

# CSCSE 636 Neural Networks (Deep Learning)

Lecture 7: Deep Learning for Computer Vision

Anxiao (Andrew) Jiang

Based on the interesting lecture of Prof. Hung-yi Lee,  
[https://www.youtube.com/watch?v=FrKWiRv254g&list=PLJV\\_el3uVTsPy9oCRY30oBPNLCo89yu49&index=19](https://www.youtube.com/watch?v=FrKWiRv254g&list=PLJV_el3uVTsPy9oCRY30oBPNLCo89yu49&index=19)


# Convolutional Neural Network (CNN)

# Why CNN for Image?

[Zeiler, M. D., *ECCV 2014*]

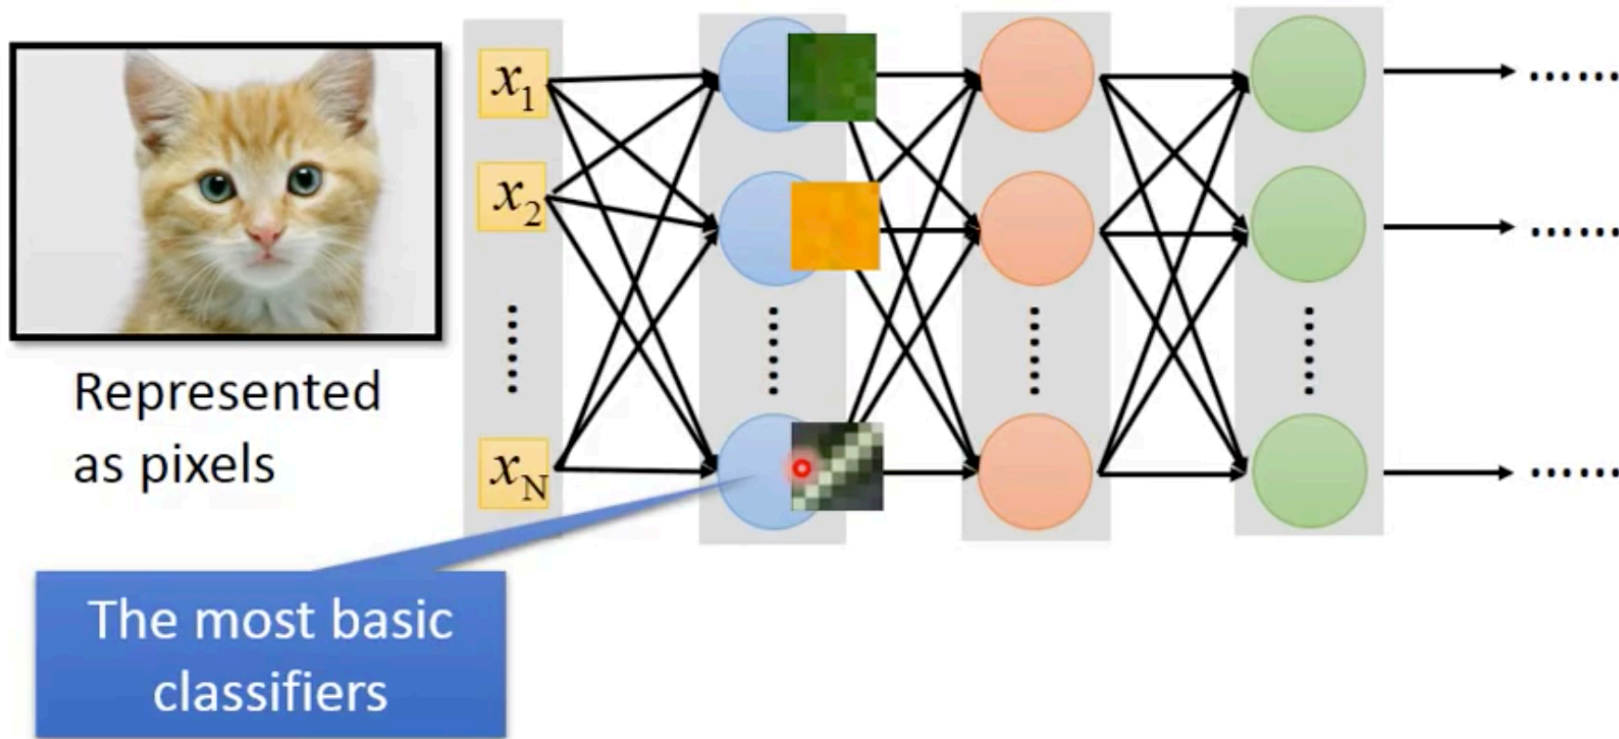


Represented  
as pixels



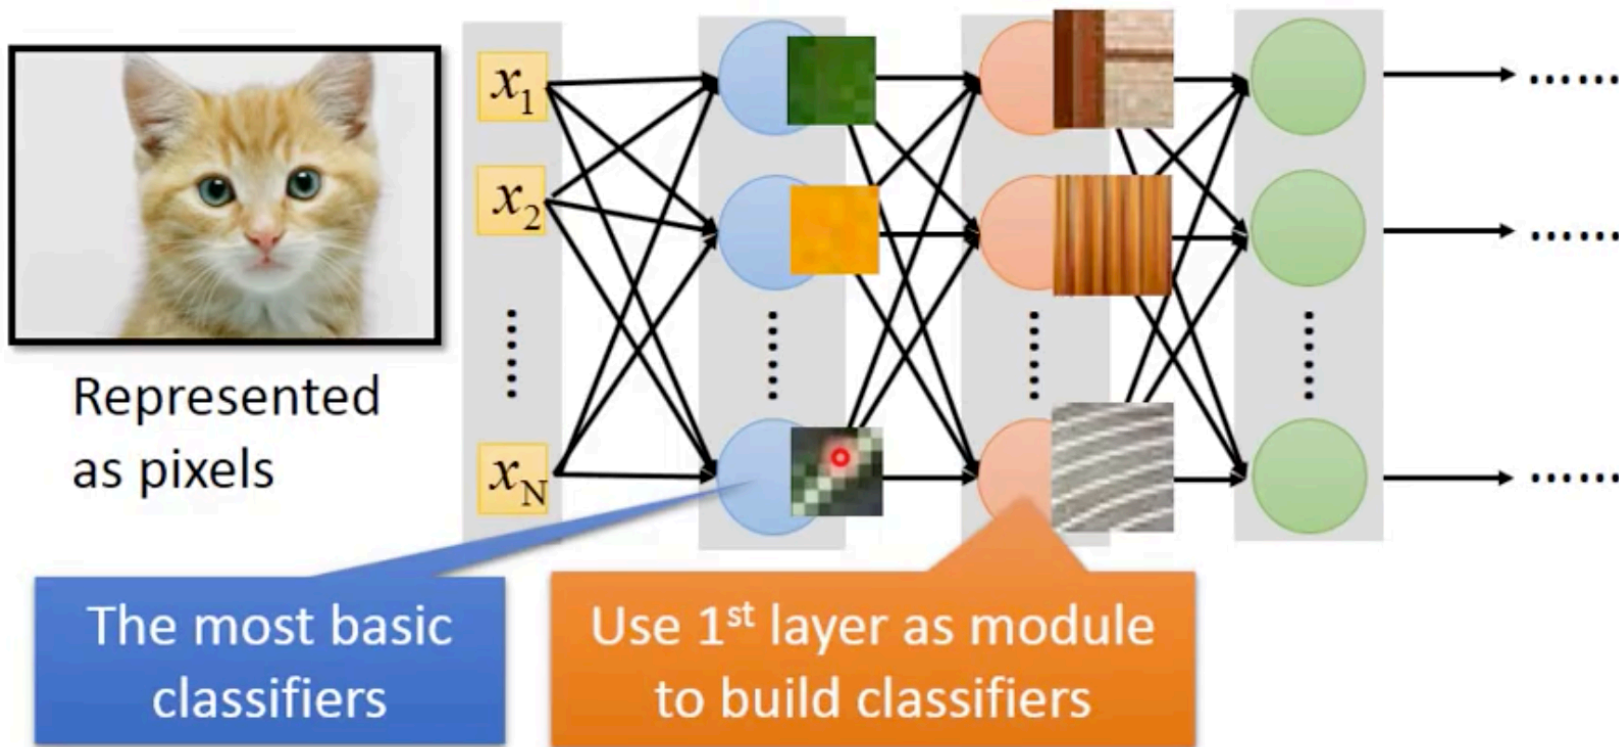
# Why CNN for Image?

[Zeiler, M. D., *ECCV 2014*]




# Why CNN for Image?

[Zeiler, M. D., *ECCV 2014*]



# Why CNN for Image?

[Zeiler, M. D., *ECCV 2014*]



Too many weights in a dense network!

# Why CNN for Image

- Some patterns are much smaller than the whole image.

A neuron does not have to see the whole image to discover the pattern.



有沒有某一個 pattern 出現

# Why CNN for Image

- Some patterns are much smaller than the whole image

A neuron does not have to see the whole image to discover the pattern.

Connecting to small region with less parameters





# Why CNN for Image

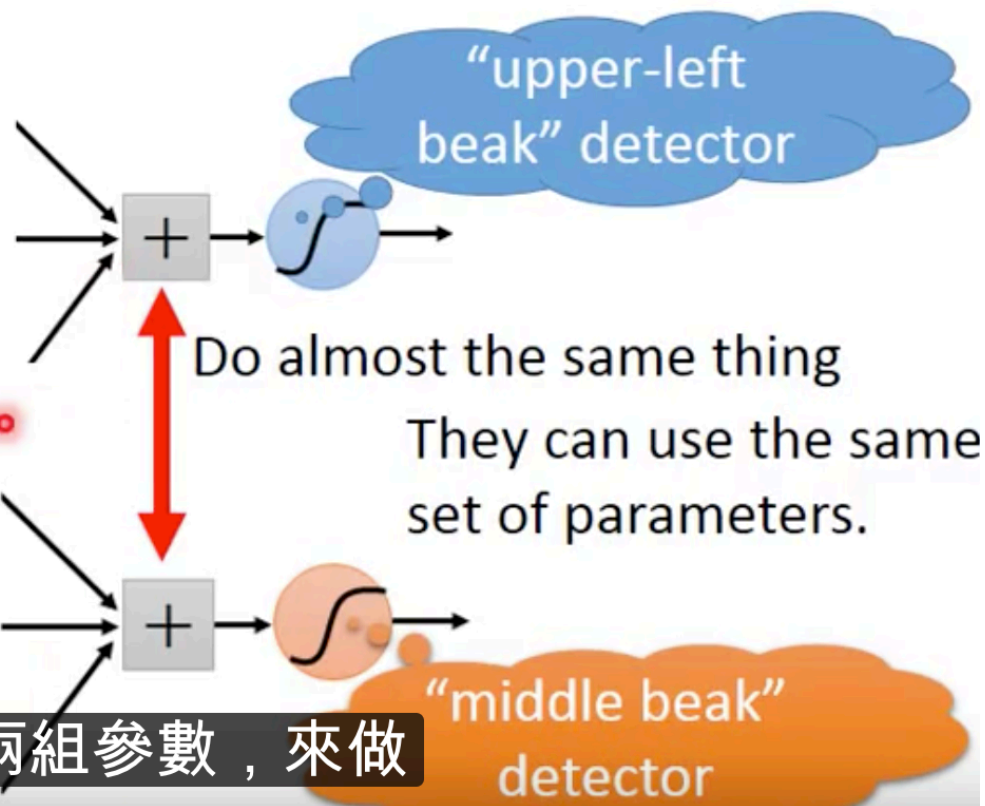
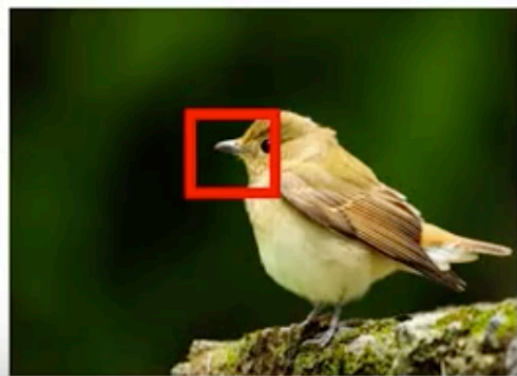
- The same patterns appear in different regions.



