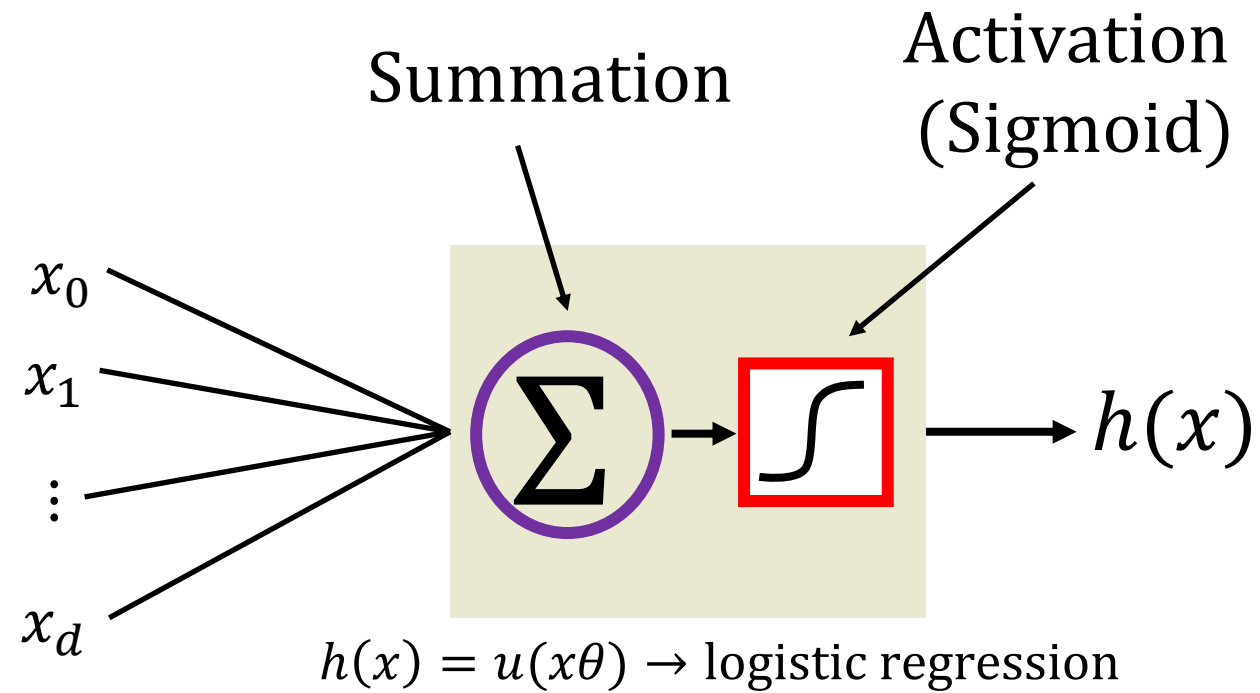


CONVOLUTIONAL NEURAL NETWORK

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 shared?



Consider learning an image:

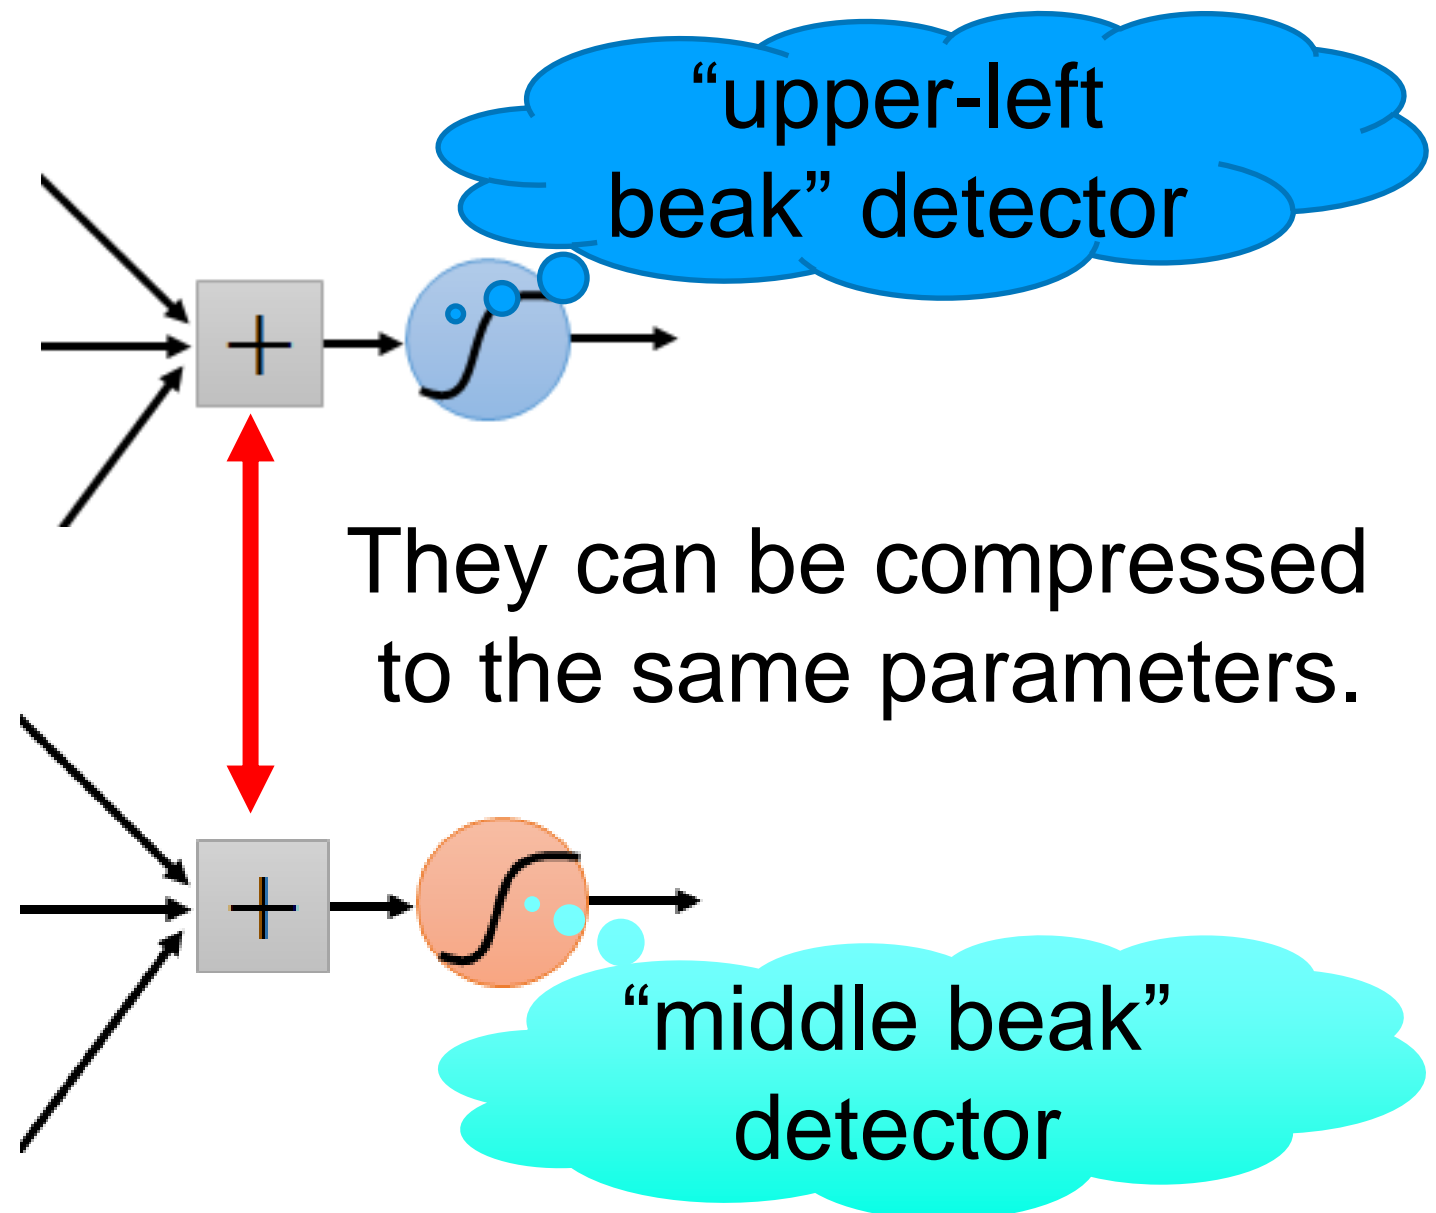
- Some patterns are much smaller than the whole image

