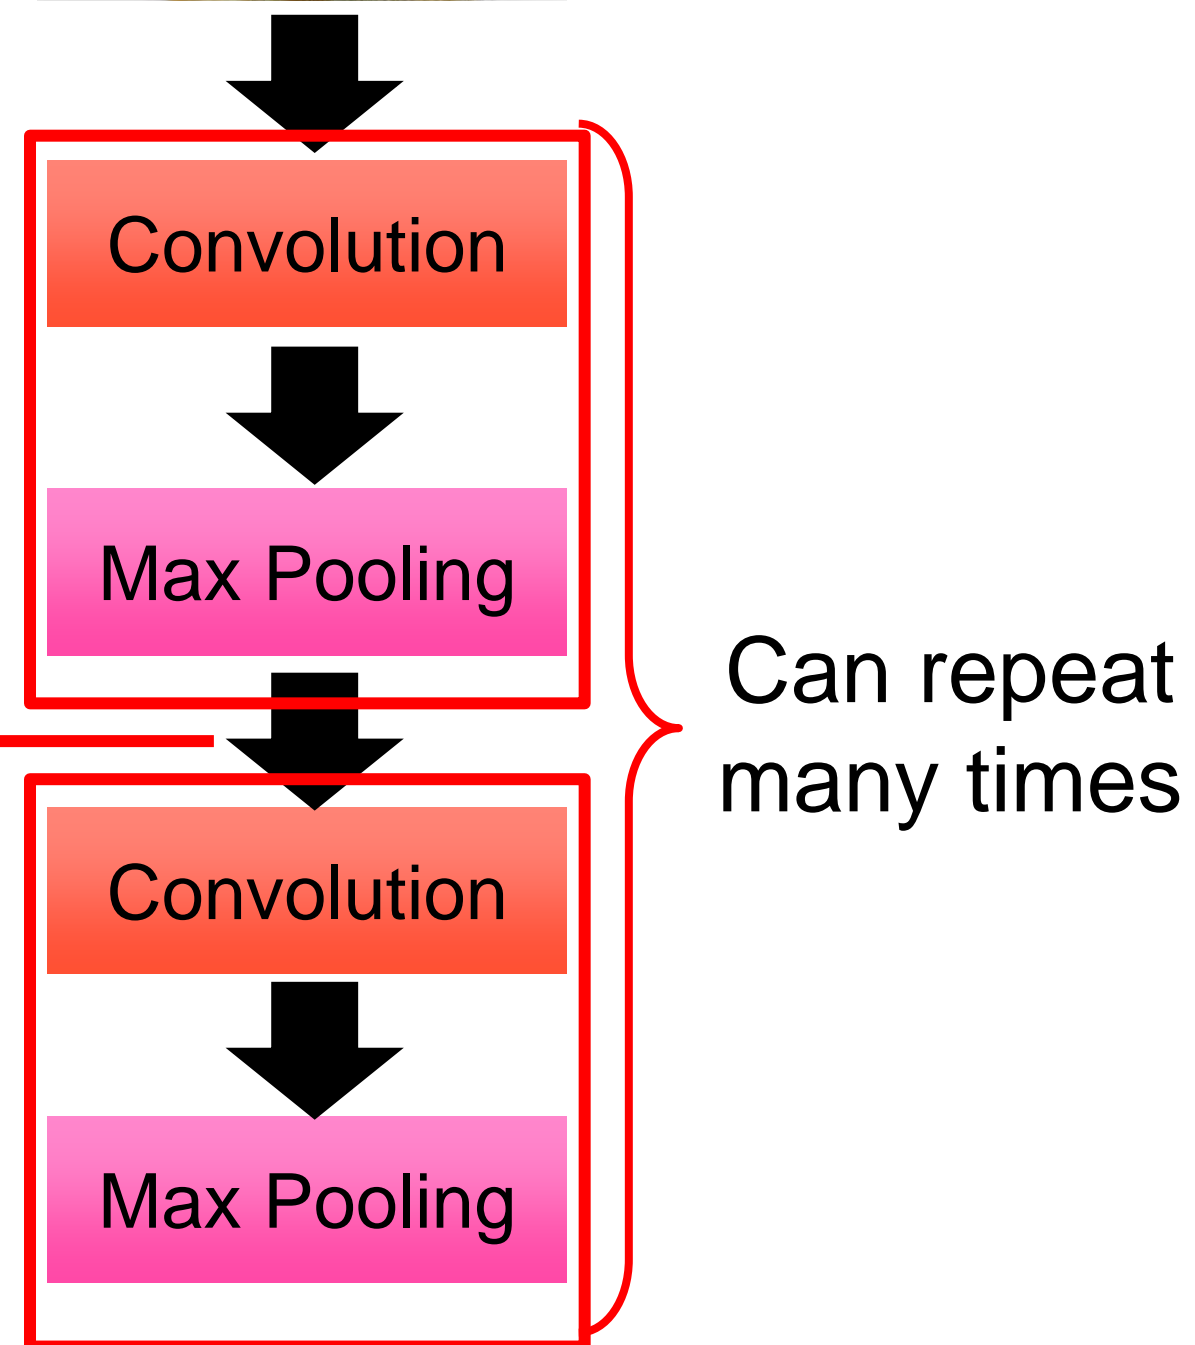
Each filter
is a channel

Example network




Smaller than the original image

The number of channels is the number of filters



Example network

cat dog



Convolution