同樣的這個 pattern

# Why CNN for Image

- The same patterns appear in different regions.



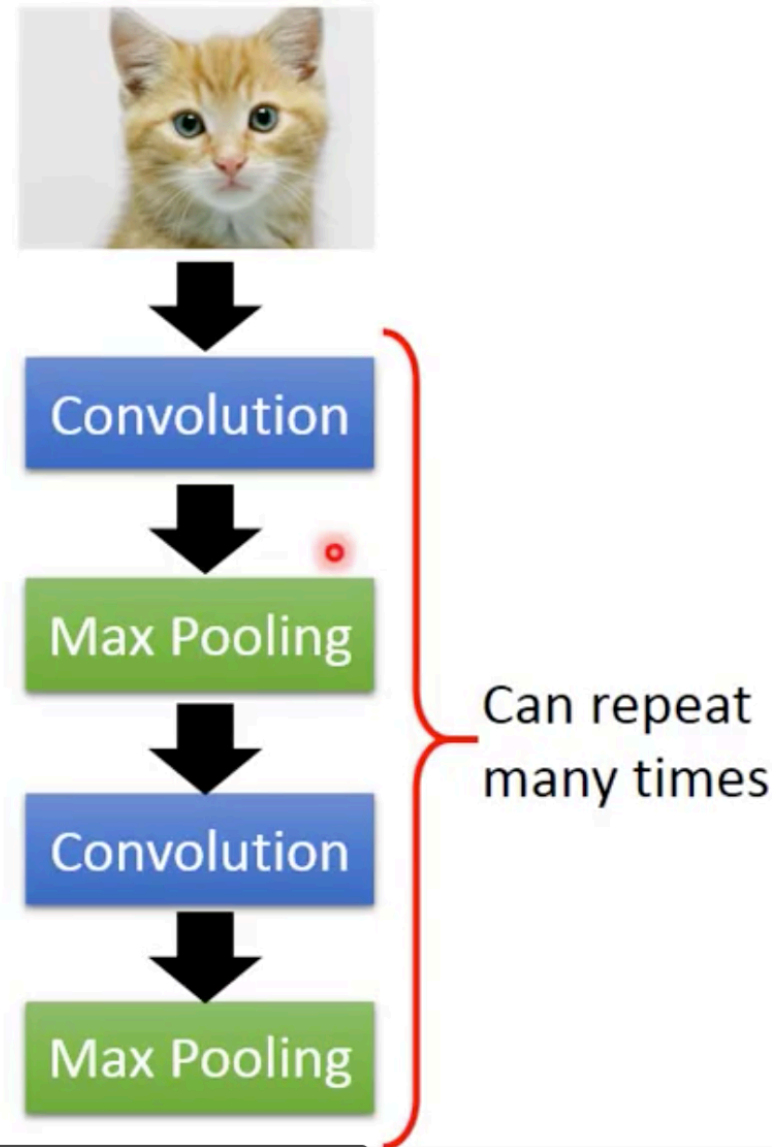
# Why CNN for Image

- Subsampling the pixels will not change the object



We can subsample the pixels to make image smaller

# The whole CNN



# The whole CNN

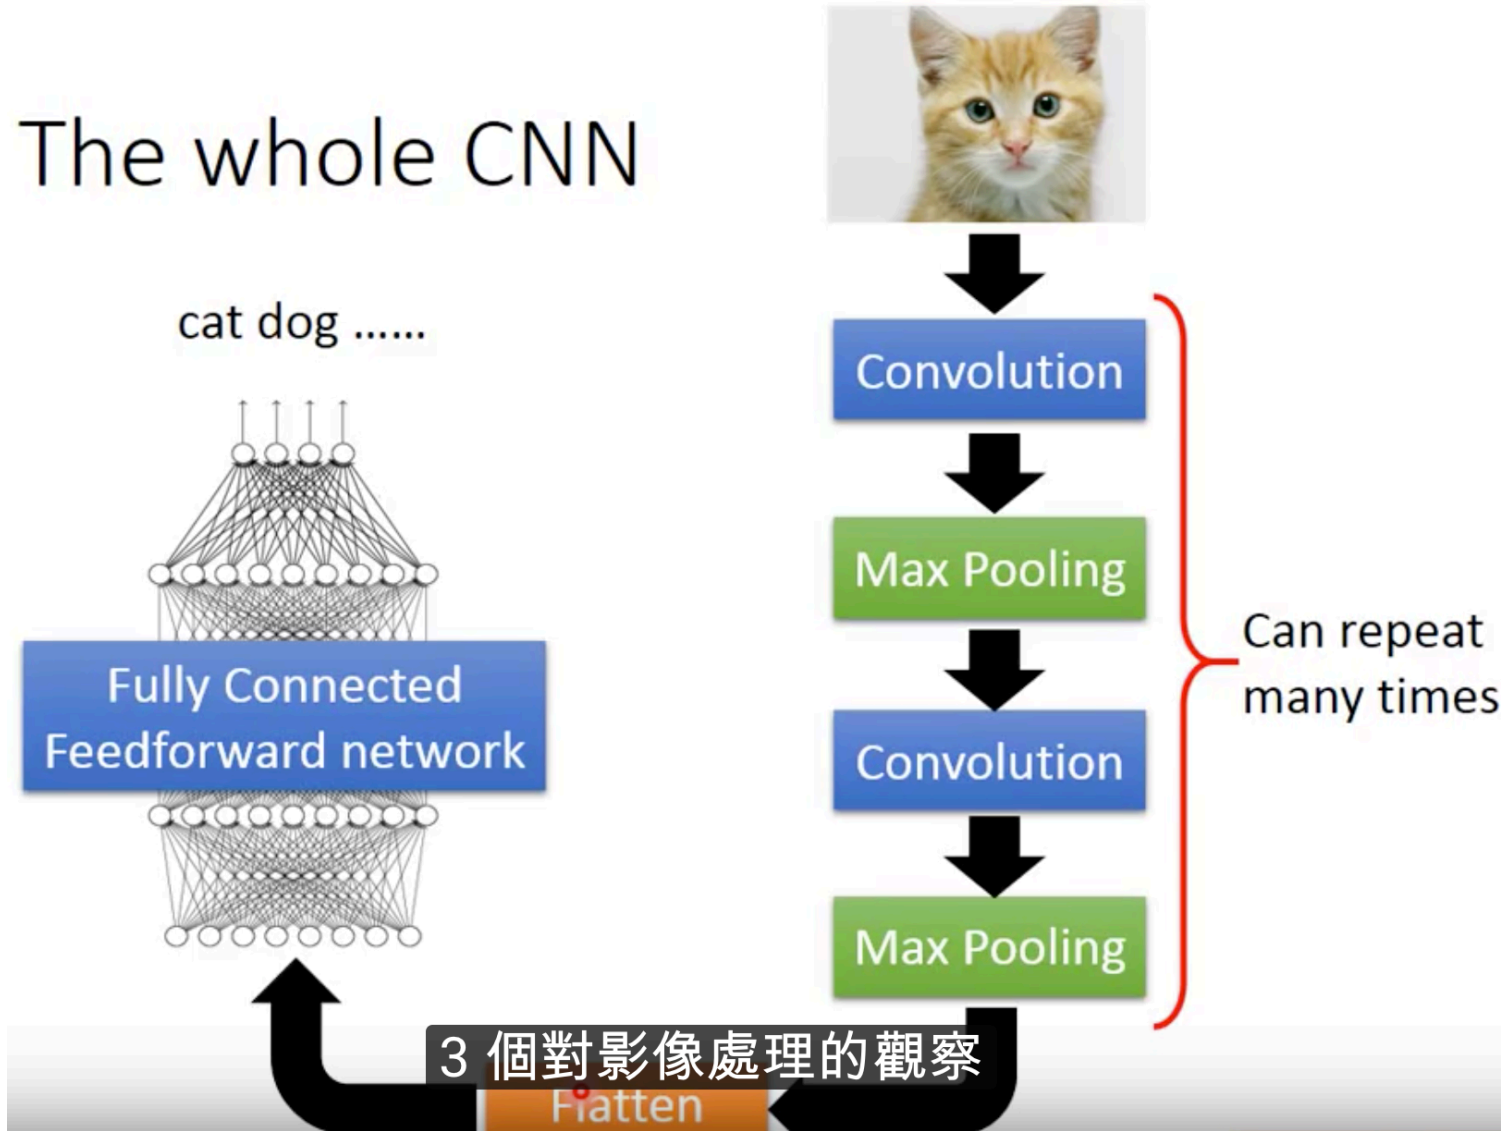


Can repeat many times

你要做一件事情叫 Flatten

Flatten

# The whole CNN



# The whole CNN



## Property 1

- Some patterns are much smaller than the whole image

## Property 2

- The same patterns appear in different regions.