Can represent a small region with fewer parameters



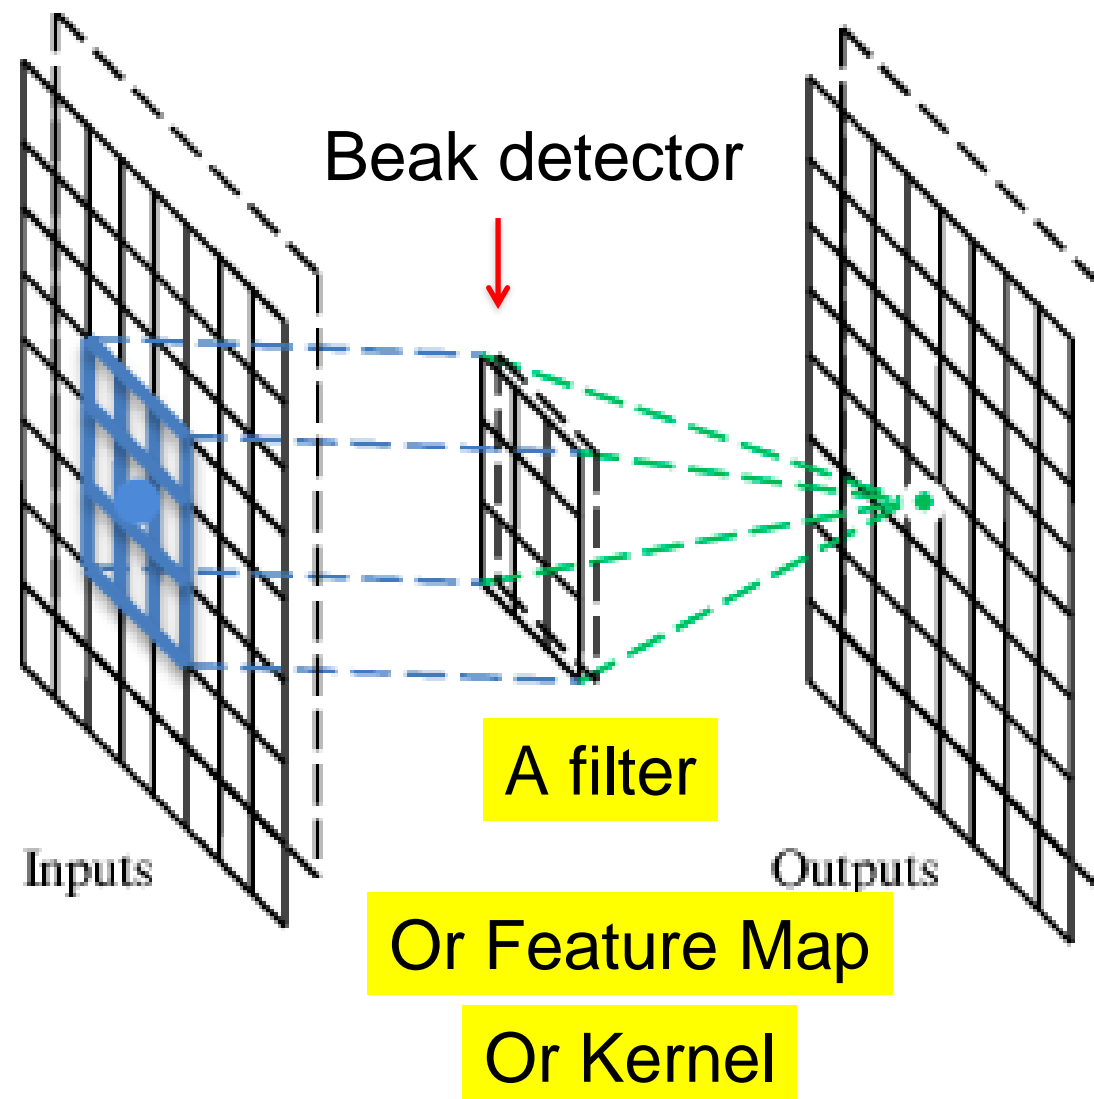
Same pattern appears in different places:
They can be compressed!

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 Kunihiro 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.

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

Filter 1

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

Filter 2

⋮ ⋮

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

Convolution

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 1



Convolution


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

Filter 1

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




Convolution – diagonal edges?


stride=1



6 x 6 image



Filter 1



Convolution - Vertical edges?