Max Pooling

A new image

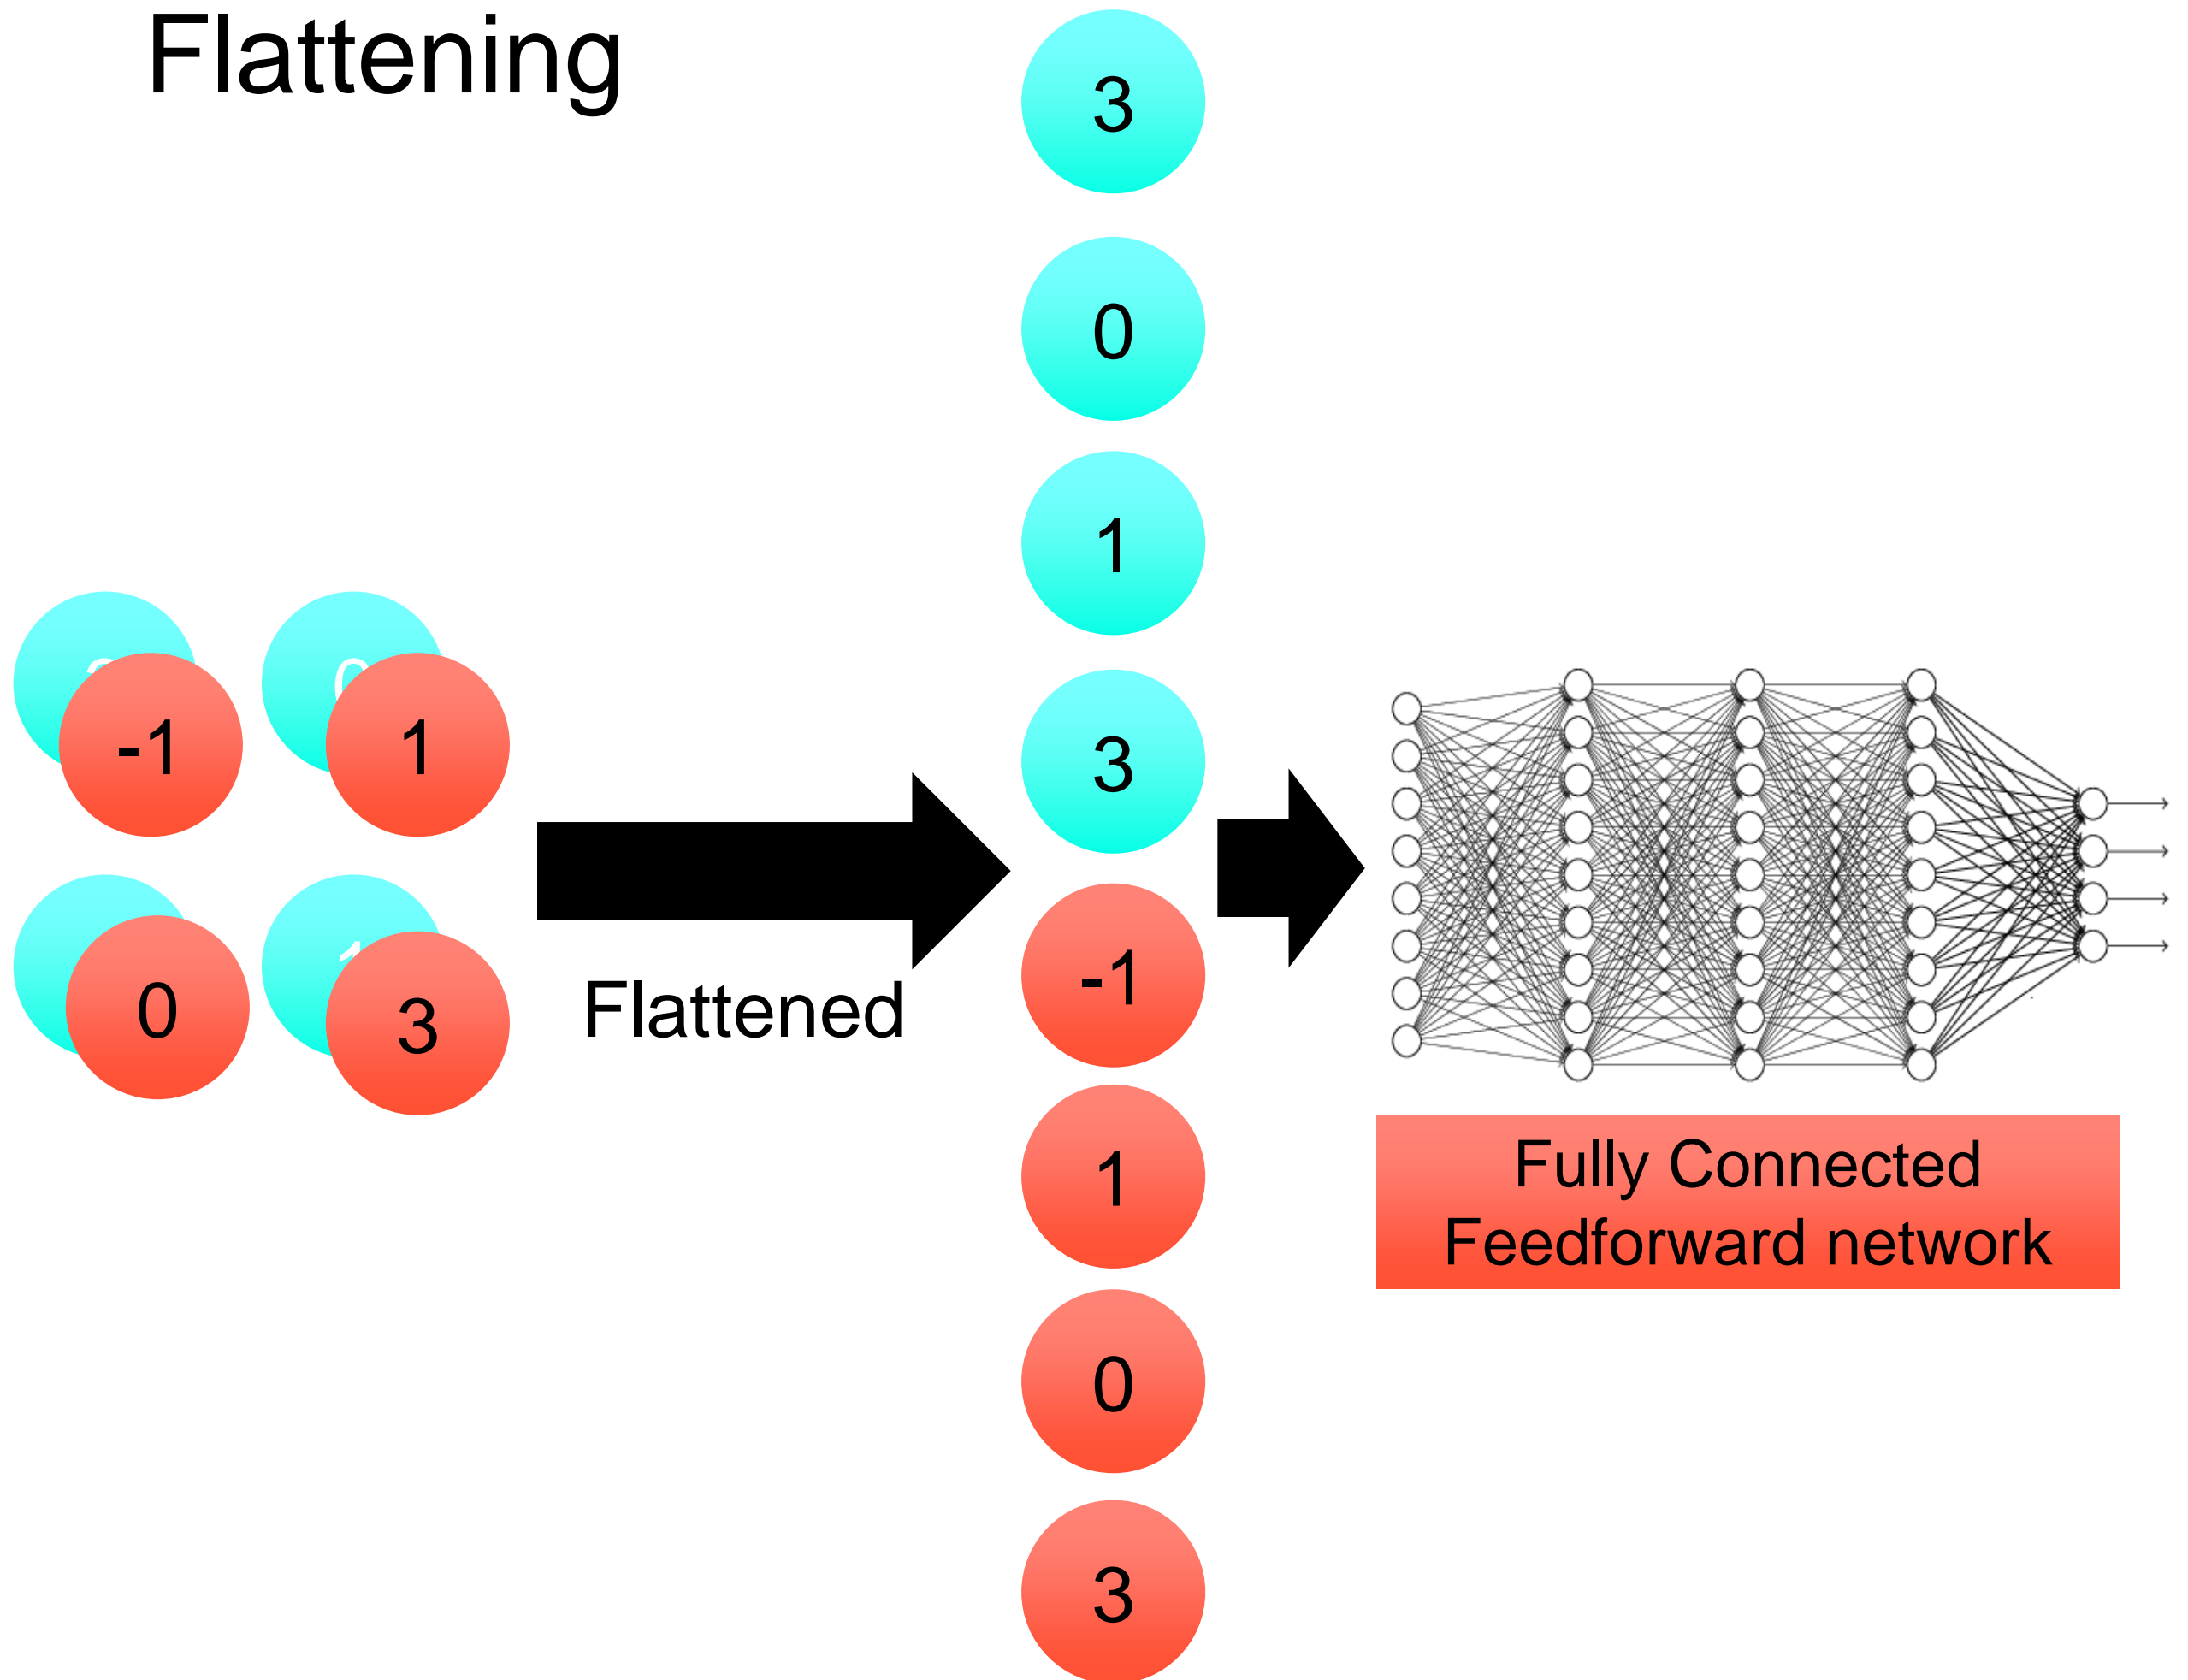
Convolution

Max Pooling

A new image

Flattened

Flattening



CNN in Keras

Only modified the *network structure* and *input format* (vector -> 3-D tensor)

```
model2.add( Convolution2D( 25,3,3,  
                           input_shape=(28,28,1)) )
```

1	-1	-1	1	-1
-1	1	-1	1	-1
-1	-1	-1	1	-1


There are
25 3x3
filters.

Input_shape = (28 , 28 , 1)