## Property 3

- Subsampling the pixels will not change the object

Convolution

Max Pooling

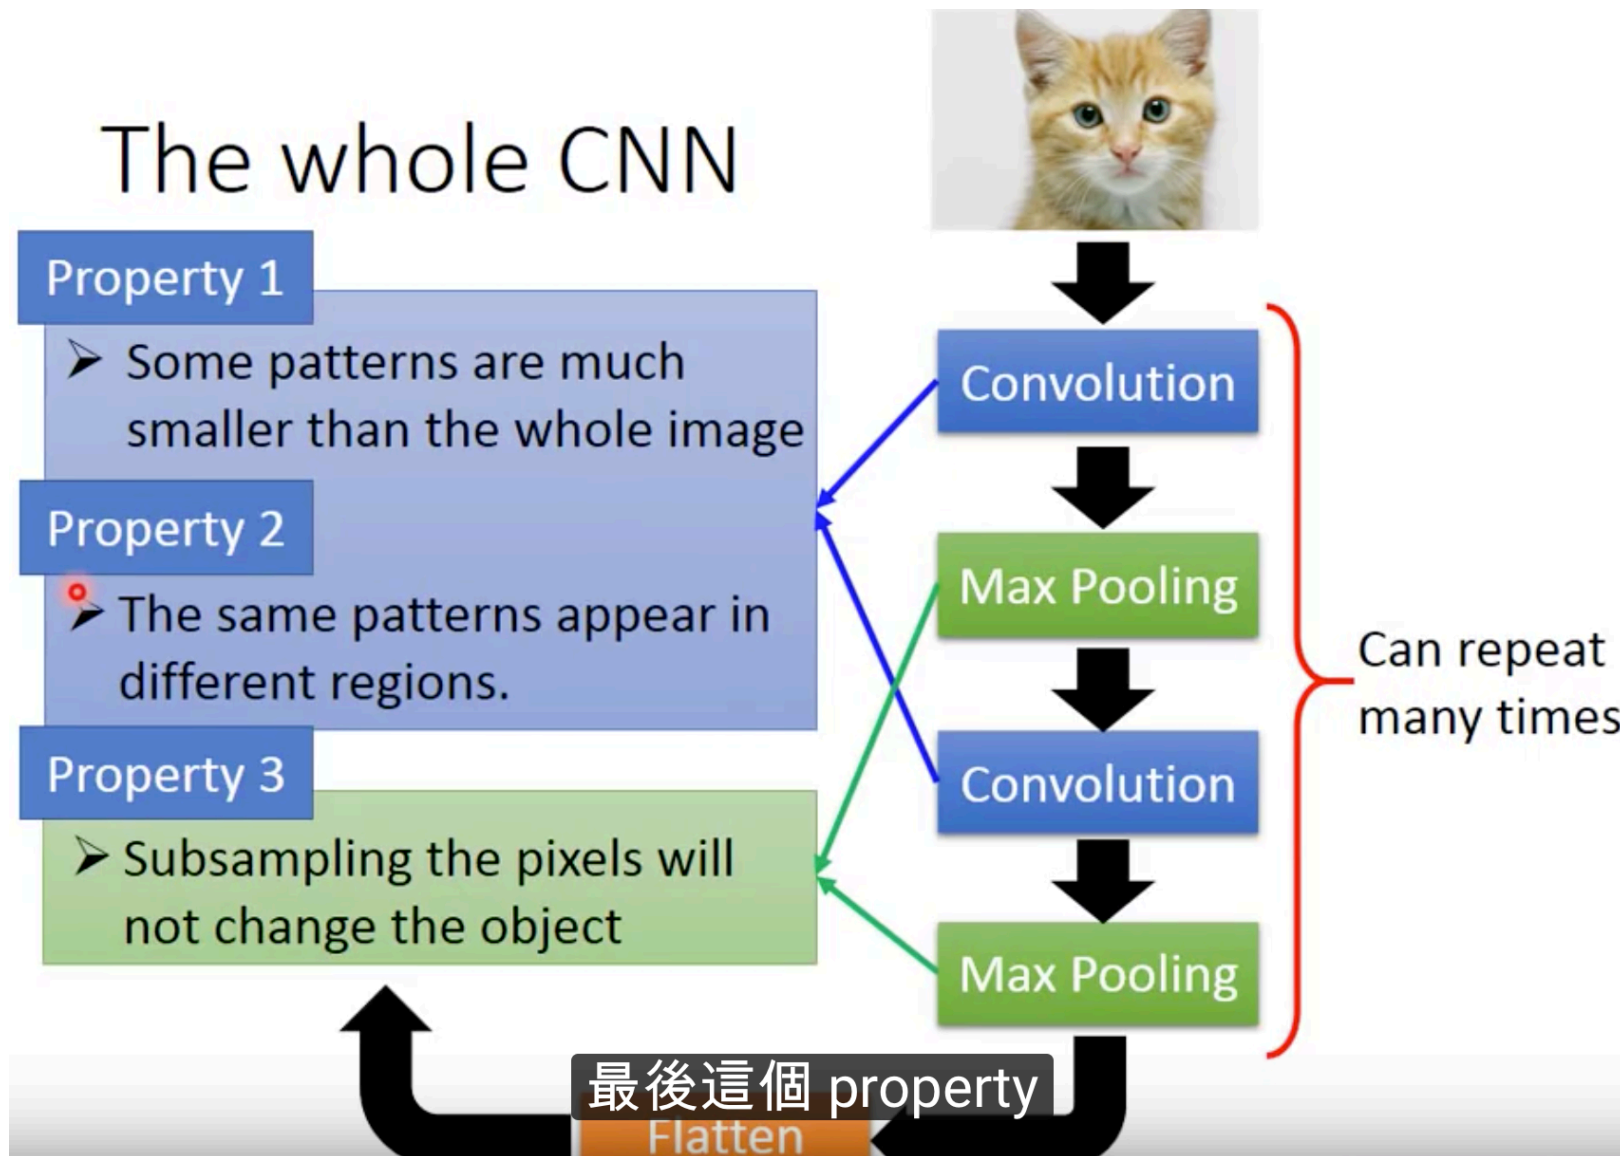
Convolution

Max Pooling

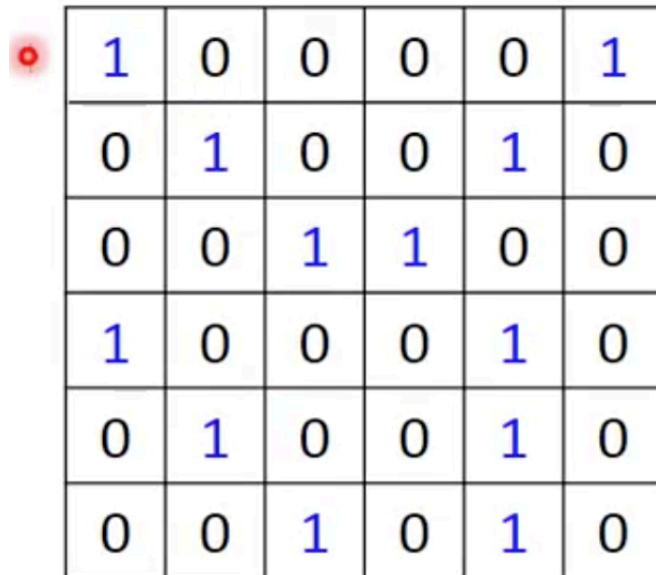
Can repeat many times

最後這個 property

Flatten



# CNN – 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

---



# CNN – 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

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

Filter 1

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

Filter 2

⋮

# CNN – Convolution

A filter corresponds to a set of neurons

Those are the network parameters to be learned.

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

Matrix

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

Filter 2

Matrix

⋮

# CNN – Convolution

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

Filter 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

How does a filter operate?

# CNN – Convolution

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

Filter 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

Do inner product (dot product)

# CNN – Convolution

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

Filter 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

3



# CNN – 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

6 x 6 image

3



# CNN – 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

6 x 6 image



# CNN – 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



We set stride = 1 in the following slides.