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

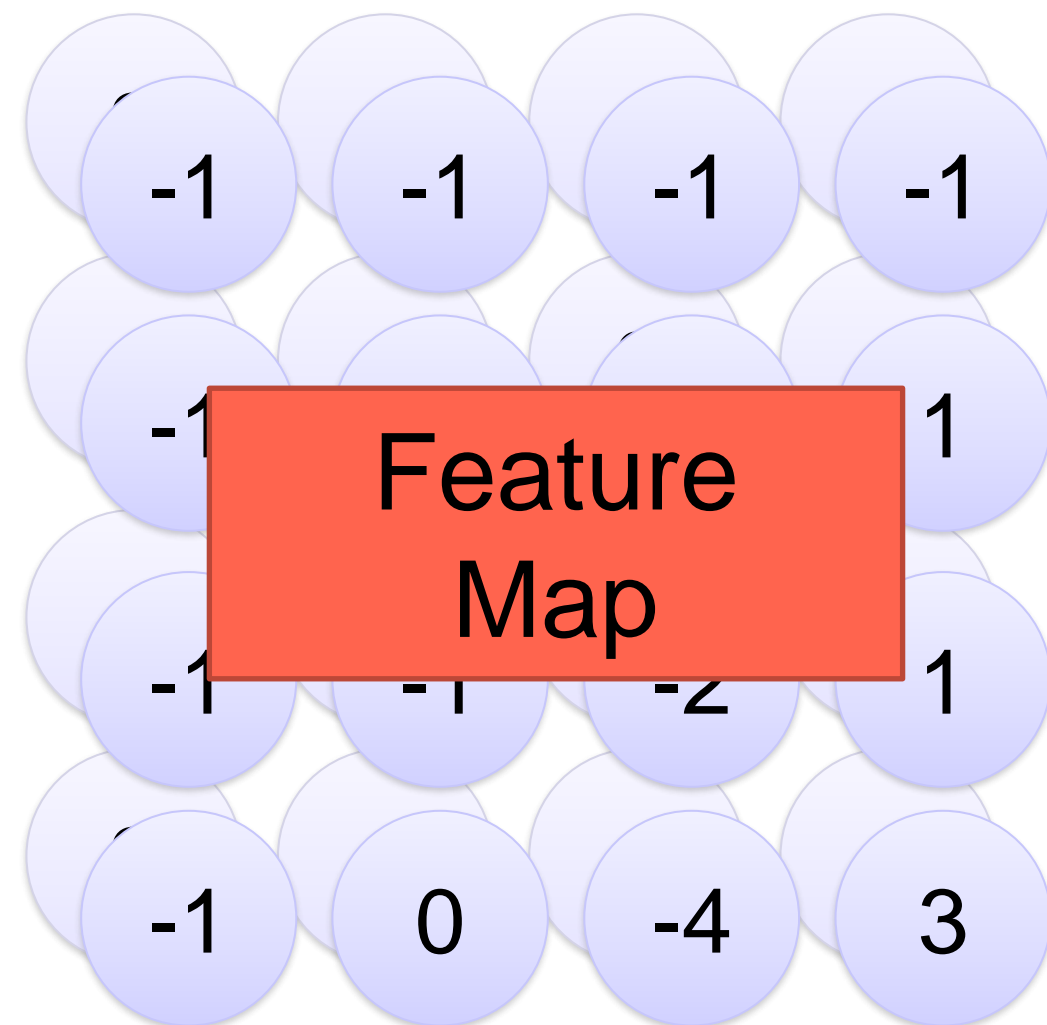
Filter 2

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

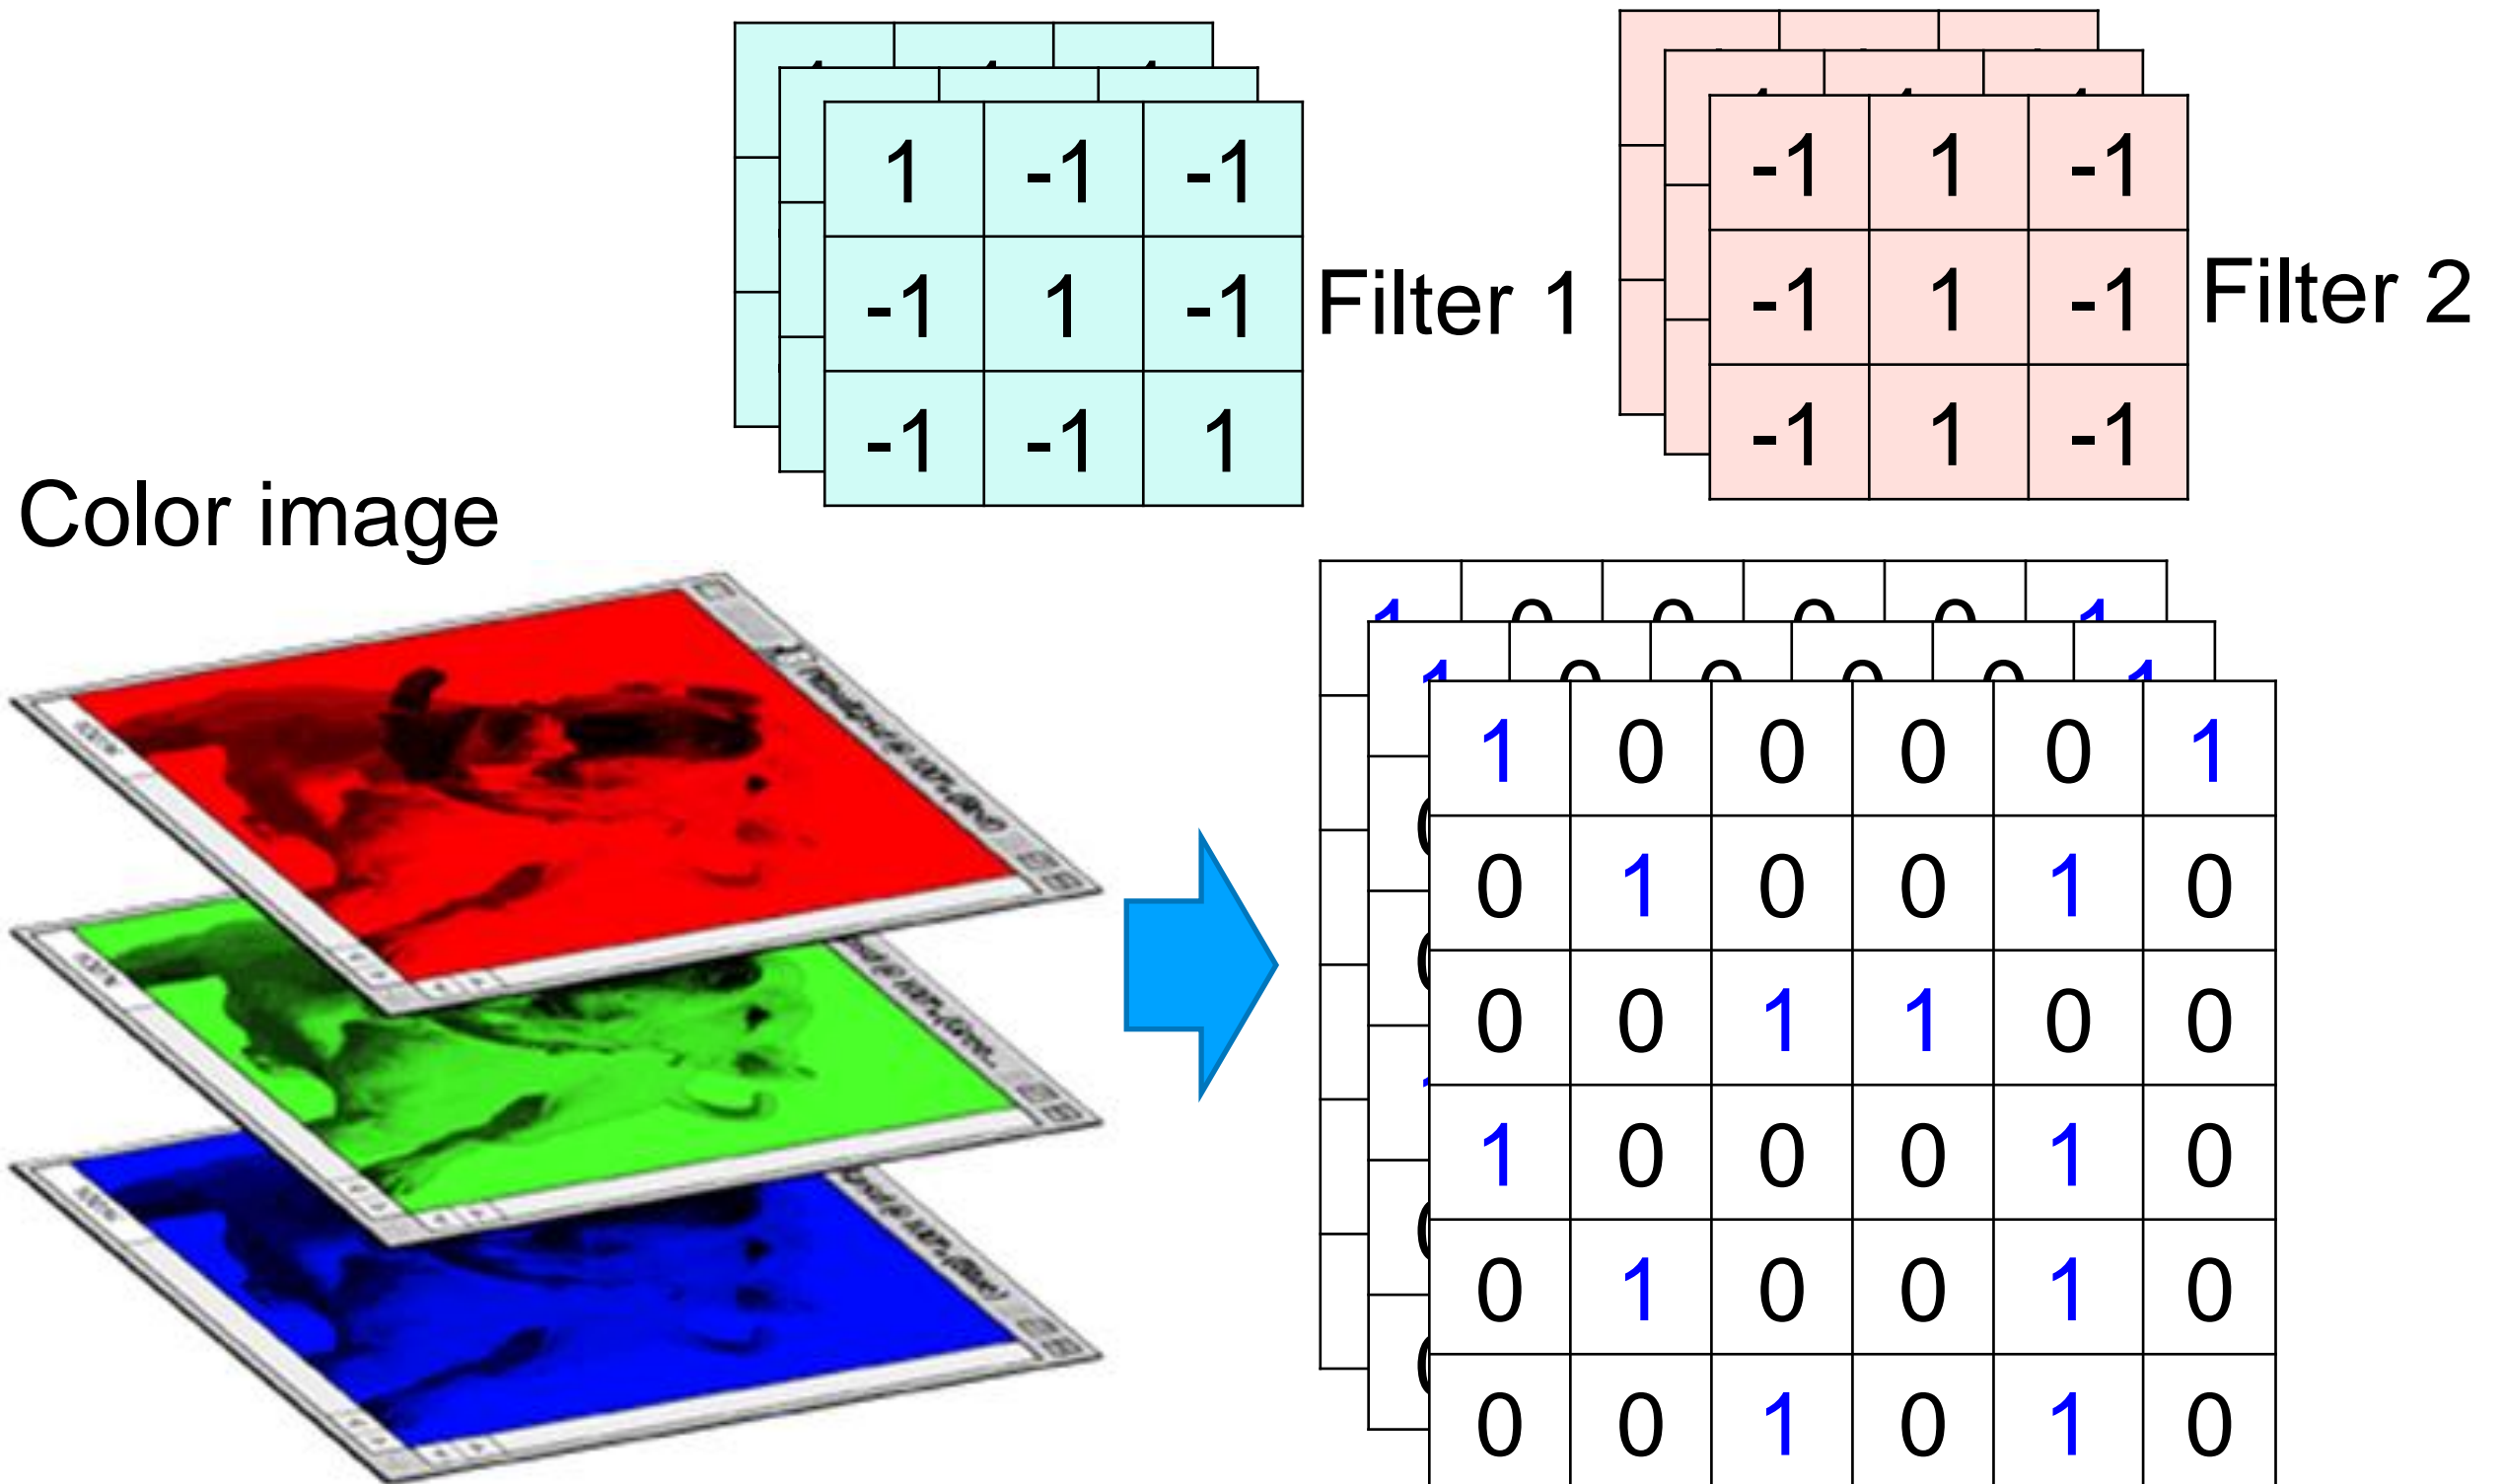
6 x 6 image

Repeat this for each filter

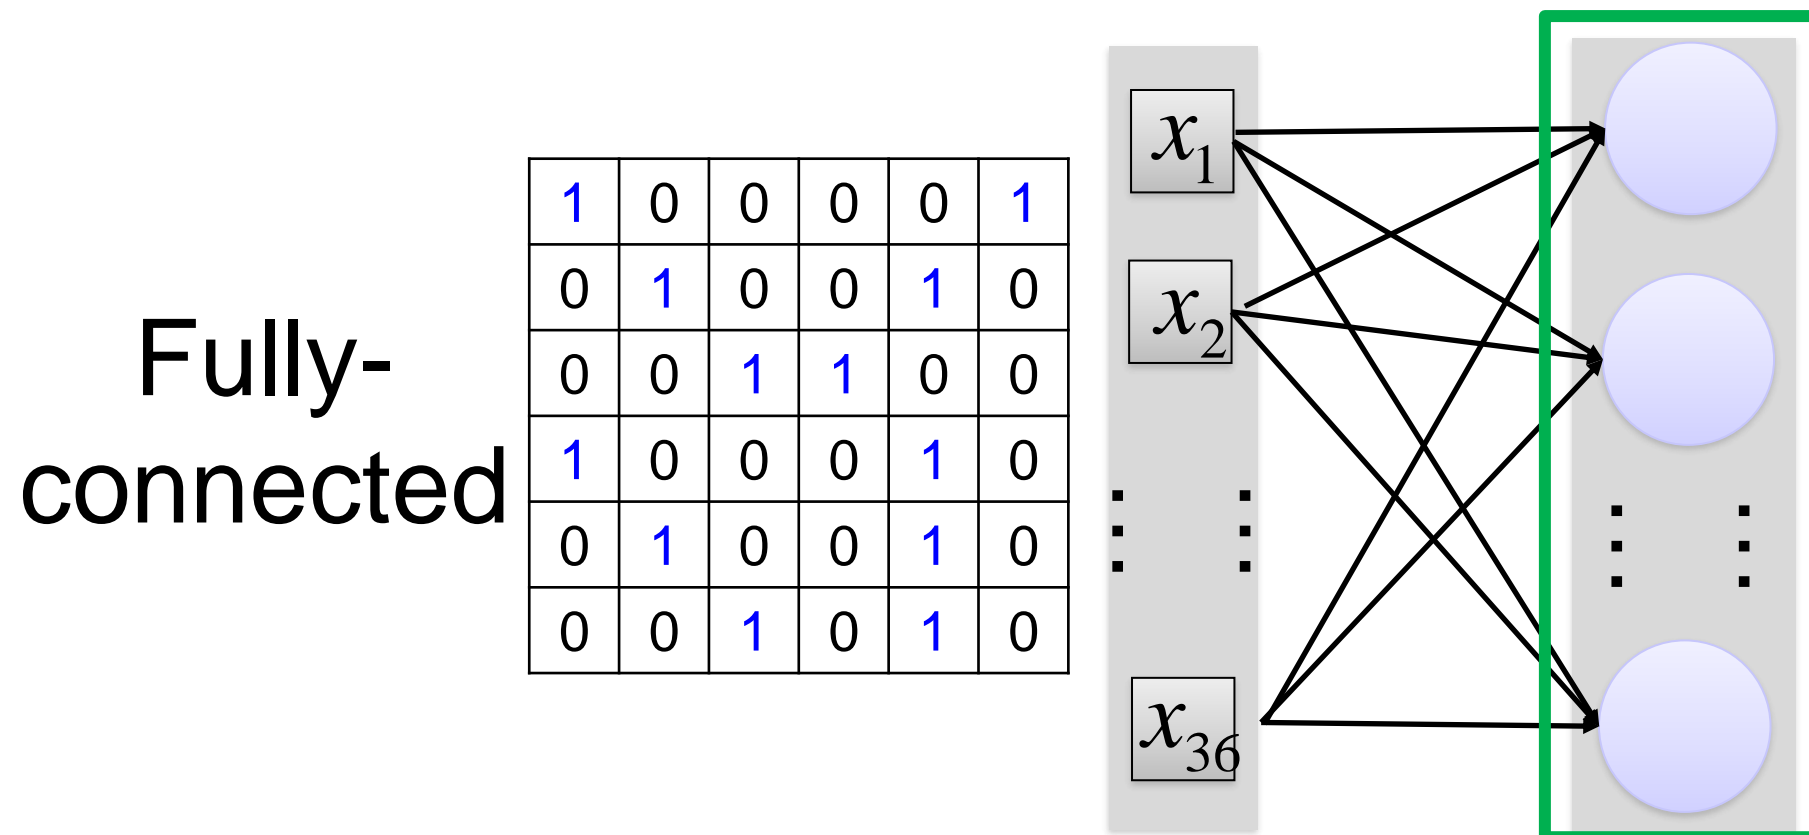



Two 4 x 4 images
Forming 2 x 4 x 4 matrix

Color image: RGB 3 channels




Convolution v.s. Fully Connected



Conventional
Fully Connected
layers
(FC layers)


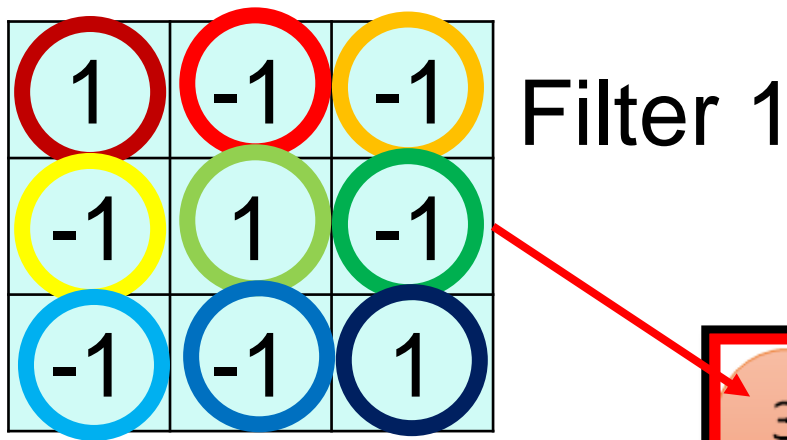
1	0	0	0	0	1
0	1	0	0	1	0
0	0	1	1	0	0
1	0	0	0	1	0
0	1	0	0	1	0
0	0	1	0	1	0

6 x 6 image