28 x 28 pixels

1: black/white, 3: RGB

```
model2.add(MaxPooling2D((2,2)))
```



input

Convolution

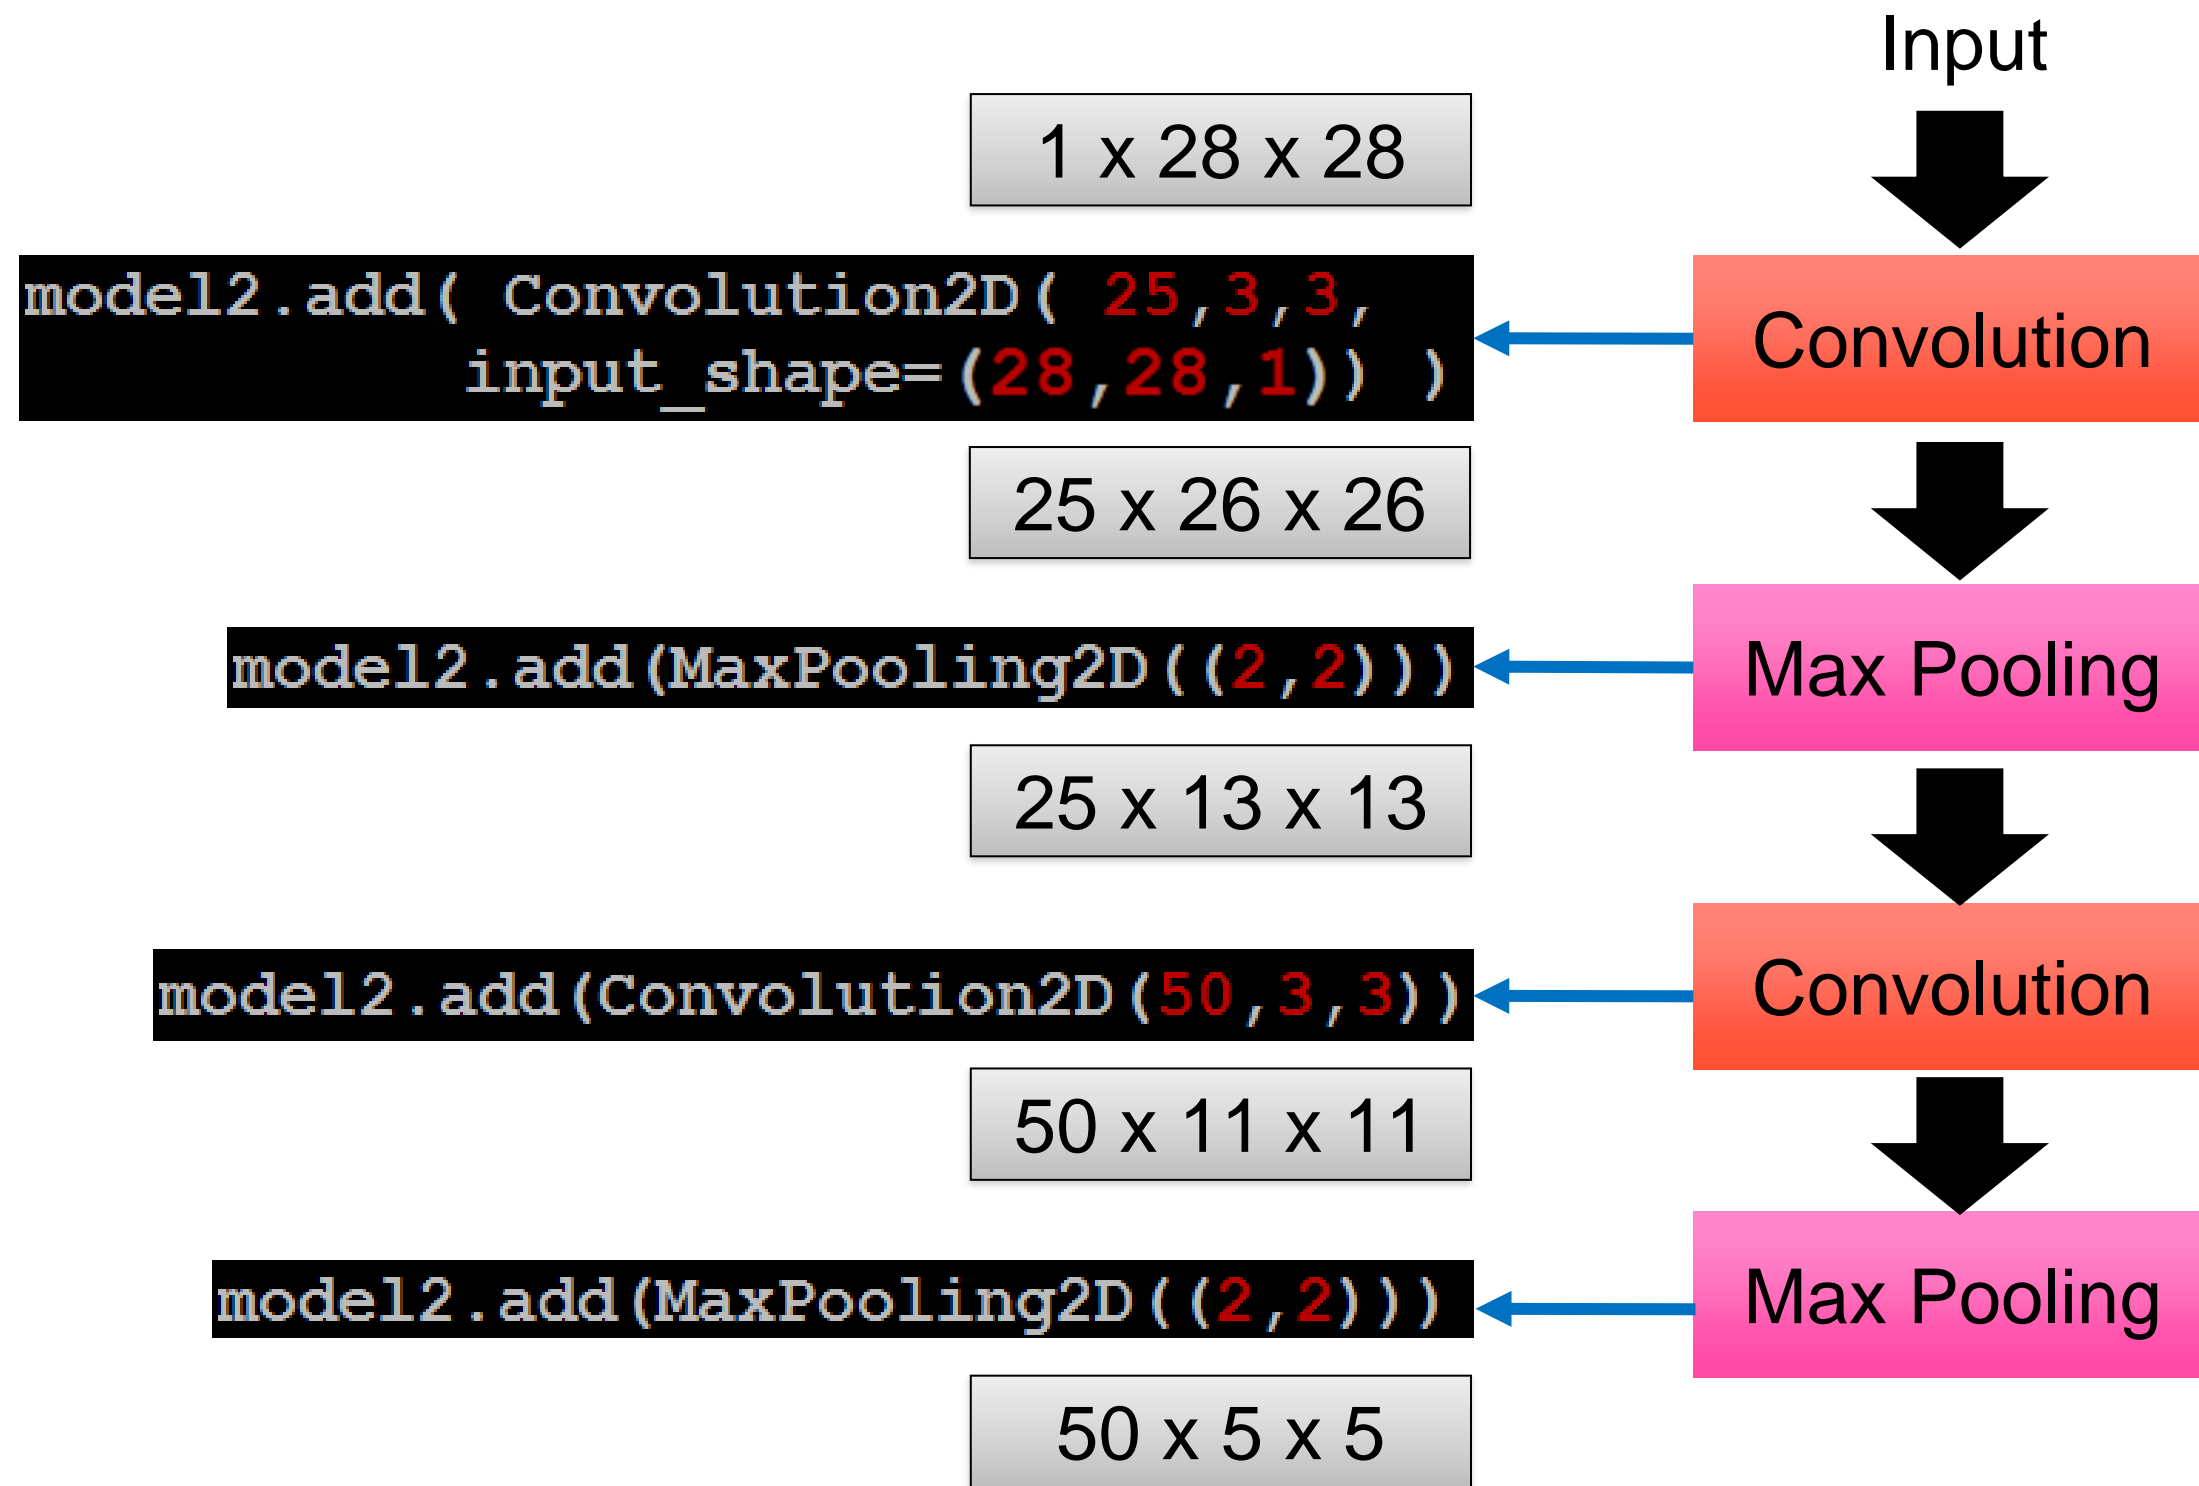
Max Pooling

Convolution