# CNN – 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

6 x 6 image



# CNN – 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

6 x 6 image



# CNN – 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

6 x 6 image



# CNN – 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

6 x 6 image



# CNN – 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

6 x 6 image

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

你就得到 -1

# CNN – 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

6 x 6 image

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

出現最大的值

# CNN – Convolution


-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



一打 filter 它會有一大堆的 filter

# CNN – Convolution

-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

Do the same process for every filter

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

比如說 這裡會有一個 filter 2



# CNN – Convolution

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

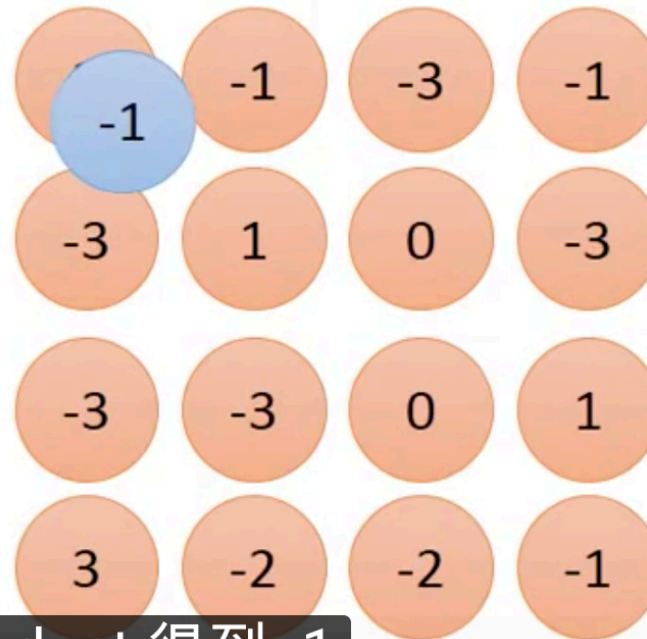
Filter 2

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

Do the same process for every filter



再做 inner product 得到 -1

# CNN – Convolution

-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

Do the same process for every filter



再挪動一次，得到 1 就這樣，以此類推

# CNN – Convolution

-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

Do the same process for every filter



# CNN – Convolution

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

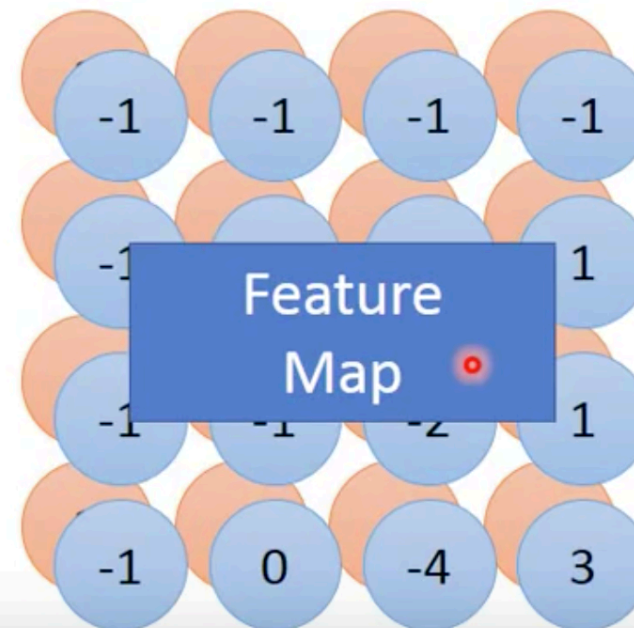
Filter 2

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


Do the same process for every filter



4 x 4 image


# CNN – Colorful image

Colorful image

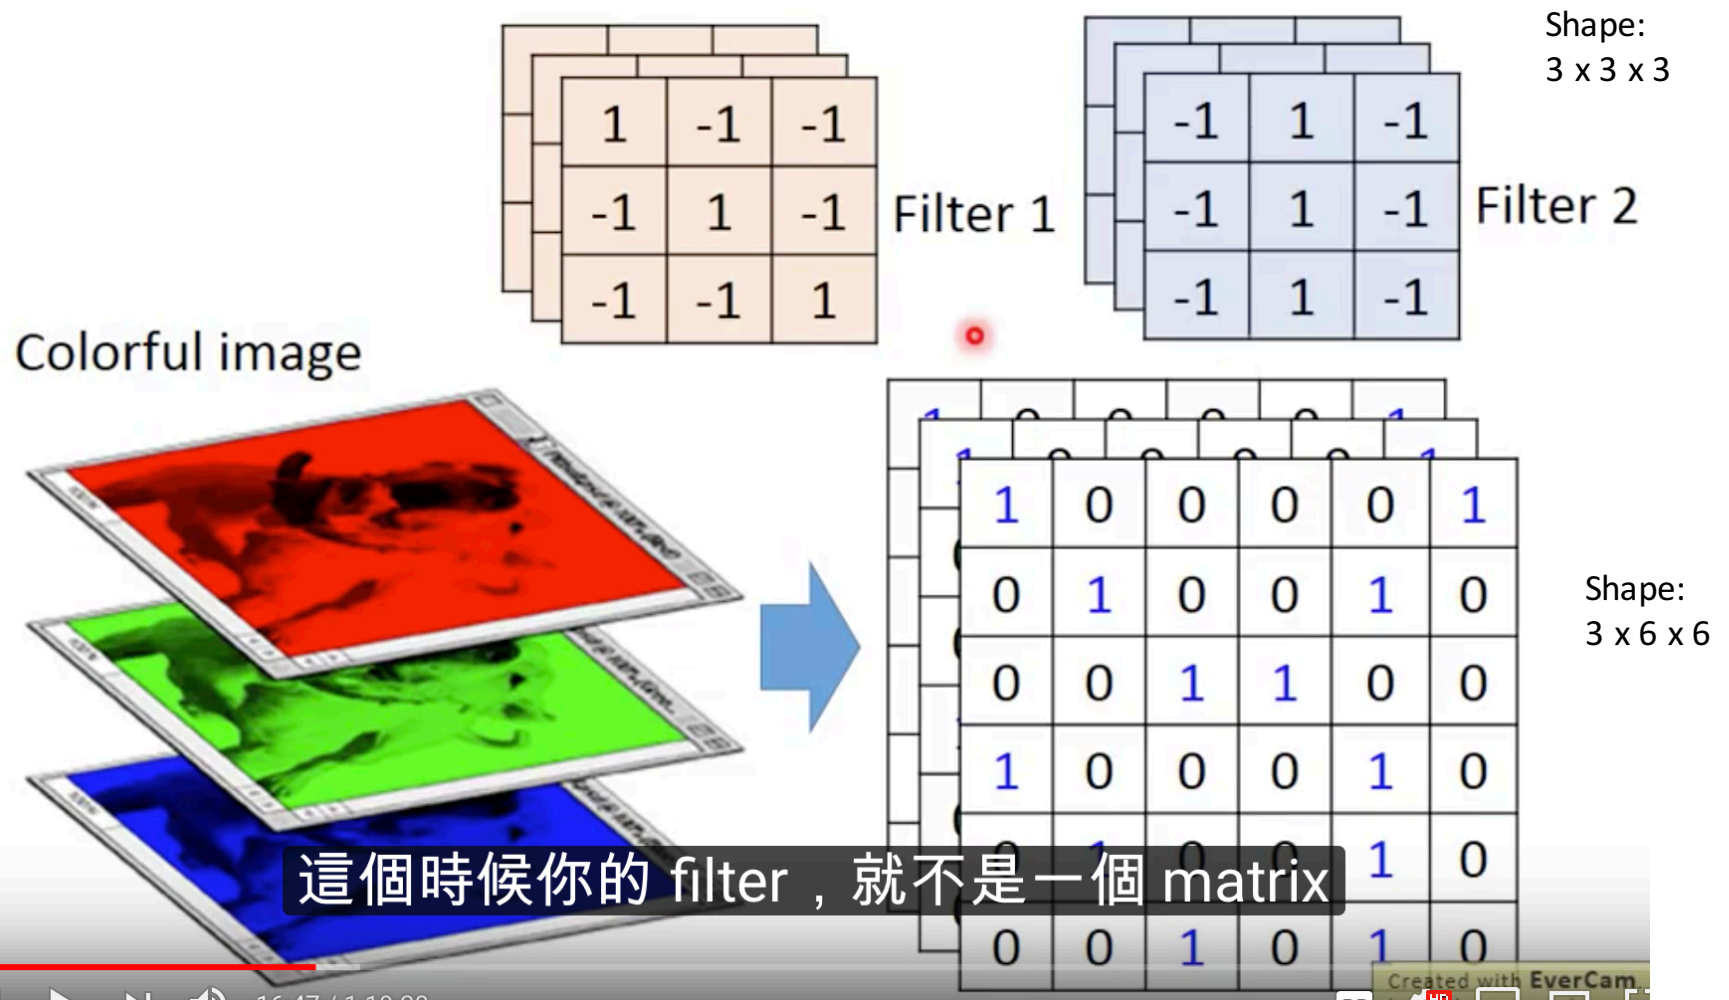


# CNN – Colorful image


Colorful image



# CNN – Colorful image



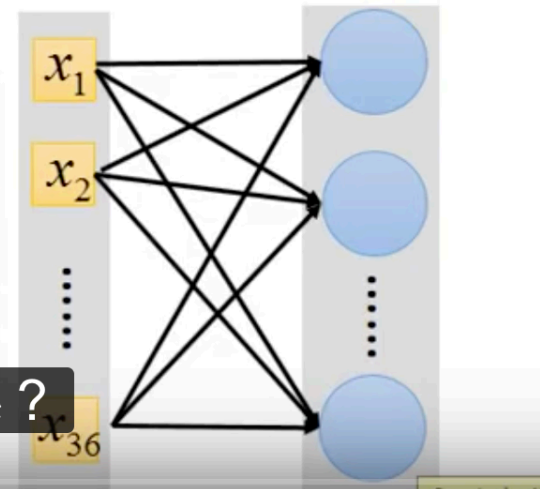
# Convolution v.s. Fully Connected



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

有什麼關係？






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

Filter 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



6 x 6 image

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

Filter 1


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

6 x 6 image

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

- 1: 1
- 2: 0
- 3: 0
- 4: 0
- ⋮
- 7: 0
- 8: 1
- 9: 0
- 10: 0
- ⋮
- 13: 0
- 14: 0
- 15: 1
- ⋮

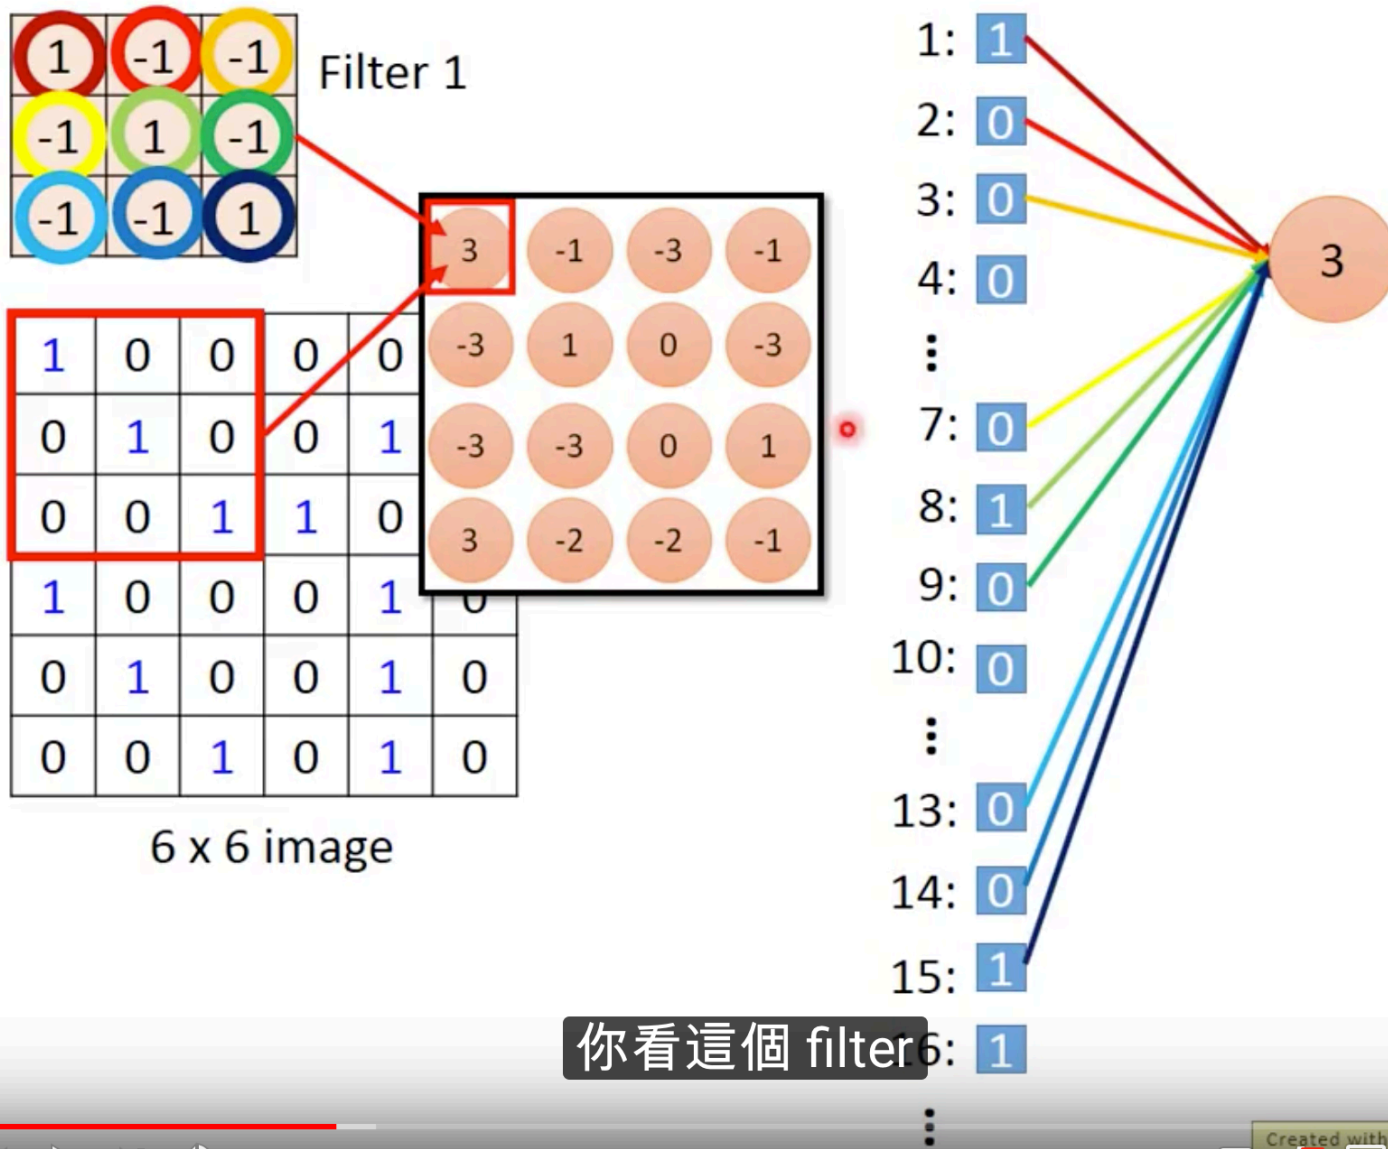
拉直變成這樣子





- 1: 1
- 2: 0
- 3: 0
- 4: 0
- ⋮
- 7: 0
- 8: 1
- 9: 0
- 10: 0
- ⋮
- 13: 0
- 14: 0
- 15: 1



然後，你有一個 neuron 的 output 是 3






Fewer parameters!


你有 36 個 pixel 當作 input

Only connect to 9 input, not fully connected




6 x 6 image

Less parameters!




發生甚麼事呢?




6 x 6 image

Less parameters!




那這個 neuron 連接到哪些 input 的 weight 呢






6 x 6 image

Less parameters!




這個框起來的地方，它正好就對應到 pixel 2, 3, 4



6 x 6 image

Less parameters!