features


1st hidden layer



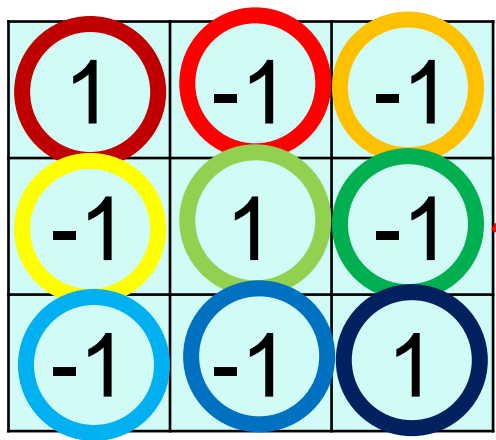
6 x 6 image




fewer parameters!




Only connect to 9 inputs, not fully connected



Filter 1




6 x 6 image



Fewer parameters



Even fewer parameters



Shared weights

Ex. constrained to be identical

An example classifier using CNNs



Can repeat many times

A red bracket on the right side of the diagram groups the first 'Convolution' and 'Max Pooling' boxes, with the text 'Can repeat many times' next to it.

Flattened

A cyan rectangular box containing the text 'Flattened'. A thick black arrow points from the bottom 'Max Pooling' box to this box, and another thick black arrow points from this box to the 'Fully Connected Feedforward network'.

Max Pooling

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

Filter 1

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

Filter 2

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

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

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





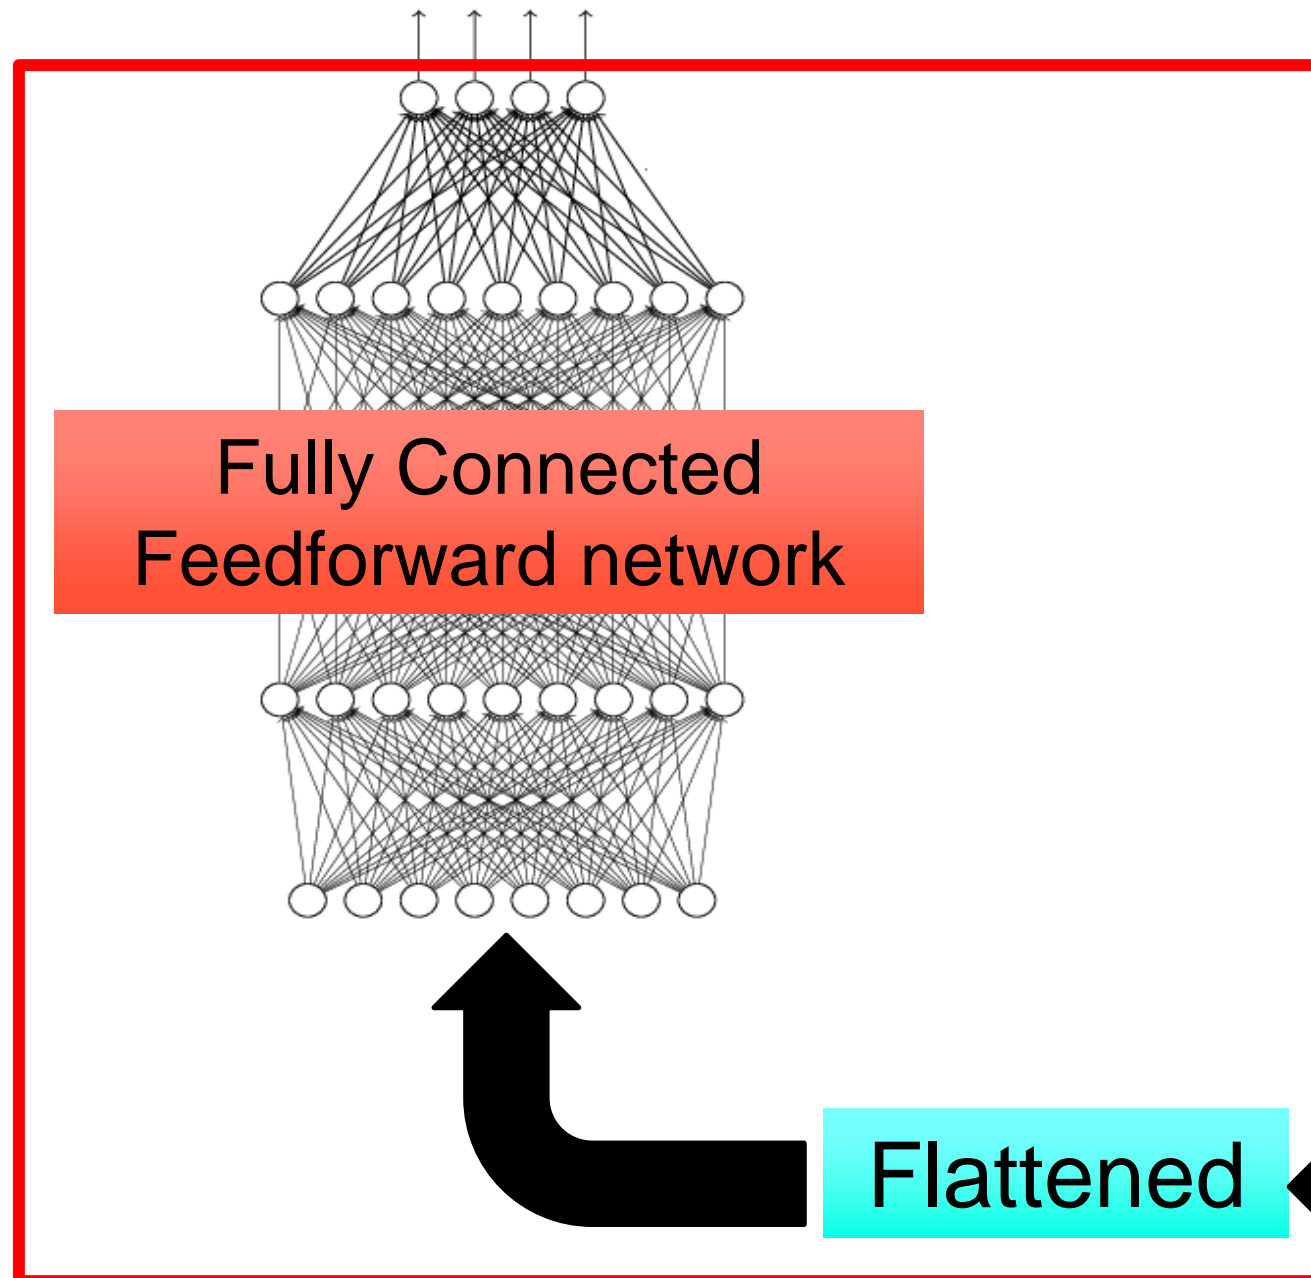
Smaller than the original image
The number of channels is the number of filters