Max Pooling

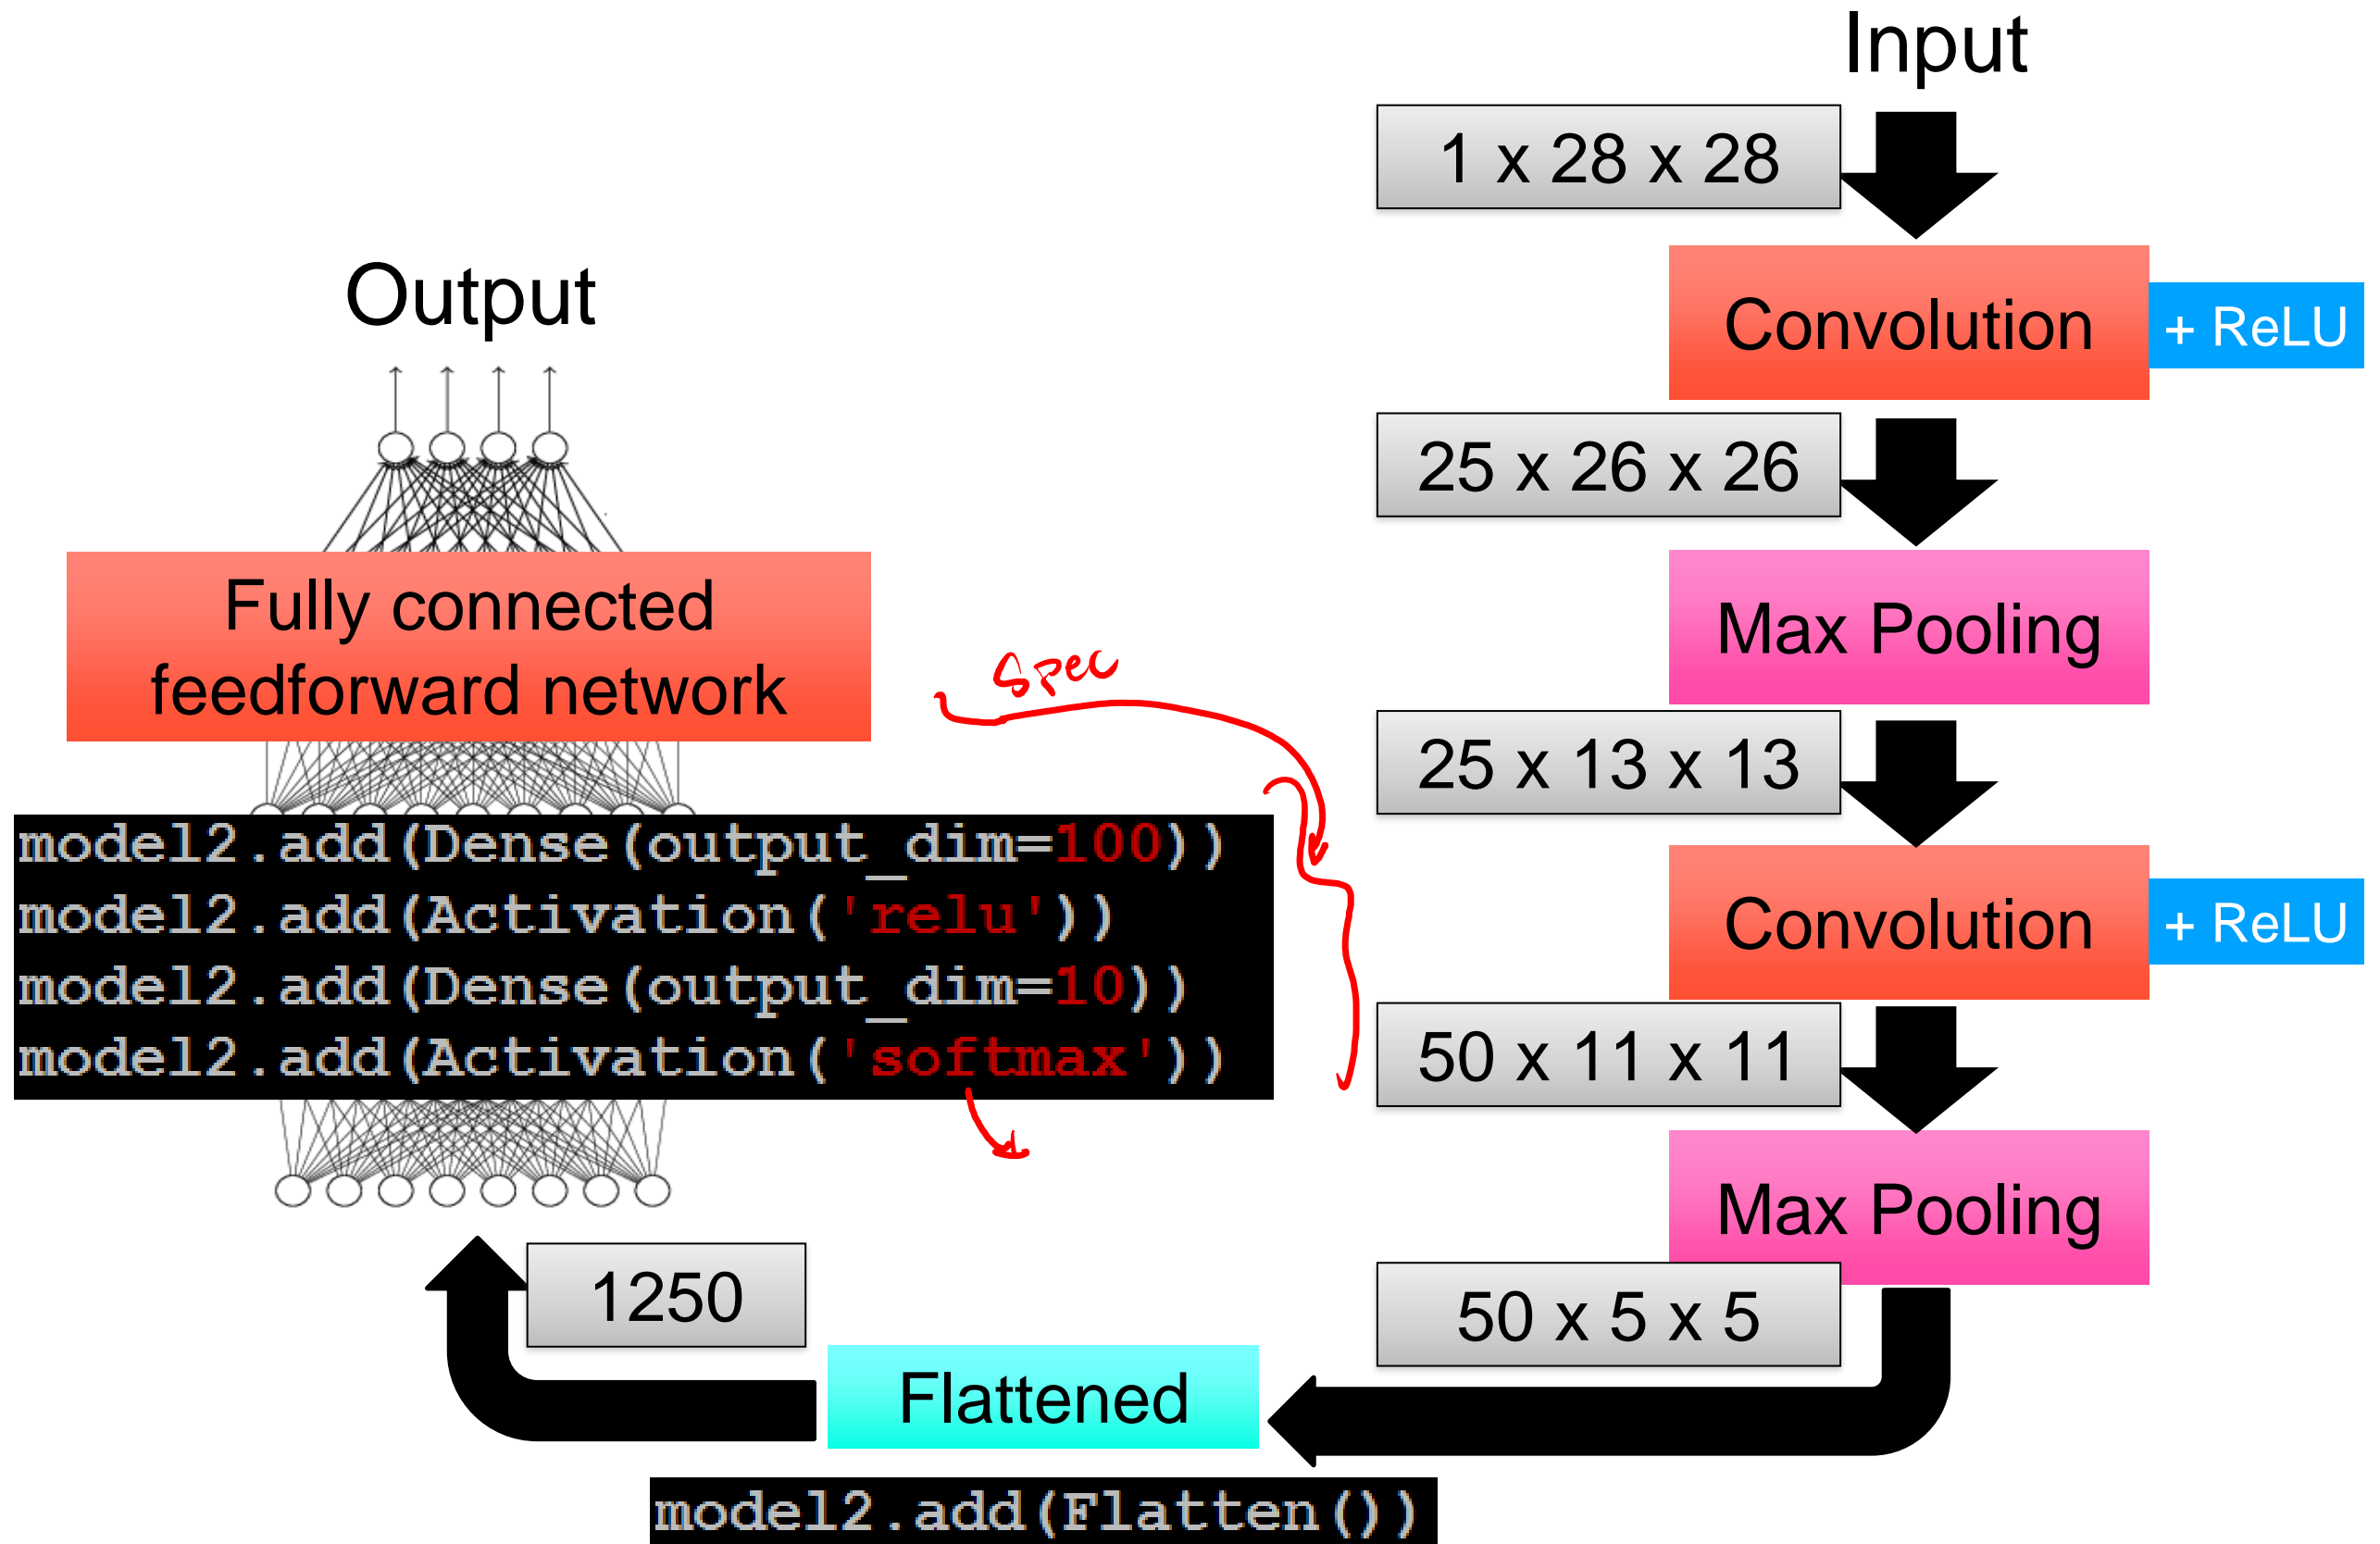
CNN in Keras

Only modified the *network structure* and *input format* (vector -> 3-D array)



CNN in Keras

Only modified the *network structure* and *input format* (vector -> 3-D array)




Number of Parameters

$25 \times 3 \times 3 + 25$ parameters


25 filters - Conv1

25: 3×3




Gray scaled-image



Convolved result for 25 filters



Result after Max Pooling




50: $3 \times 3 \times 25$

50 filters - Conv2


$50 \times 3 \times 3 \times 25 + 50$ parameters

50

Convolved result for 50 filters





Result after Max Pooling



neurons after flattening: $50 \times 5 \times 5 = 1250$

Flattening



10 CNN Architecture