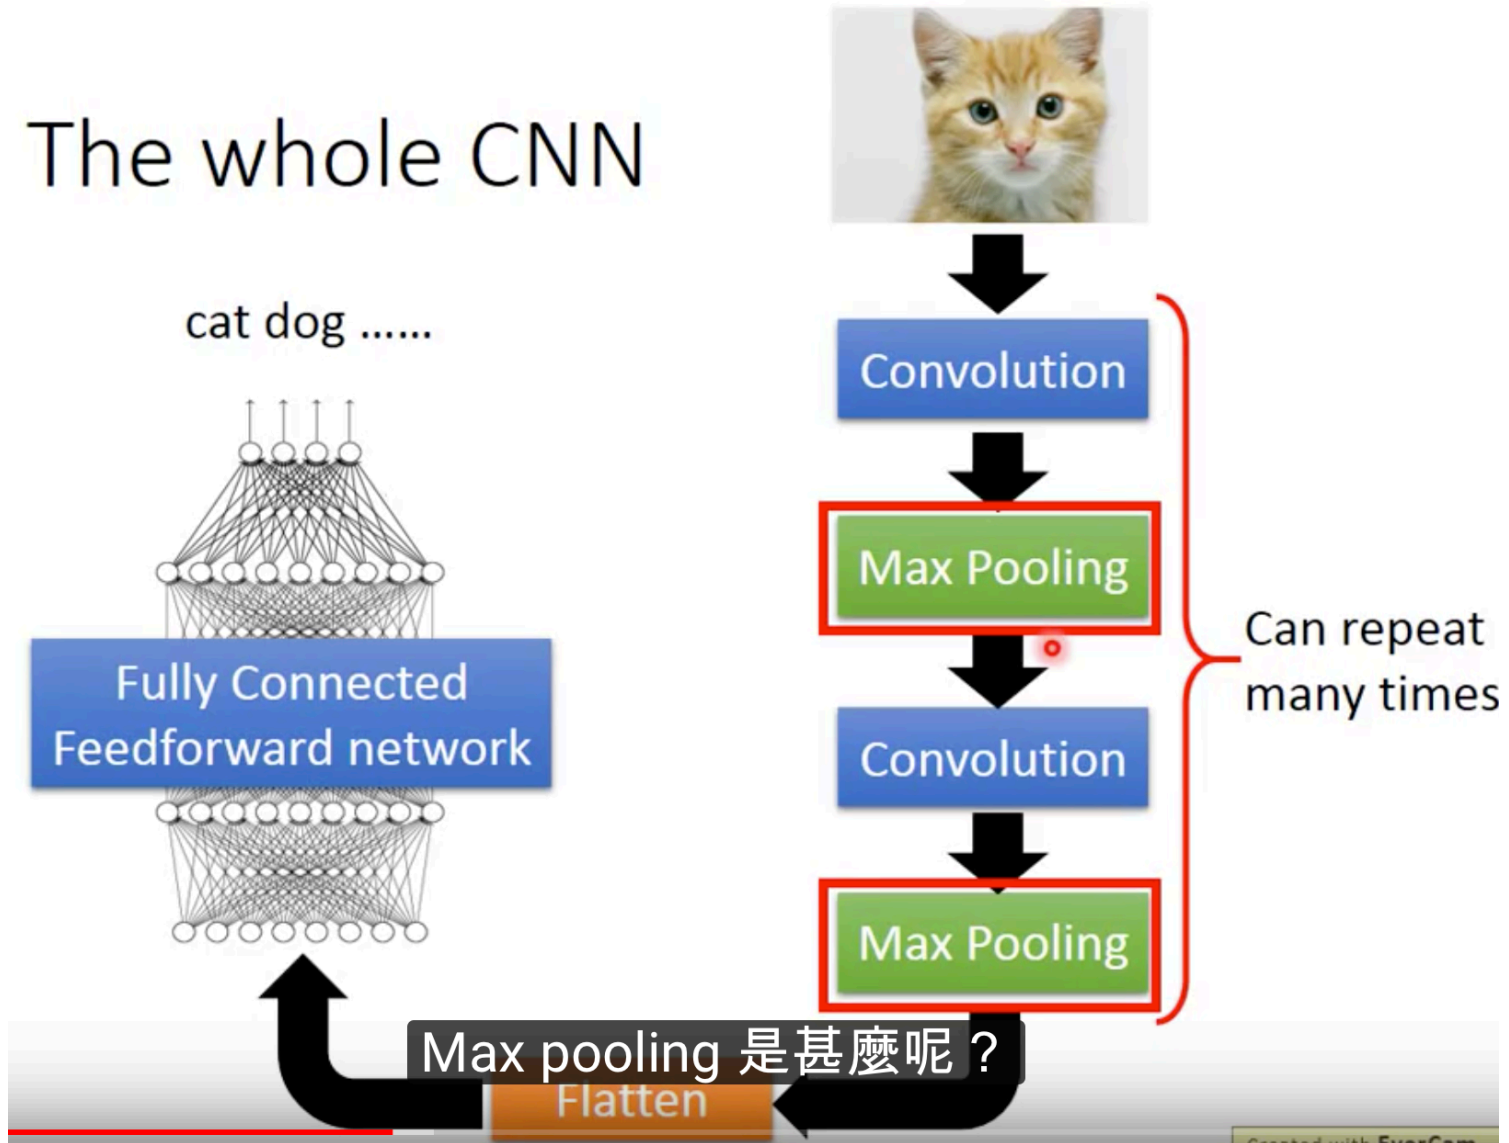
但是，當我們做這個



How to train shared weight?  
Average their gradients.