Can repeat many times

Example network

cat dog



A new image

A green rectangular box containing the text 'A new image' in white, positioned to the right of the 'Max Pooling' box.



A new image

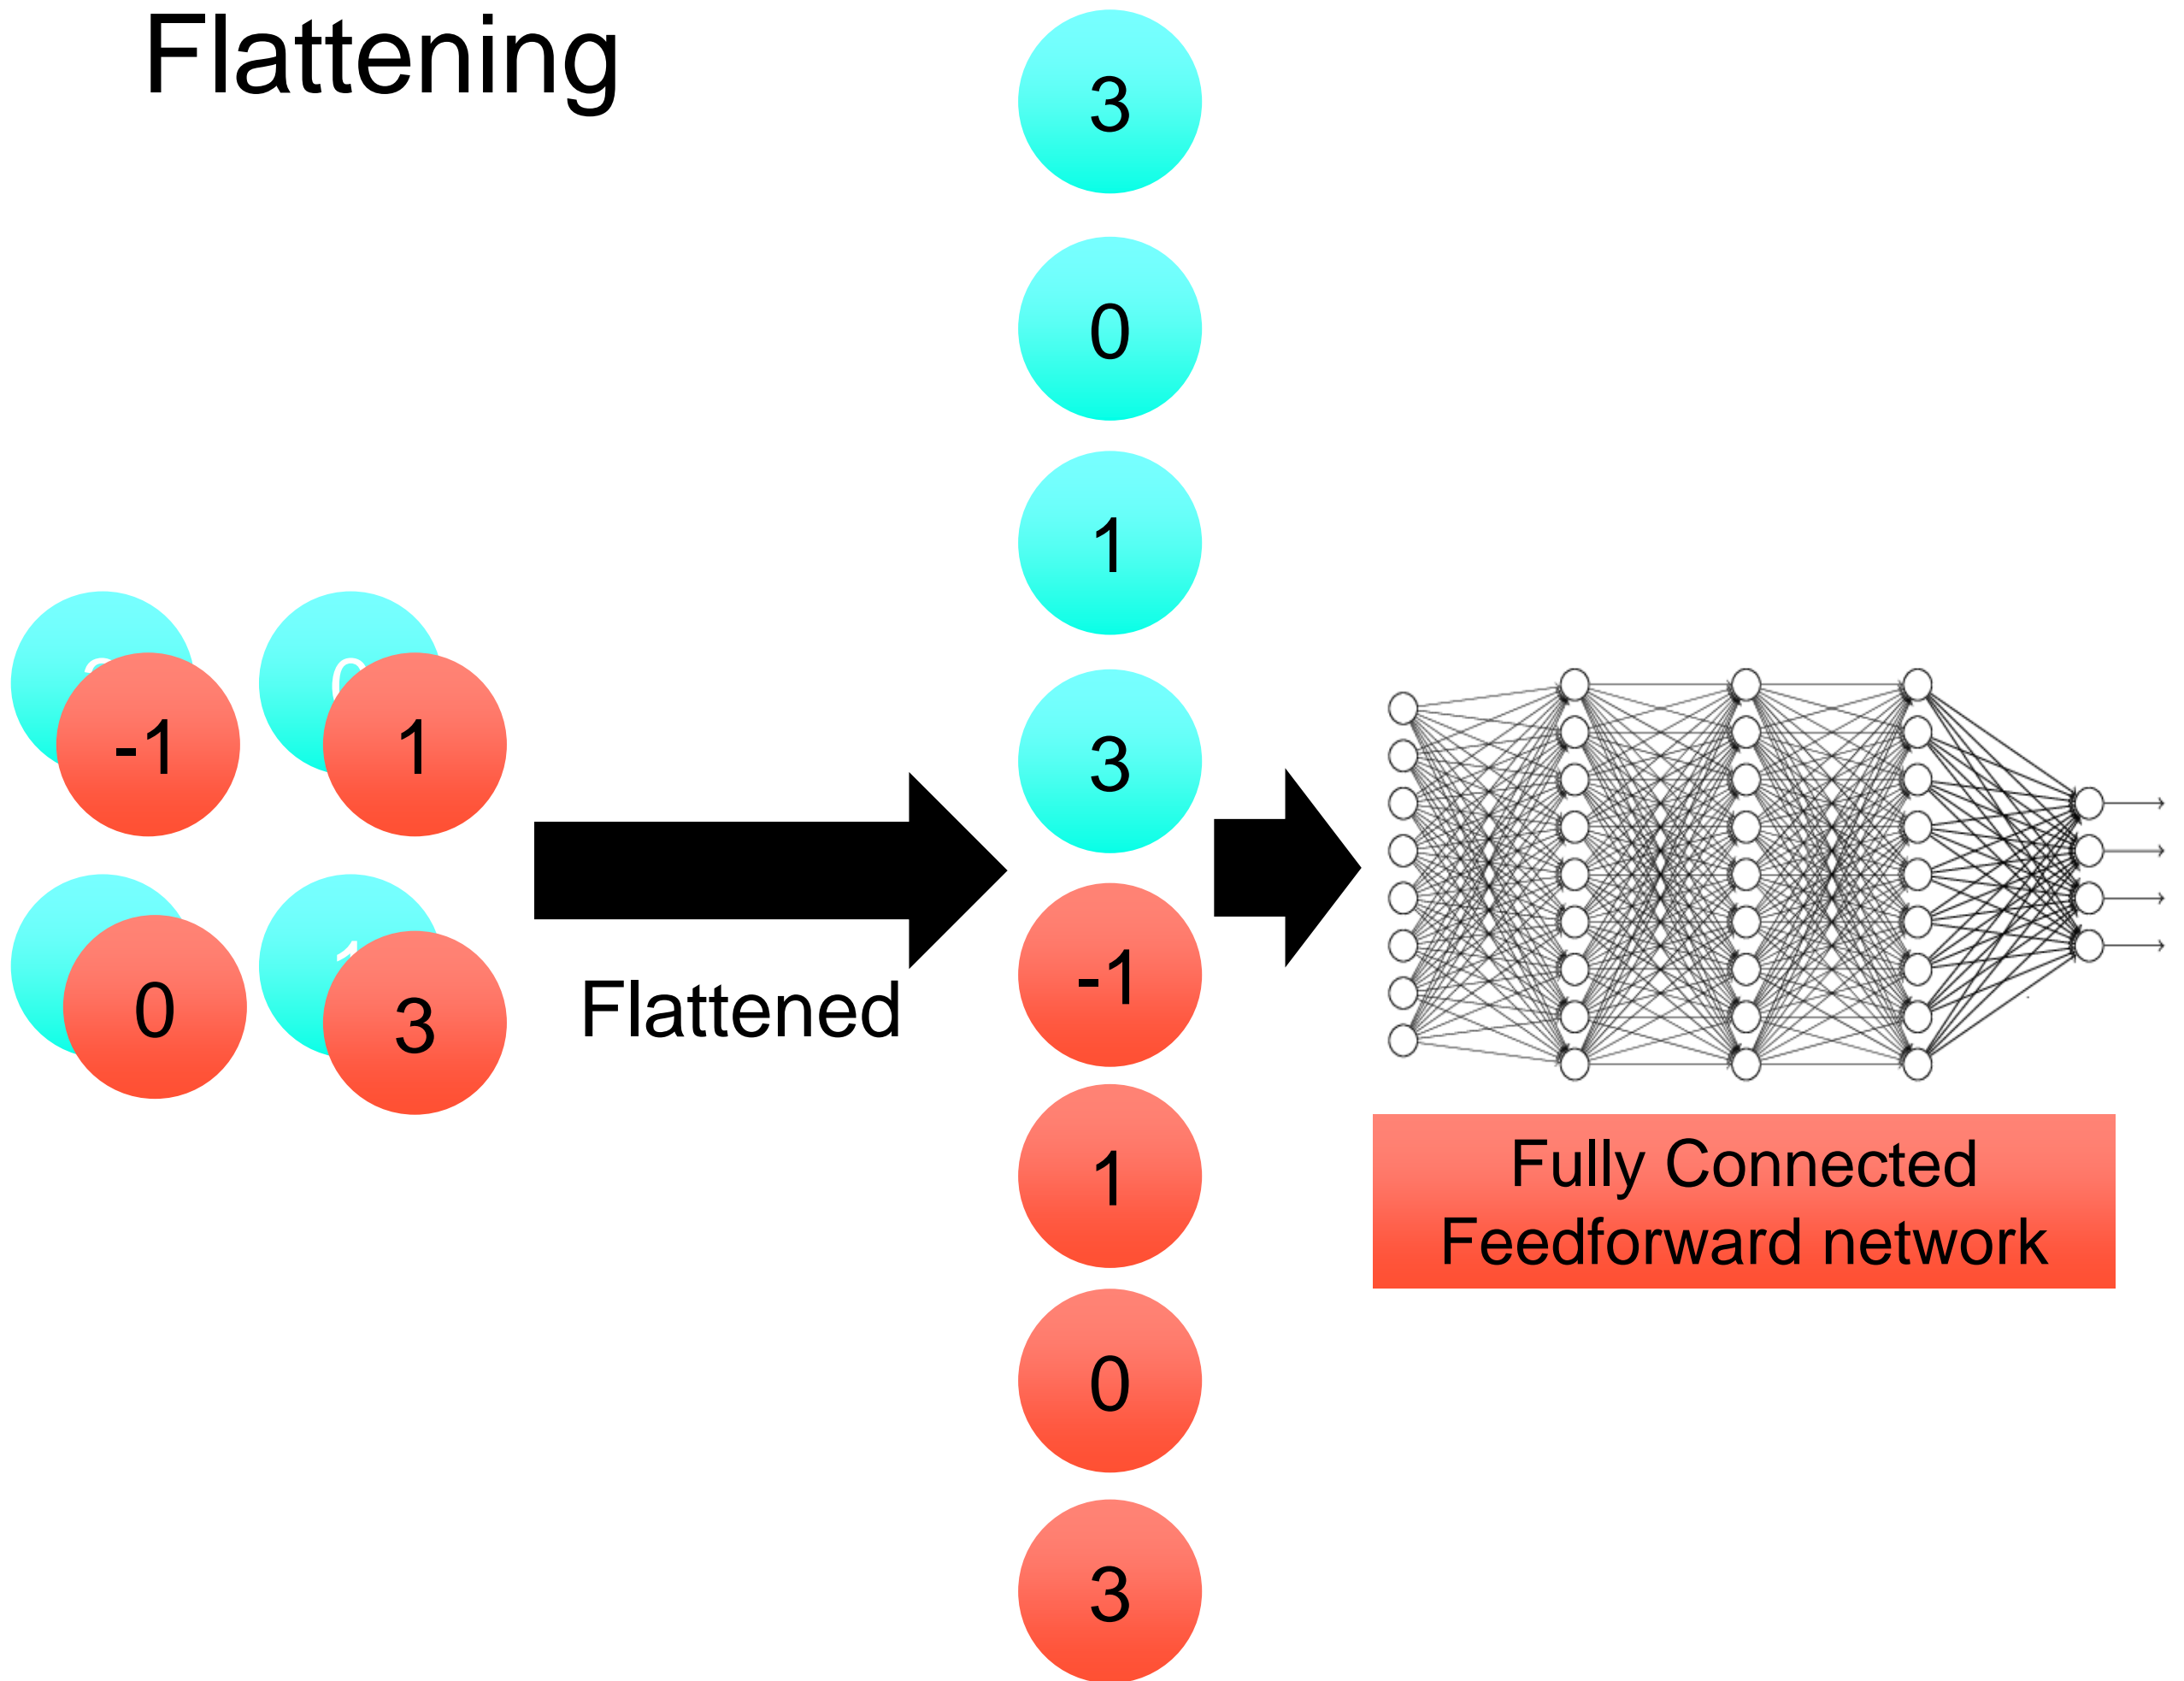
A green rectangular box containing the text 'A new image' in white, positioned to the right of the 'Max Pooling' box.