Shared weights

# The whole CNN




# CNN – Max Pooling

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

Filter 1

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

Filter 2



接下來，我們做甚麼事呢？

# CNN – 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

-1	-1	-1	-1
-1	-1	-2	1

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

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

所以，3, -1, 3, 1, 我就選 3

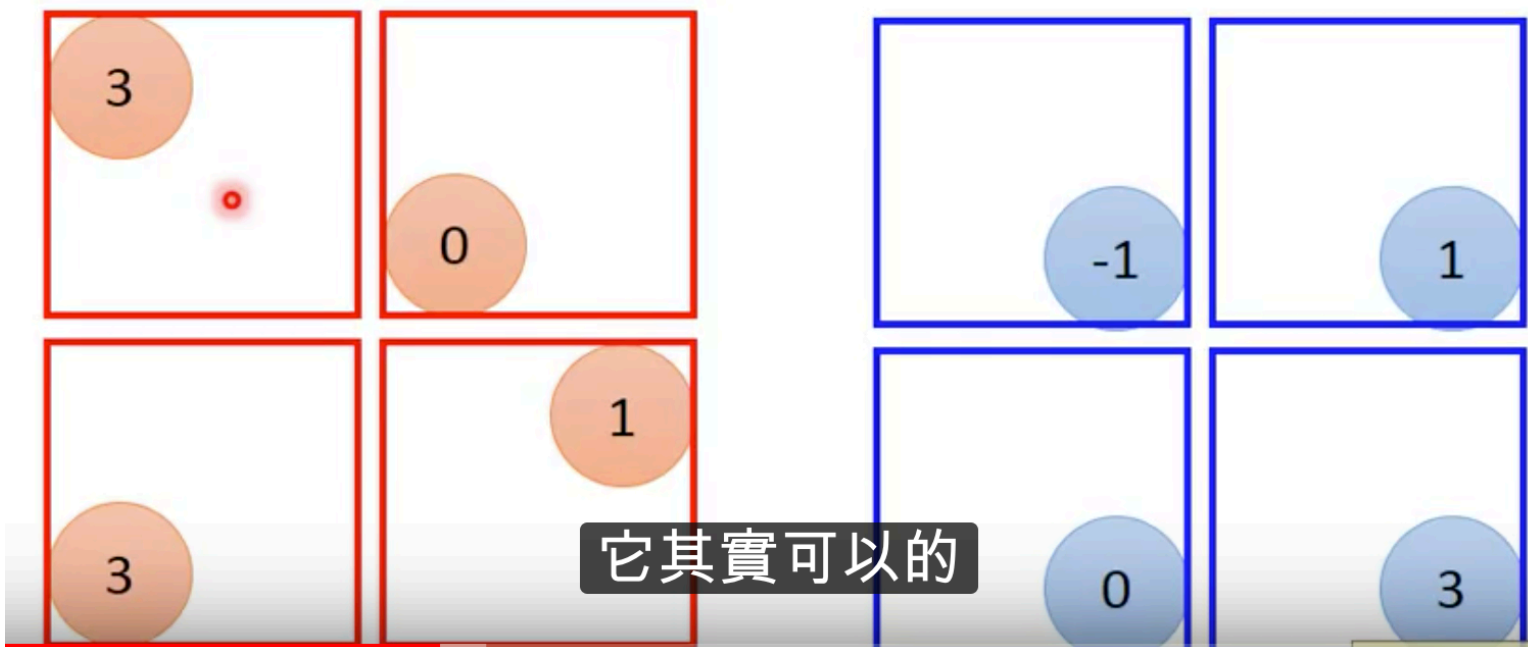
# CNN – Max Pooling

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

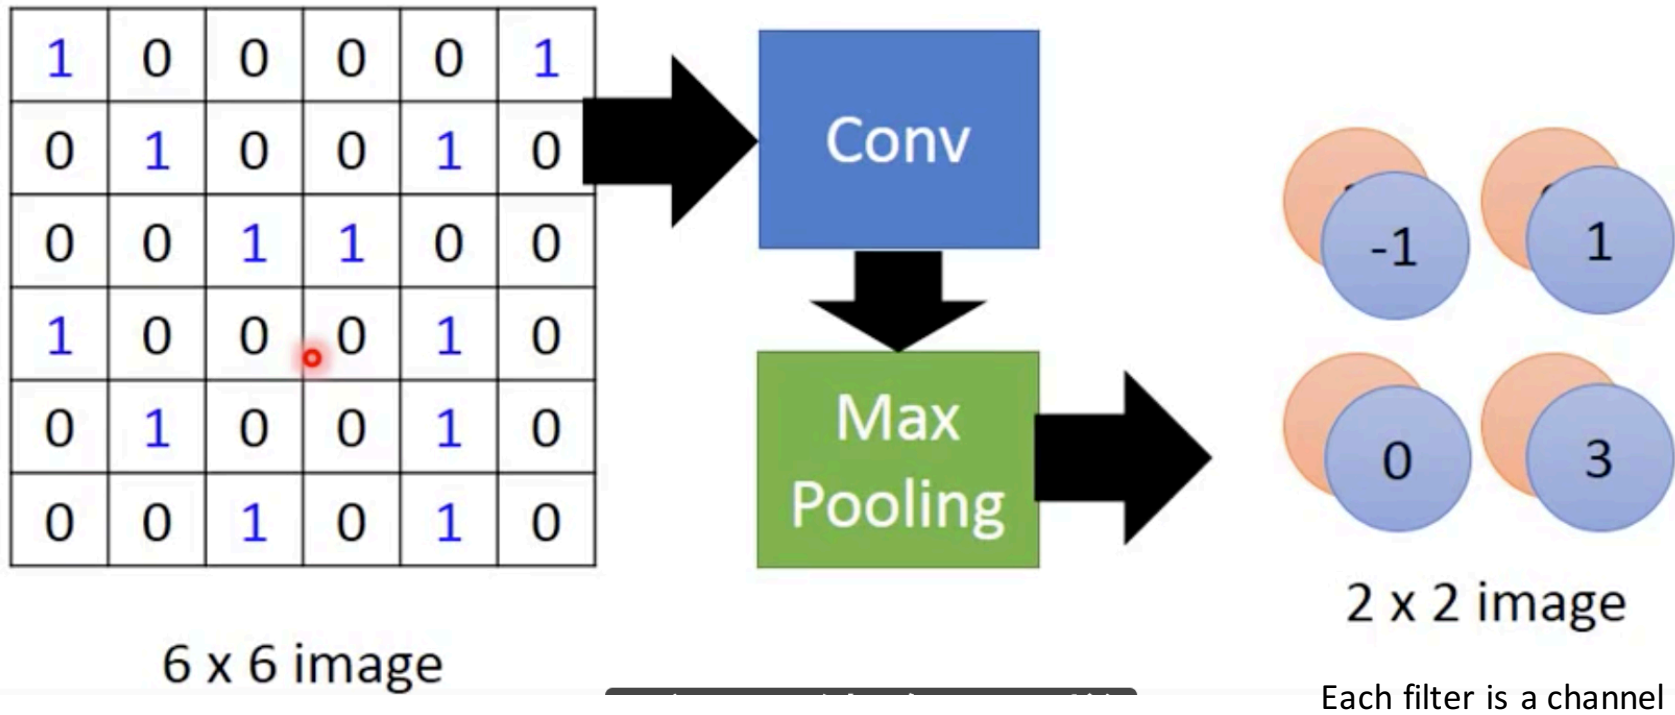
Filter 1

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


Filter 2



# CNN – Max Pooling




# The whole CNN



A new image

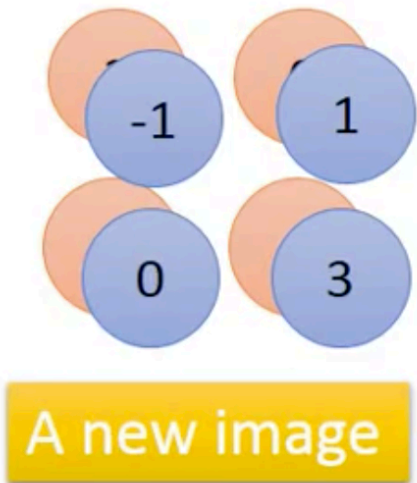
Smaller than the original image



Can repeat many times

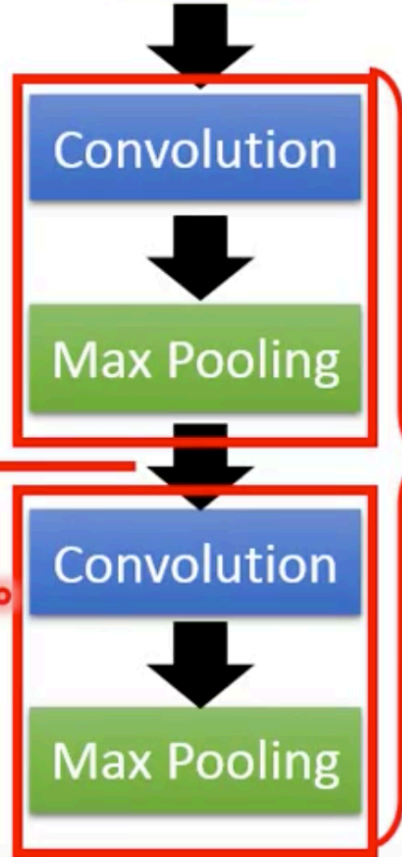


# The whole CNN



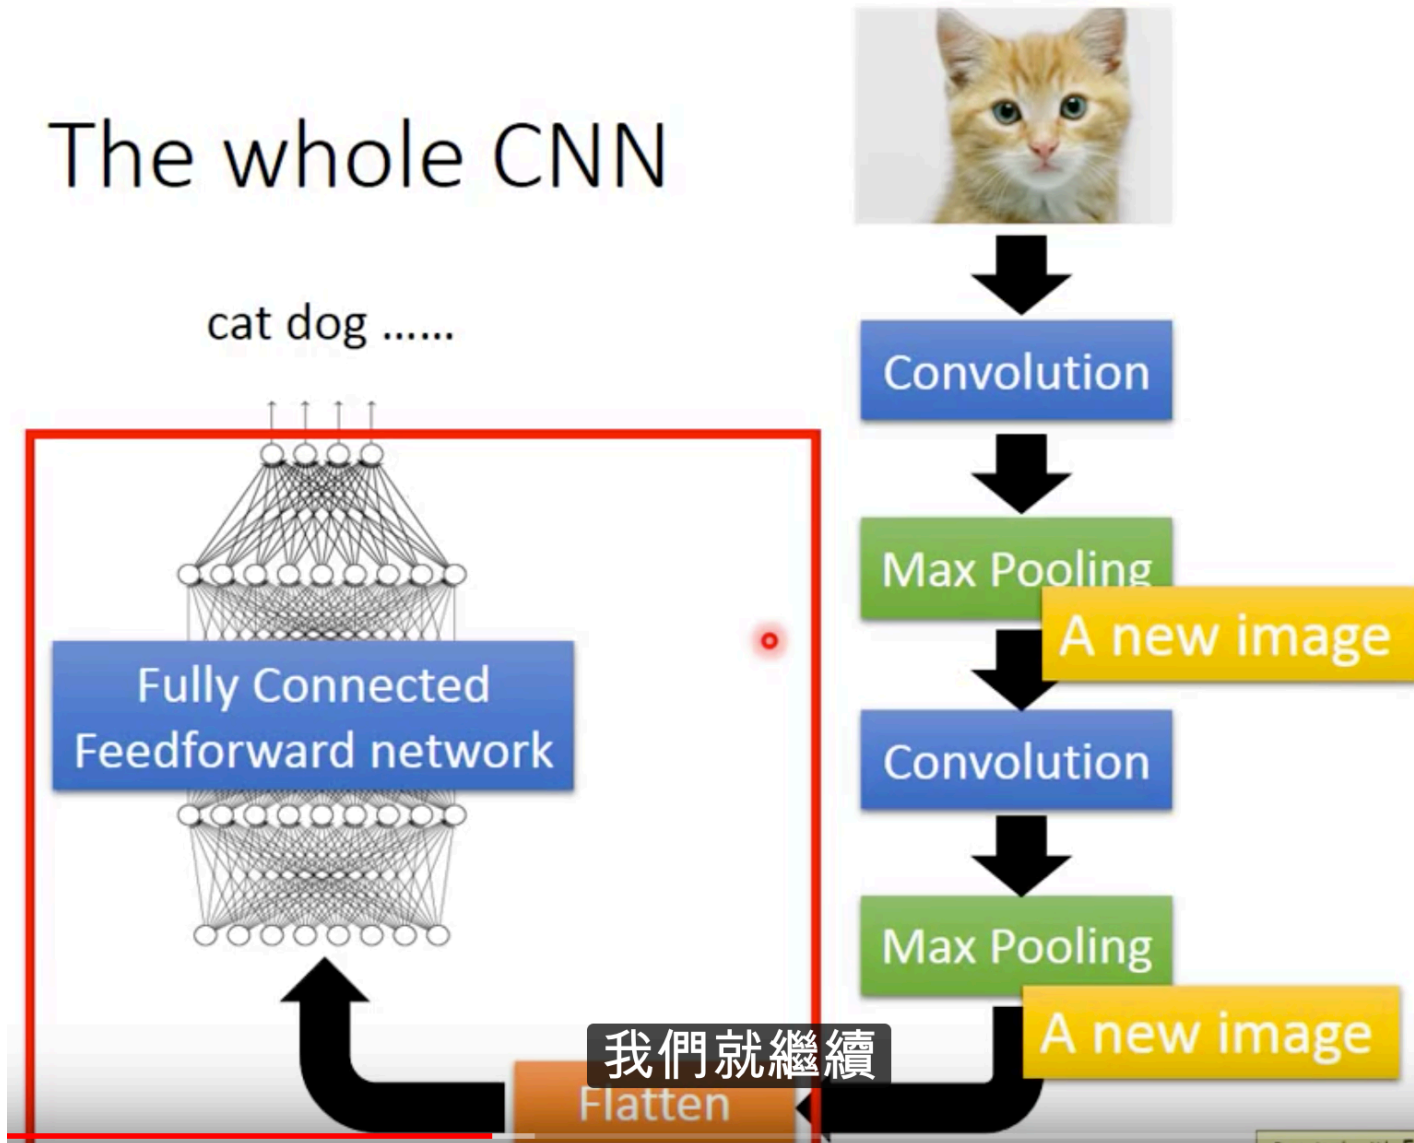
Smaller than the original image

The number of the channel is the number of filters

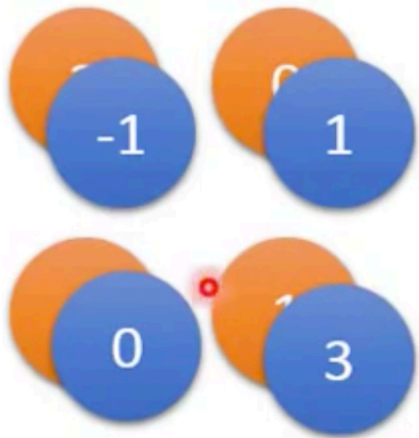


Can repeat many times

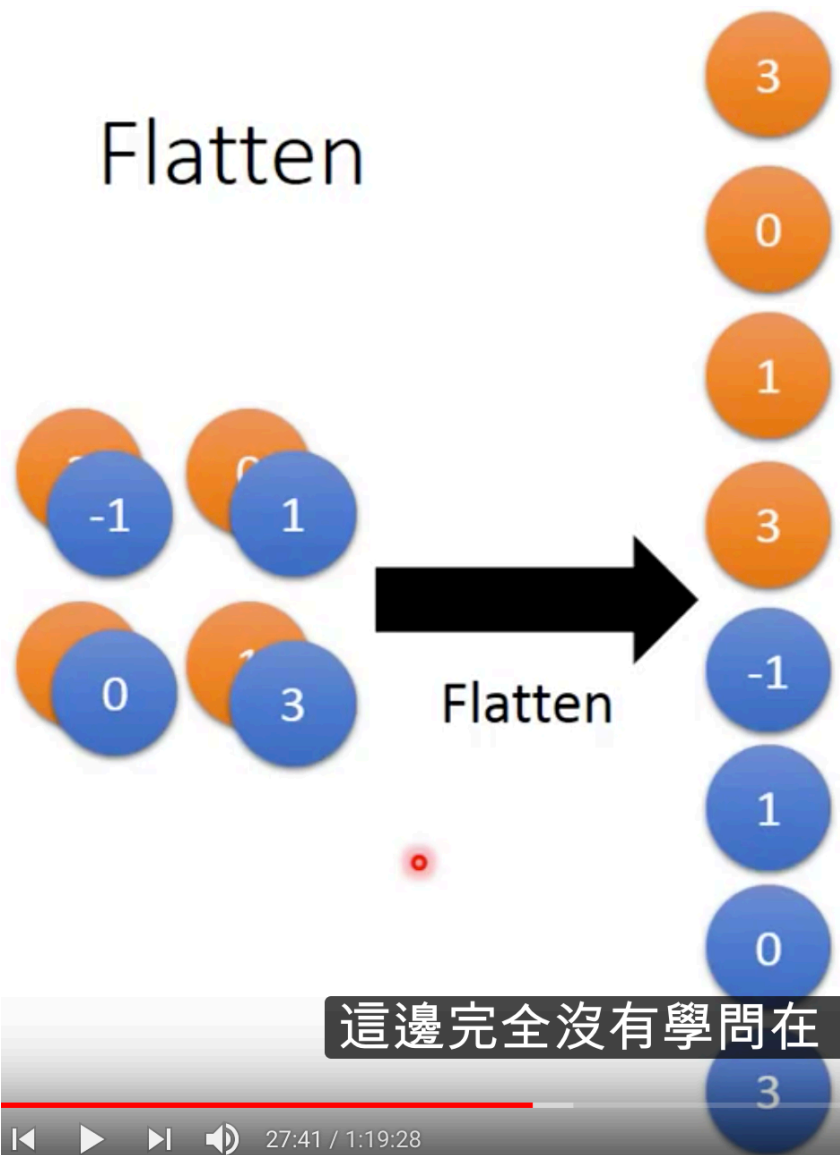
# The whole CNN



# Flatten

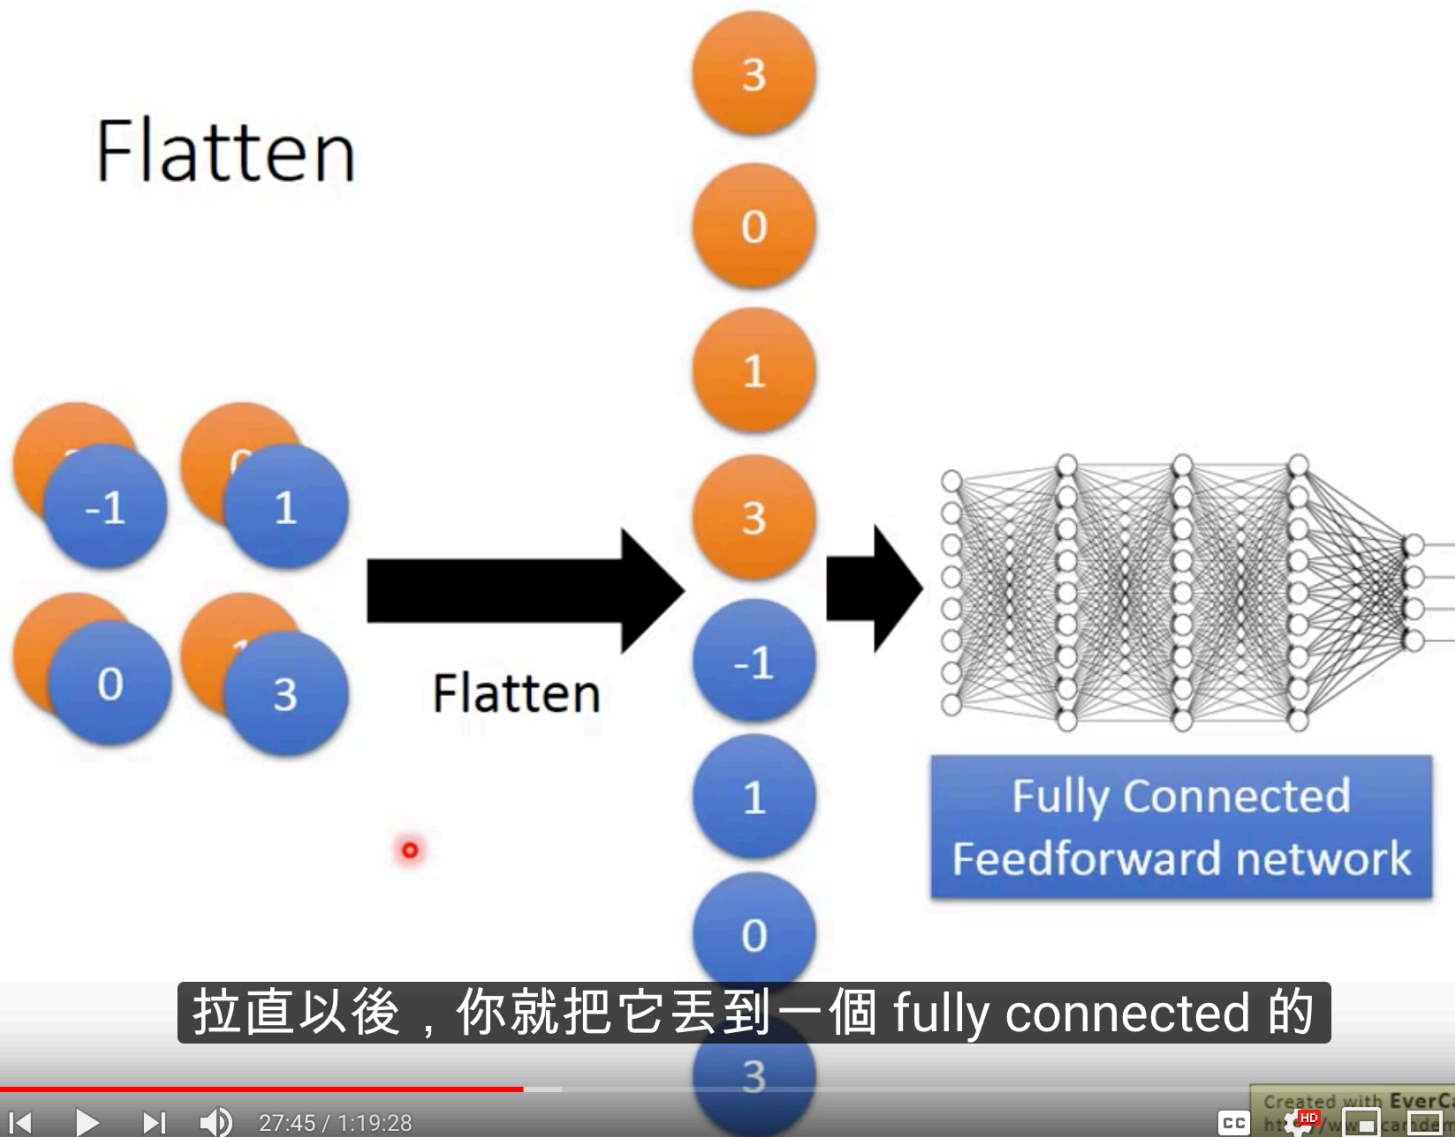


# Flatten



這邊完全沒有學問在

# Flatten



拉直以後，你就把它丟到一個 fully connected 的

## CNN in Keras

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



## CNN in Keras

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

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

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

..... There are 25  
3x3 filters.

Input\_shape = ( 1, 28, 28 )

1: black/weight, 3: RGB, 28 x 28 pixels

input

Convolution

Max Pooling


Convolution

Max Pooling

# CNN in Keras

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


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



..... There are 25 3x3 filters.

Input\_shape = ( 1, 28, 28 )  
1: black/weight, 3: RGB 28 x 28 pixels

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



我們把 2\*2 的這個





## CNN in Keras

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

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


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

..... There are 25  
3x3 filters.

Input\_shape = ( 1, 28, 28 )