Flattened

A cyan rectangular box containing the text 'Flattened' in black, positioned to the left of the 'A new image' box.

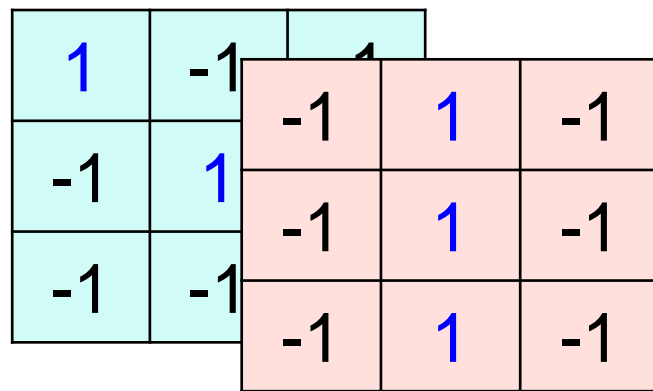
Flattening



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)) )
```





There are **25 3x3** filters.

Input_shape = (28 , 28 , 1)

28 x 28 pixels

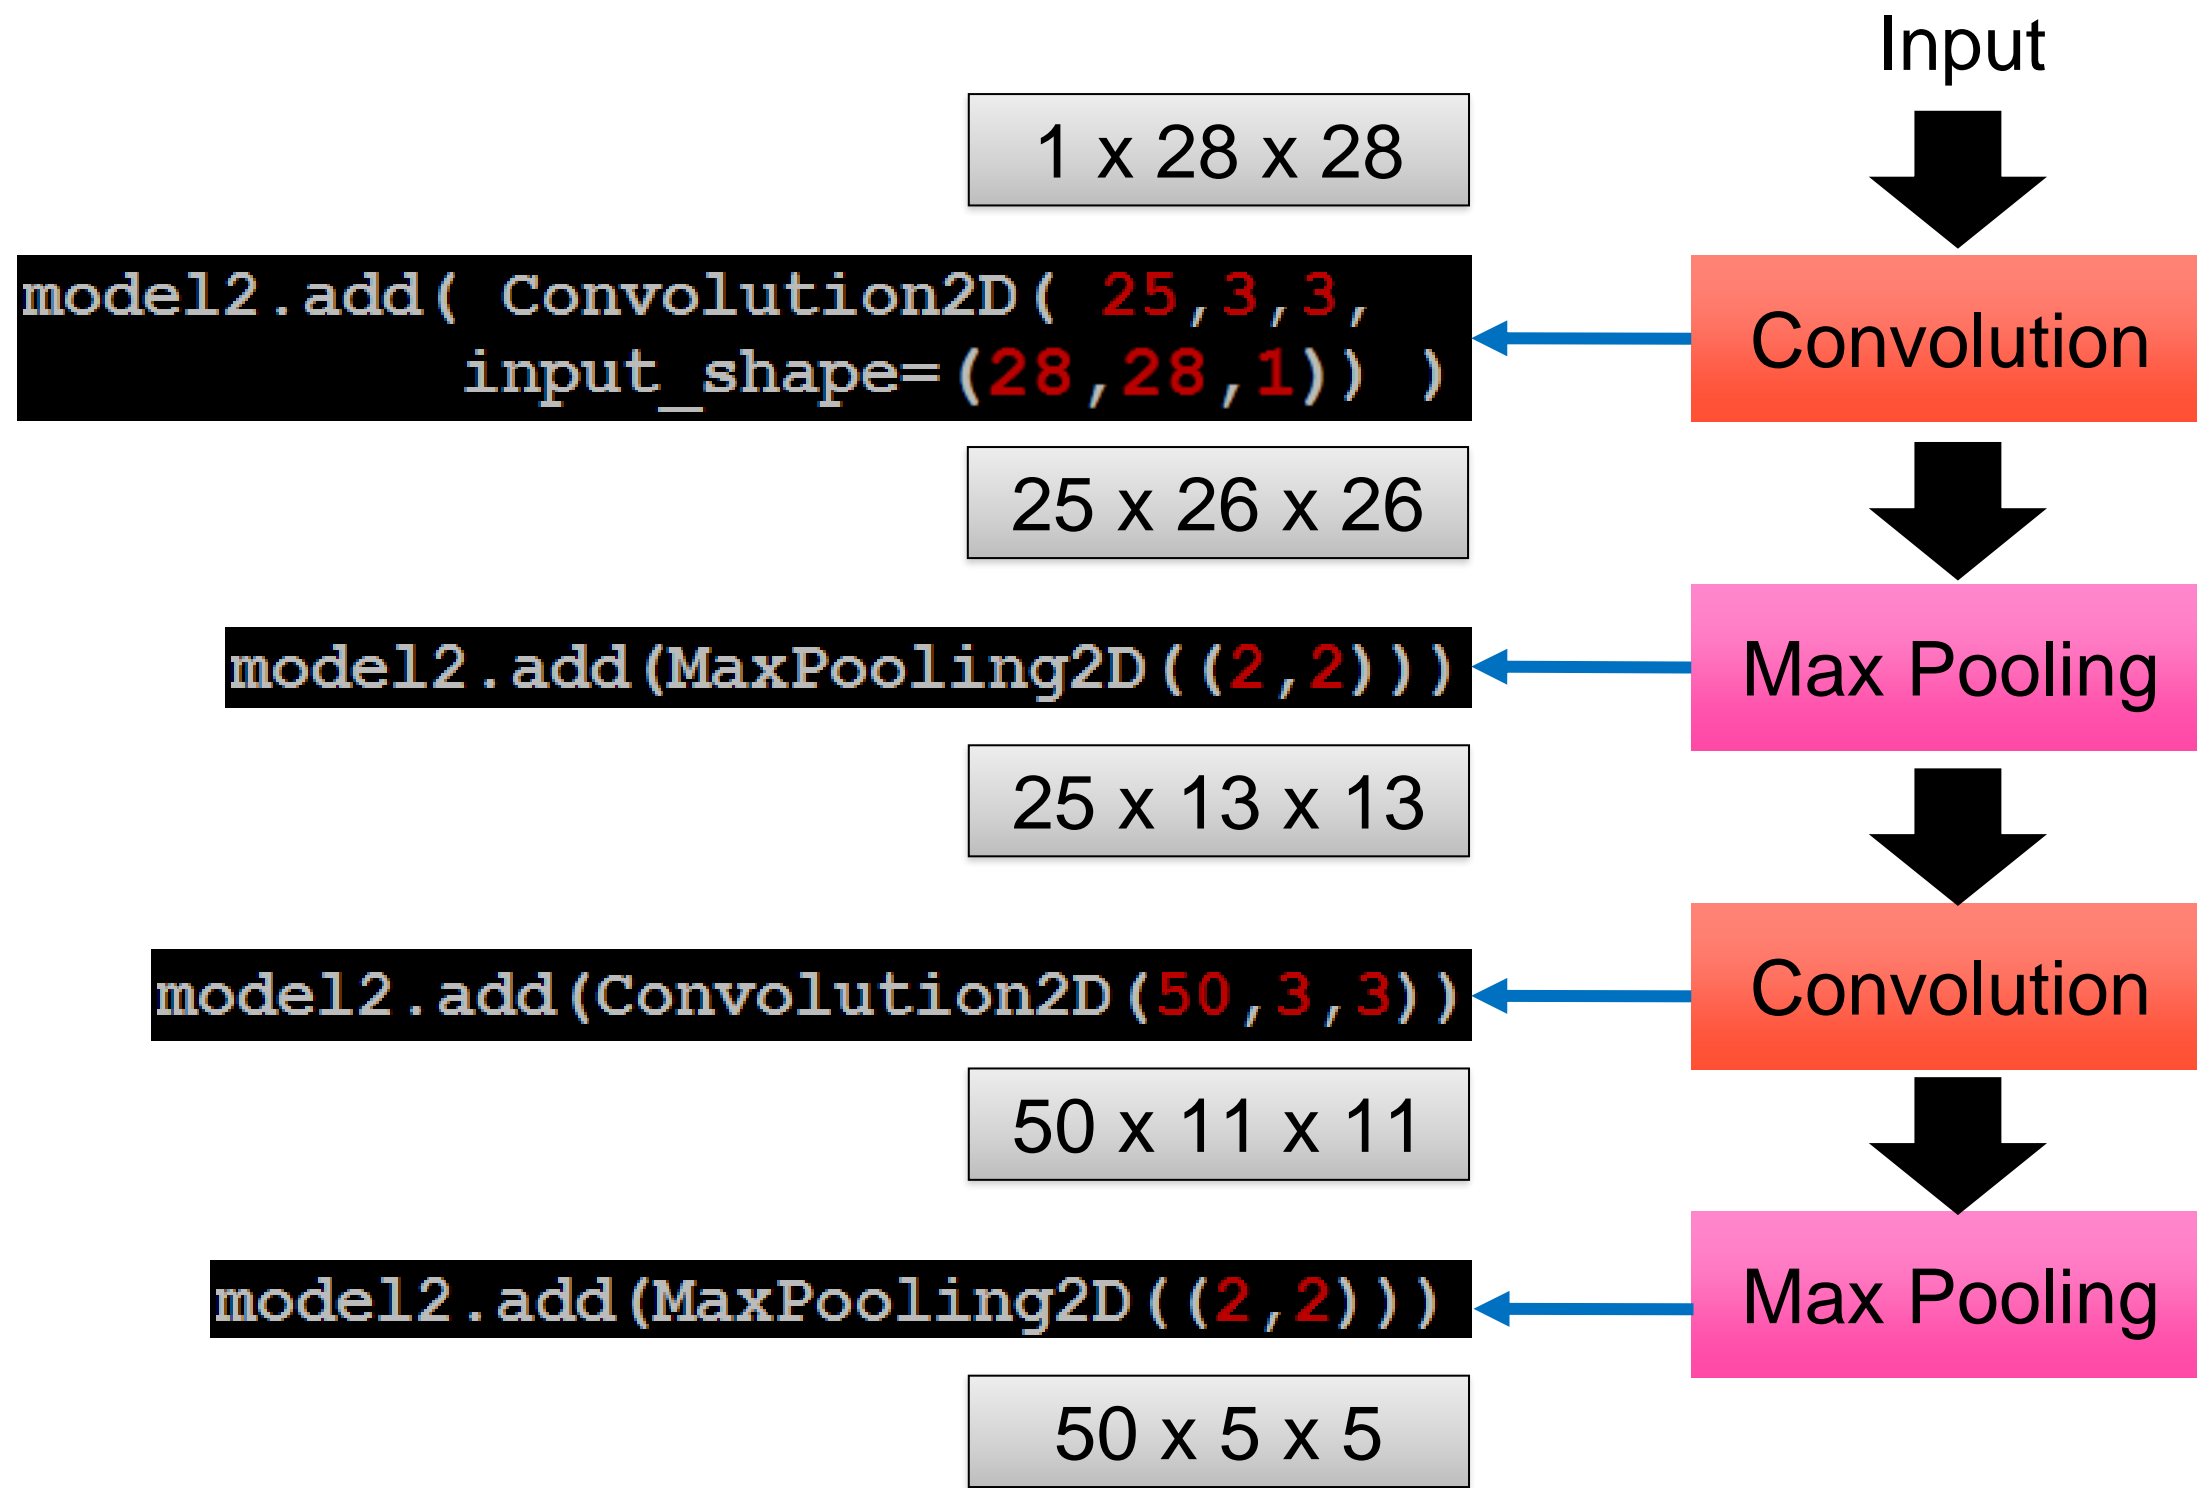
1: black/white, 3: RGB

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



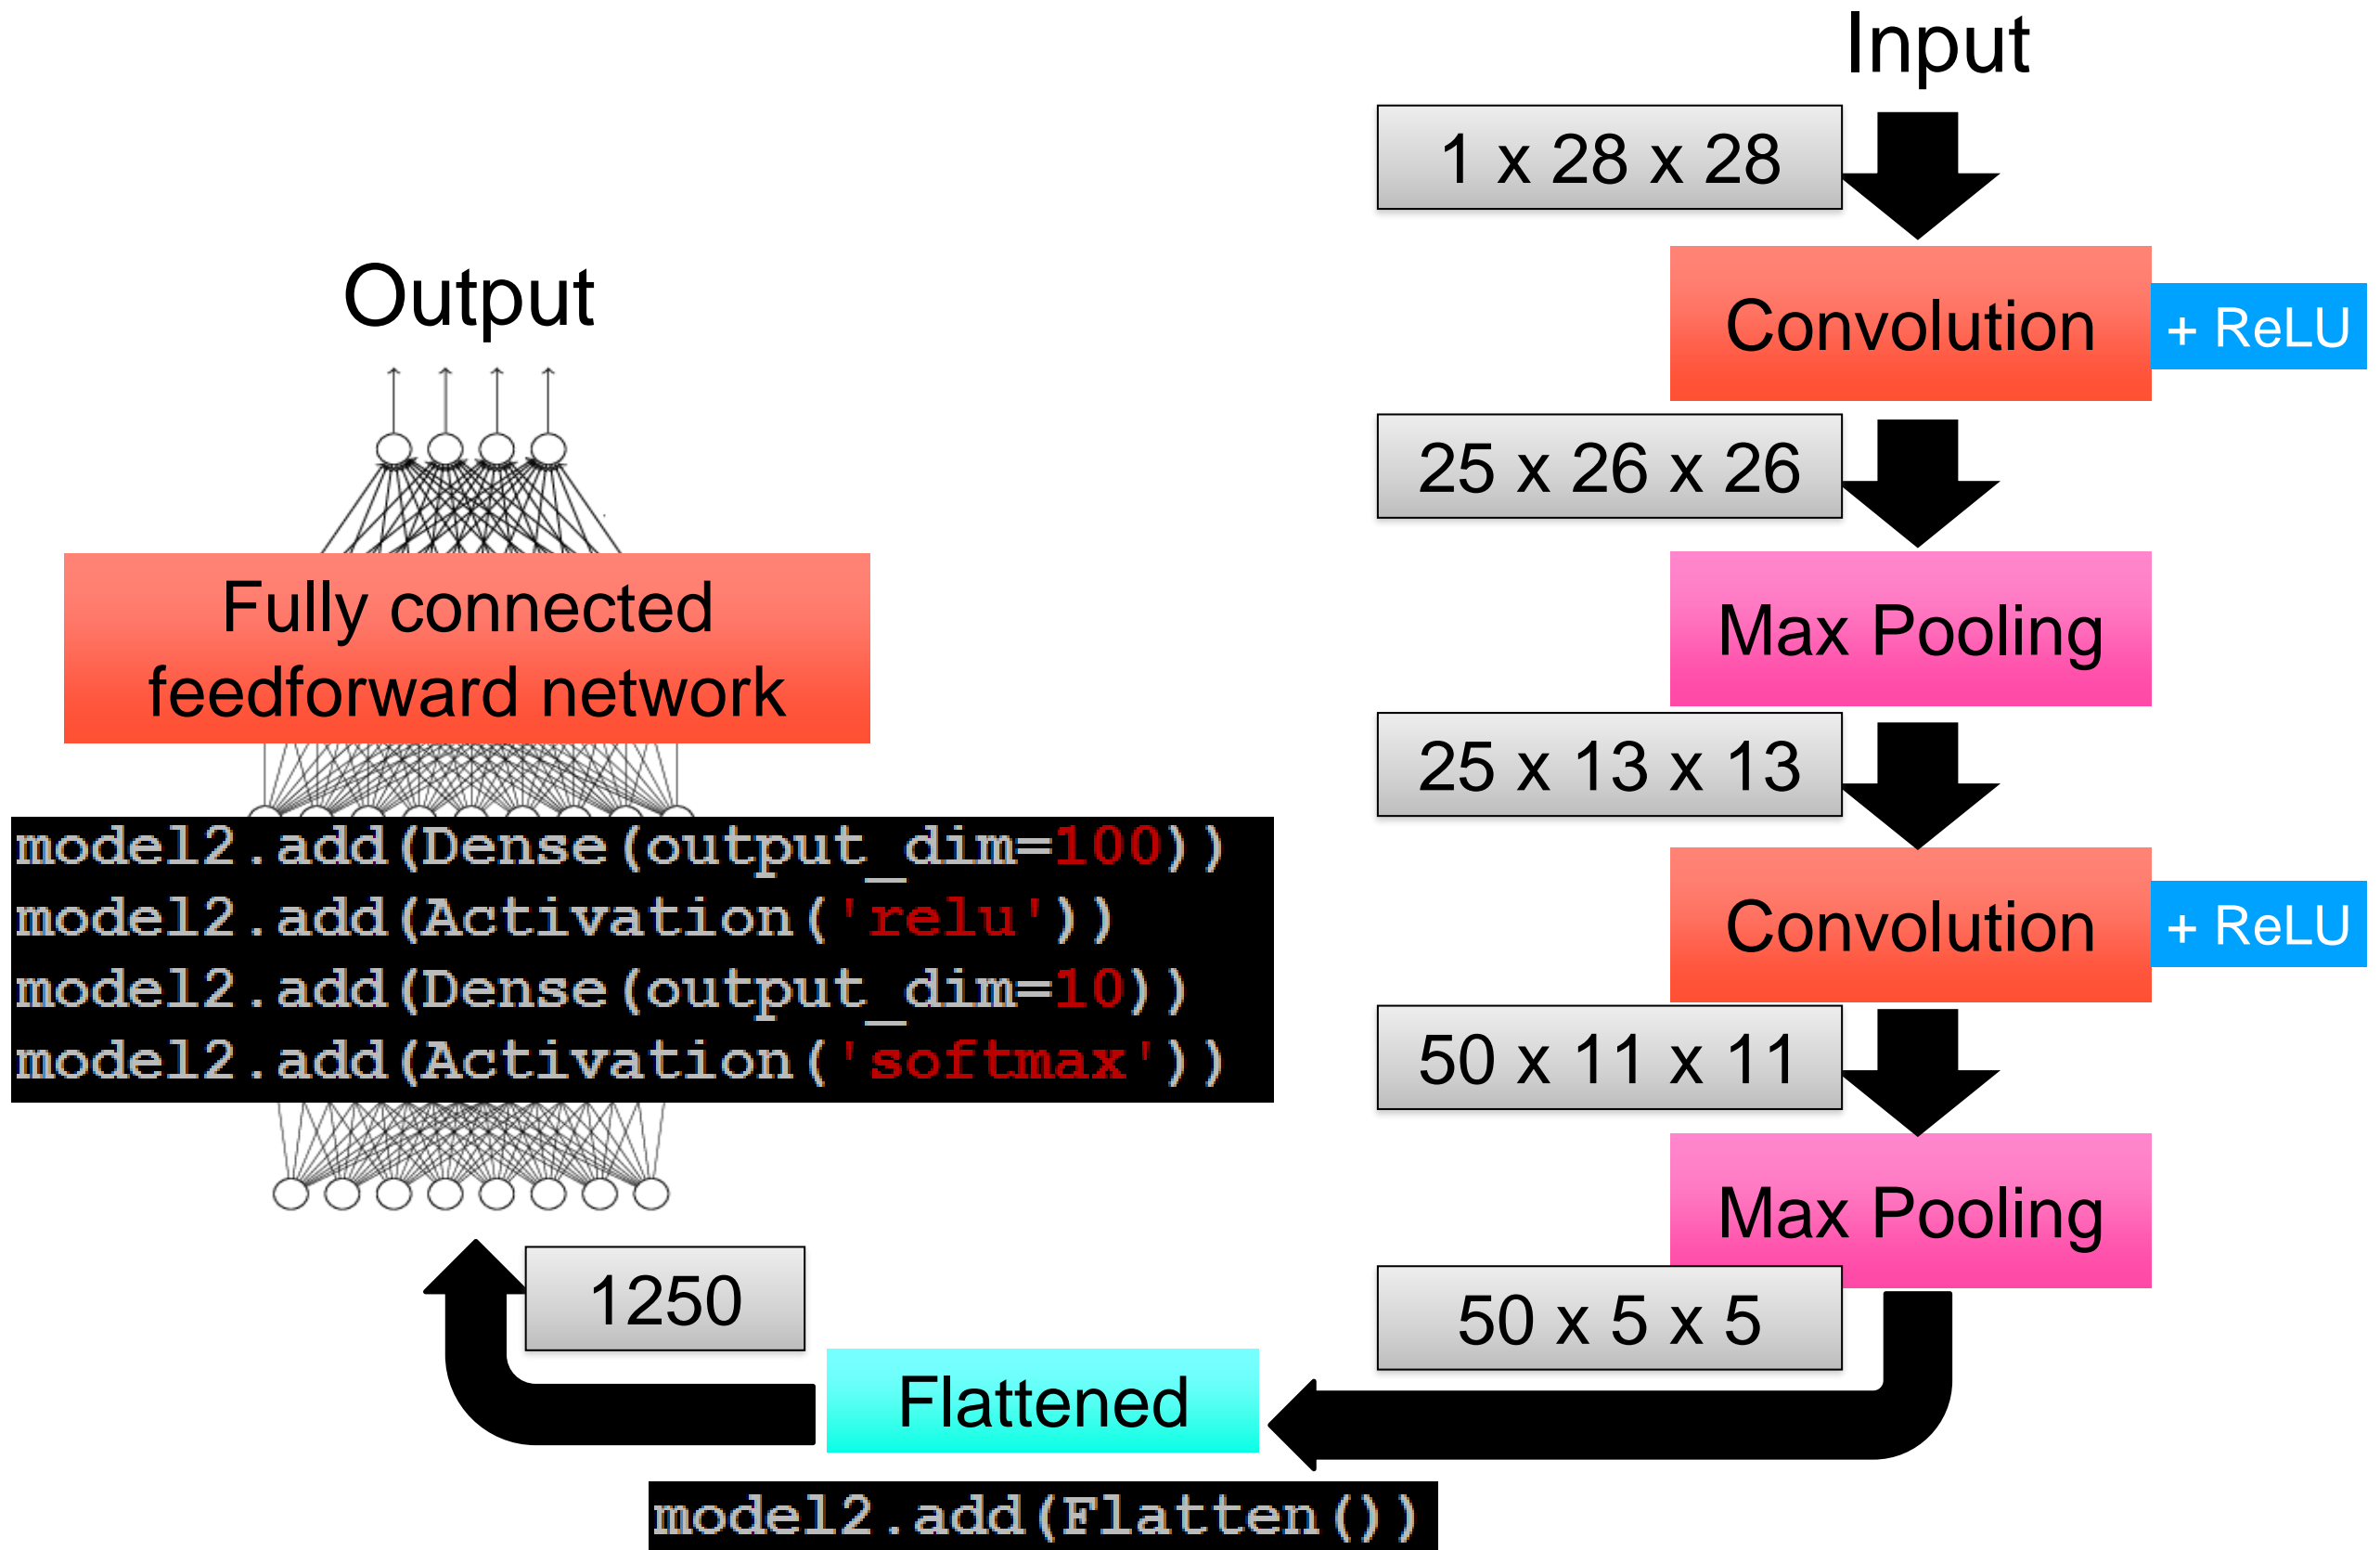
CNN in Keras

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)*



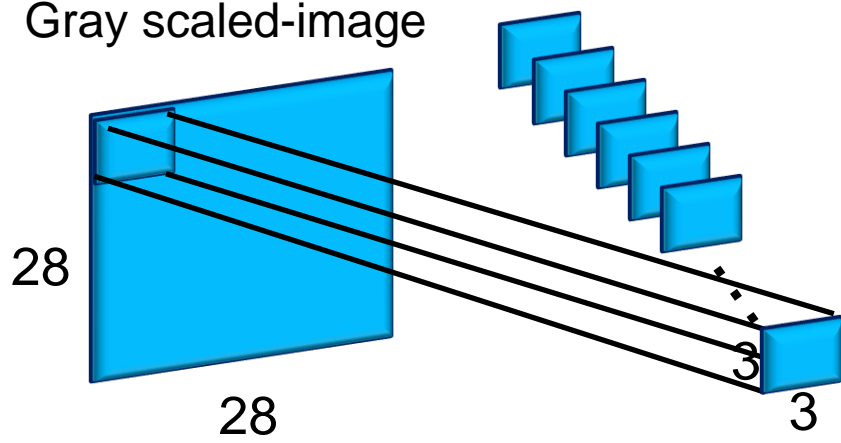
Number of Parameters

25X3X3+25 parameters


25 filters - Conv1

25: 3X3


Gray scaled-image




Convolved result for 25 filters



Result after Max Pooling





Result after Max Pooling



50: 3X3X25

50 filters - Conv2

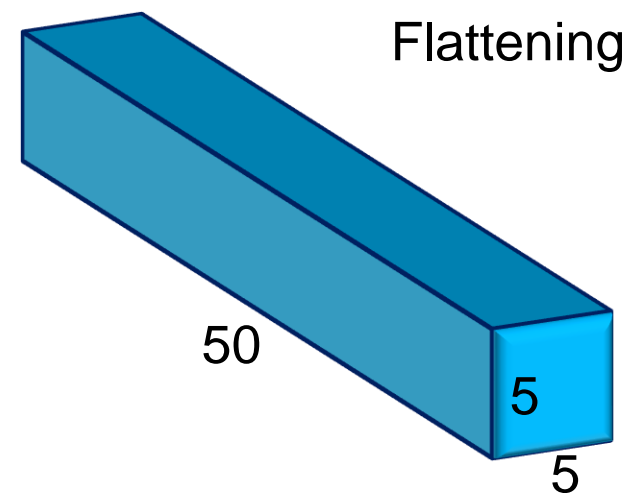
Convolved result for 50 filters




50X3X3X25+50 parameters

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

100 neurons



Flattening



FC layer
#parameters=1250*10

FC layer
#parameters=10*100

→ Prediction

10 CNN Architecture