1: black/weight, 3: RGB 28 x 28 pixels

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



input

Convolution

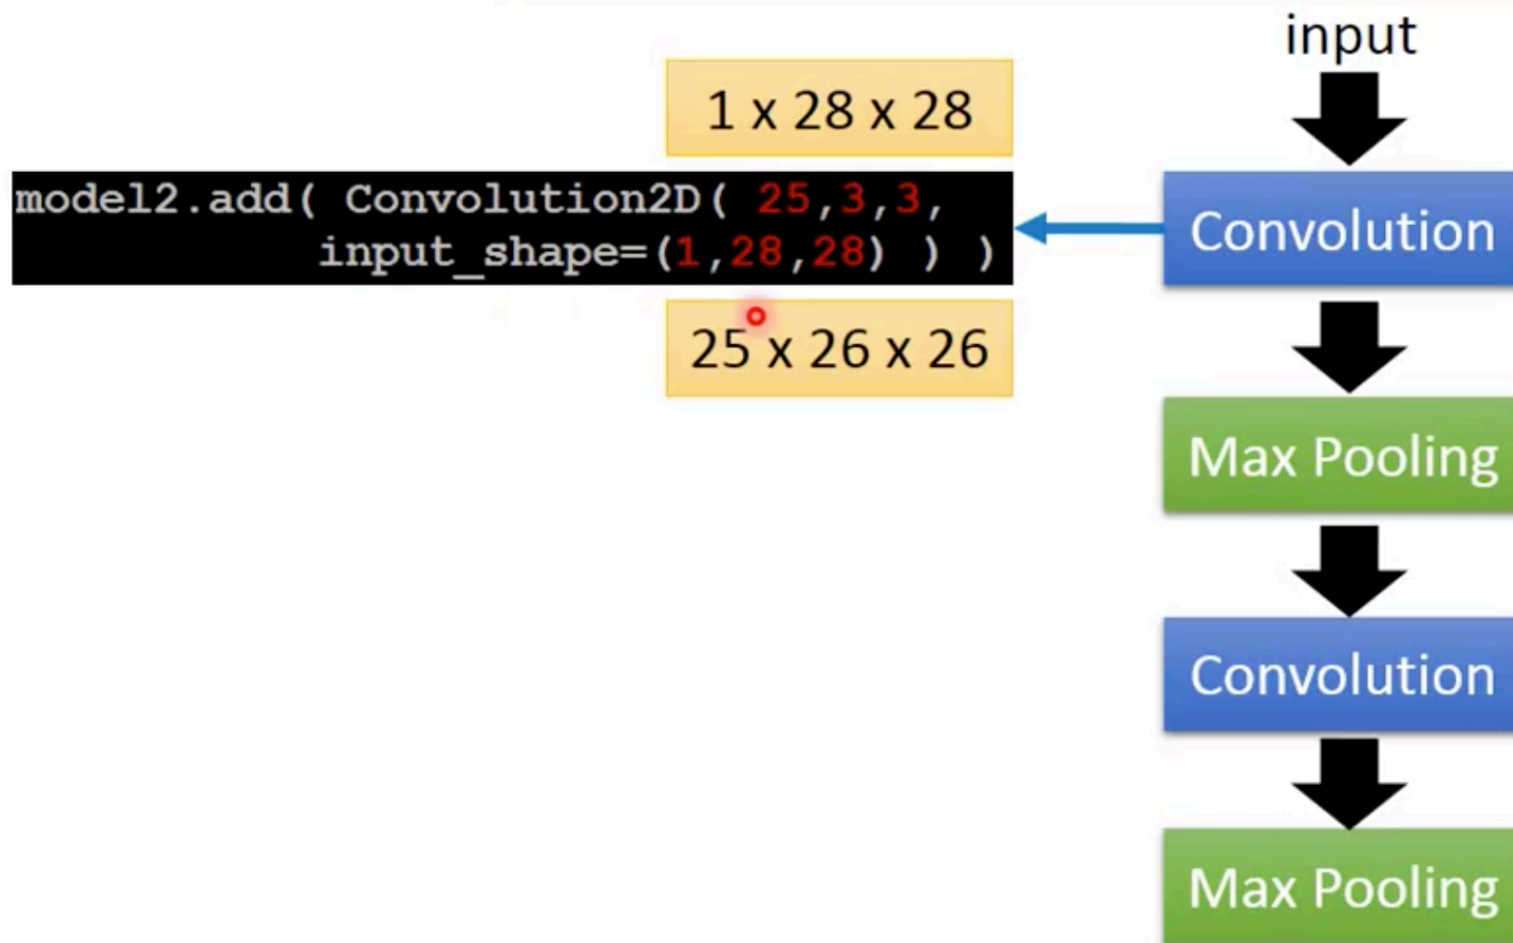
Max Pooling

Convolution

Max Pooling

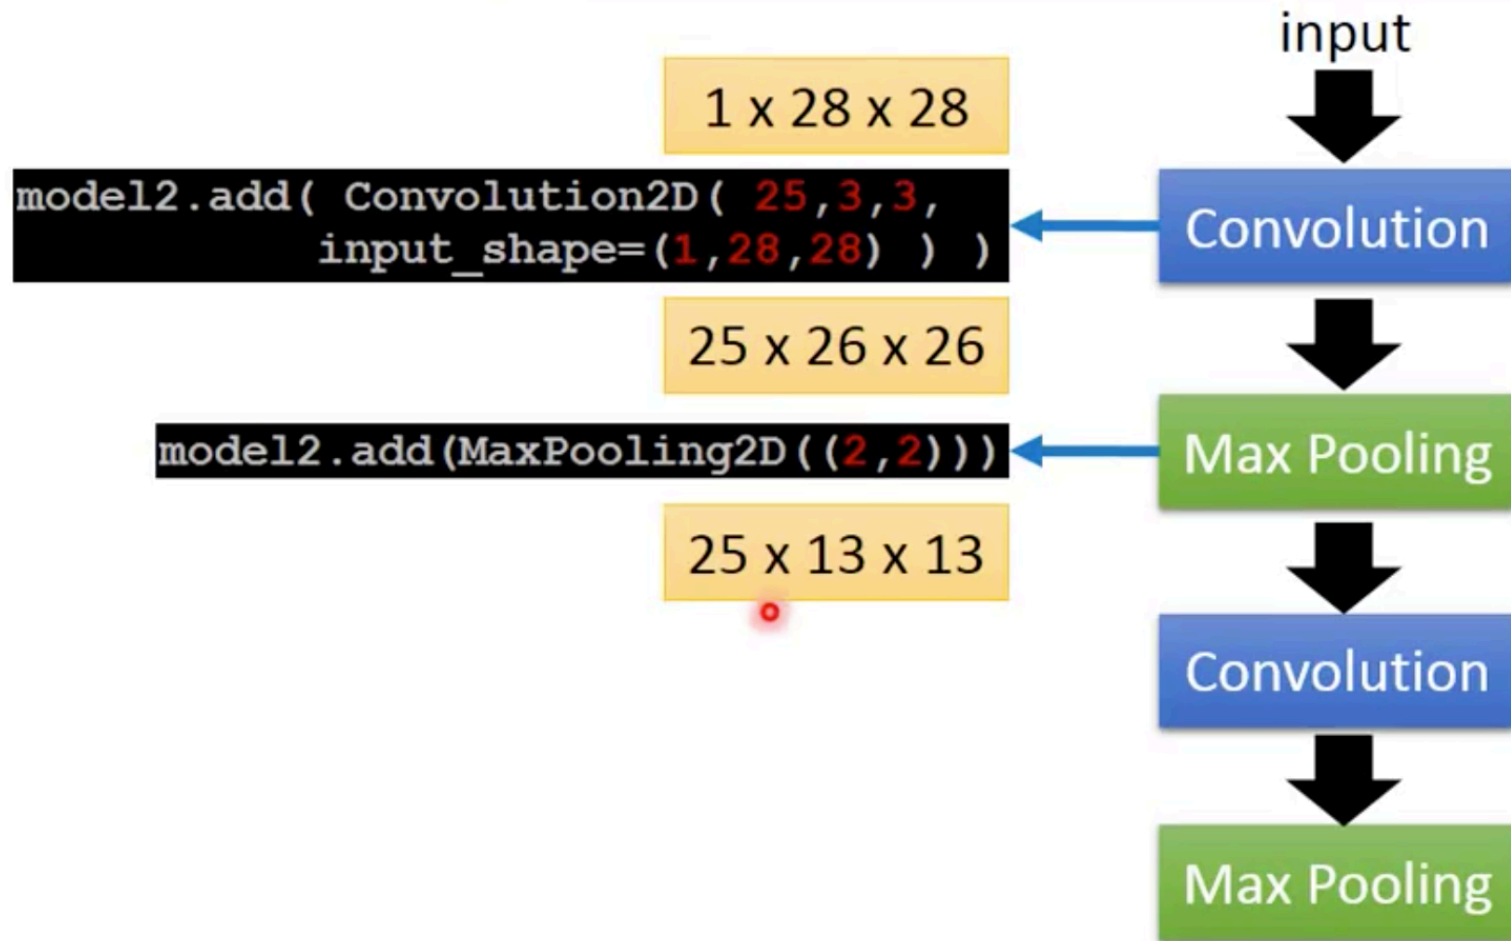
## CNN in Keras

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



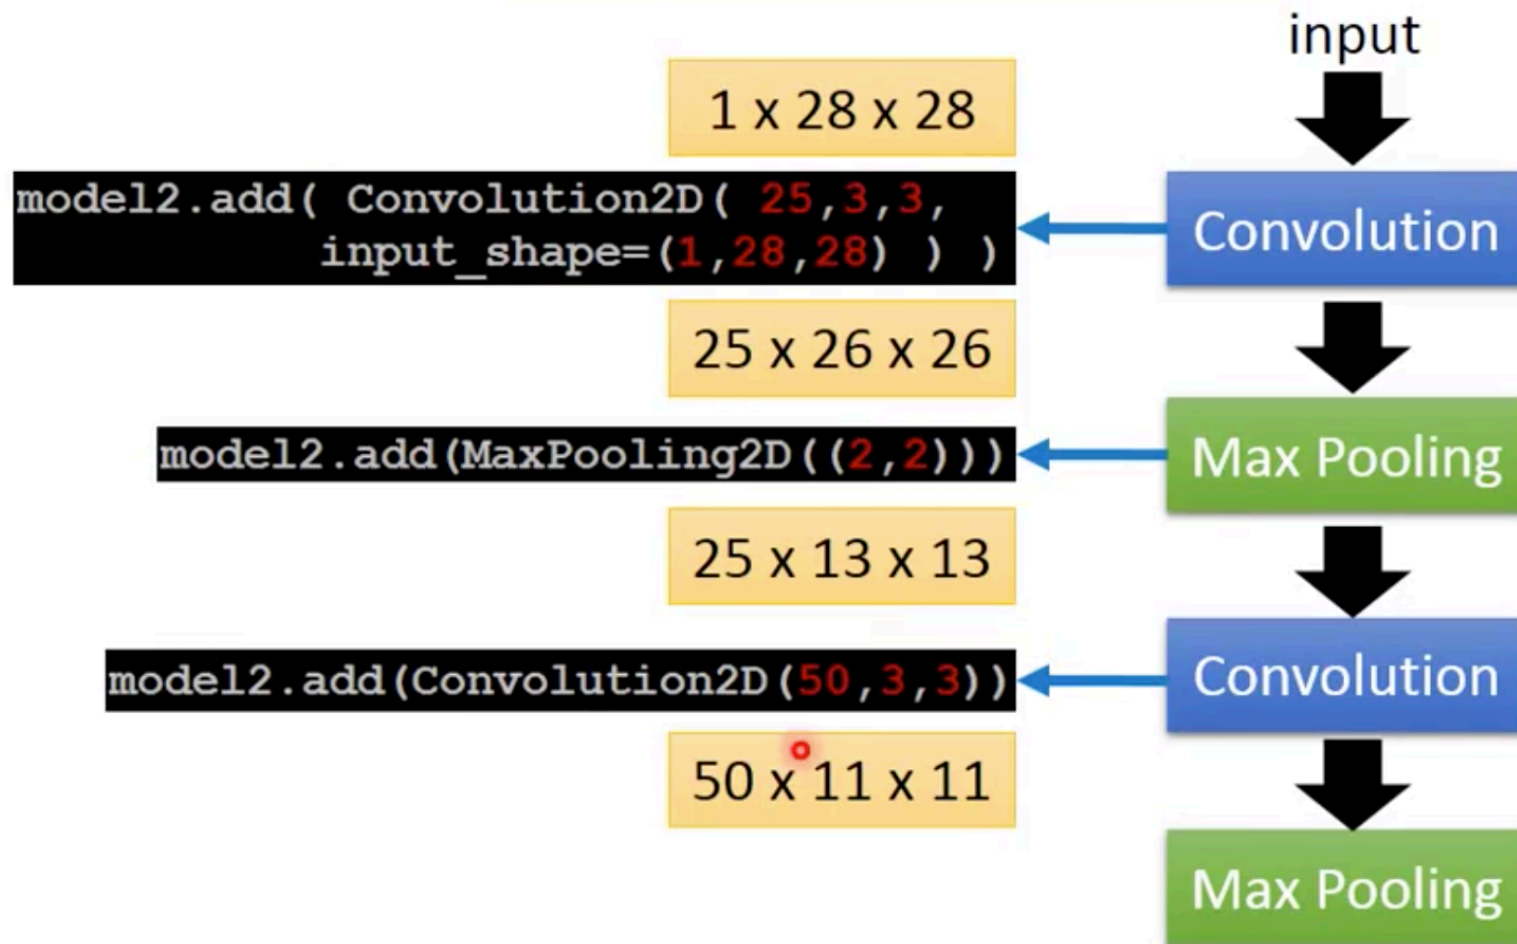
## CNN in Keras

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



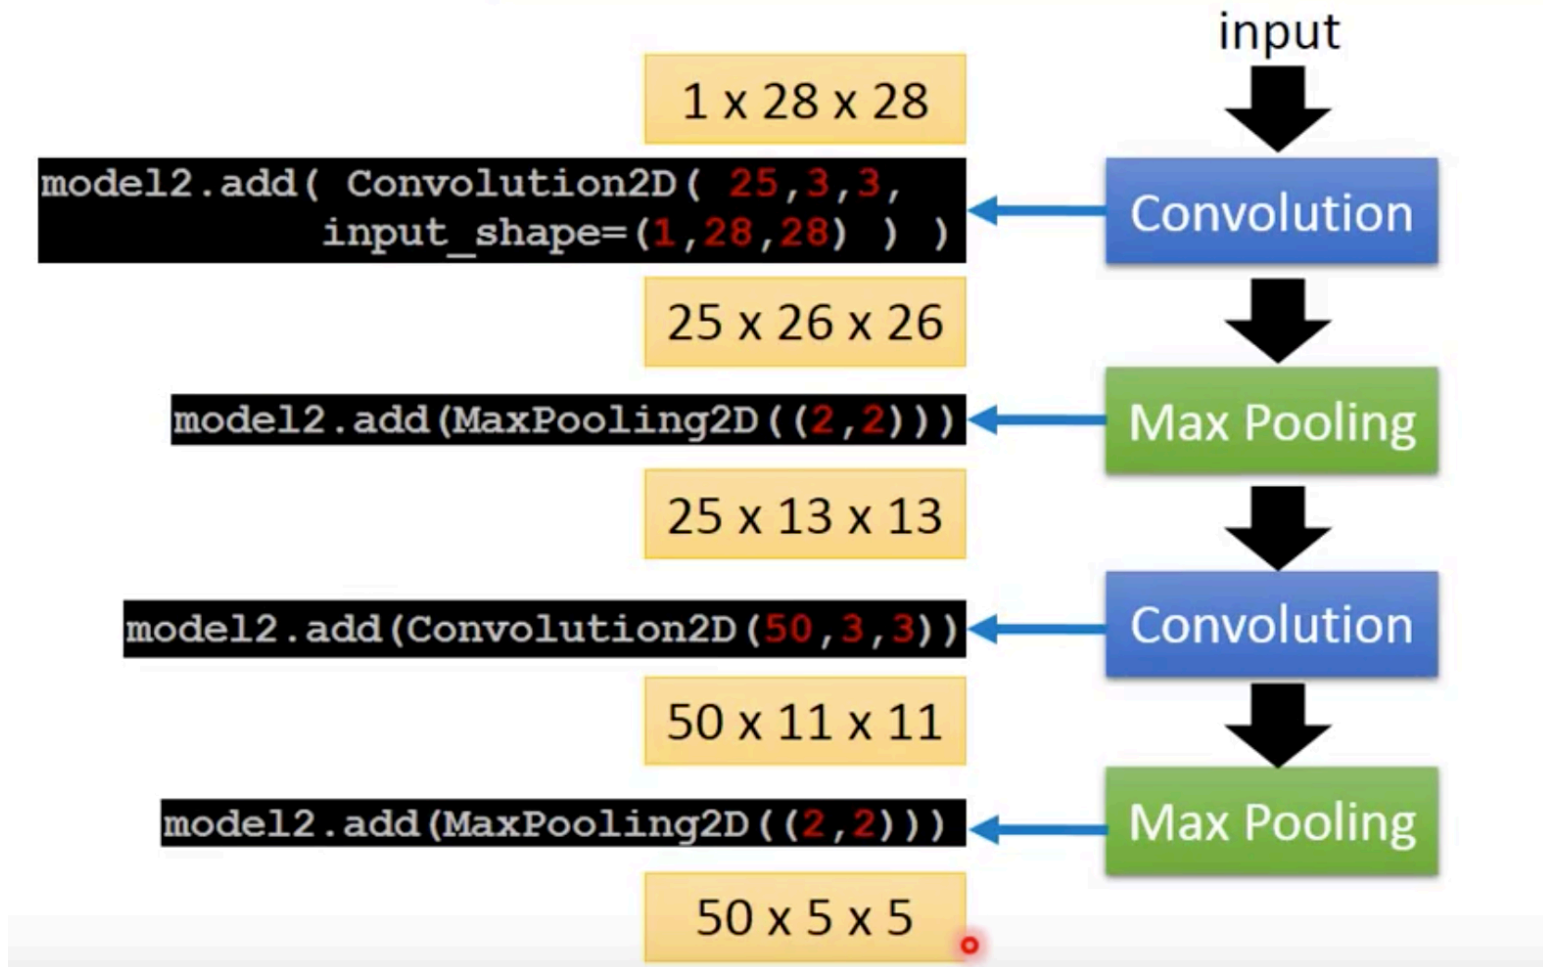
## CNN in Keras

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



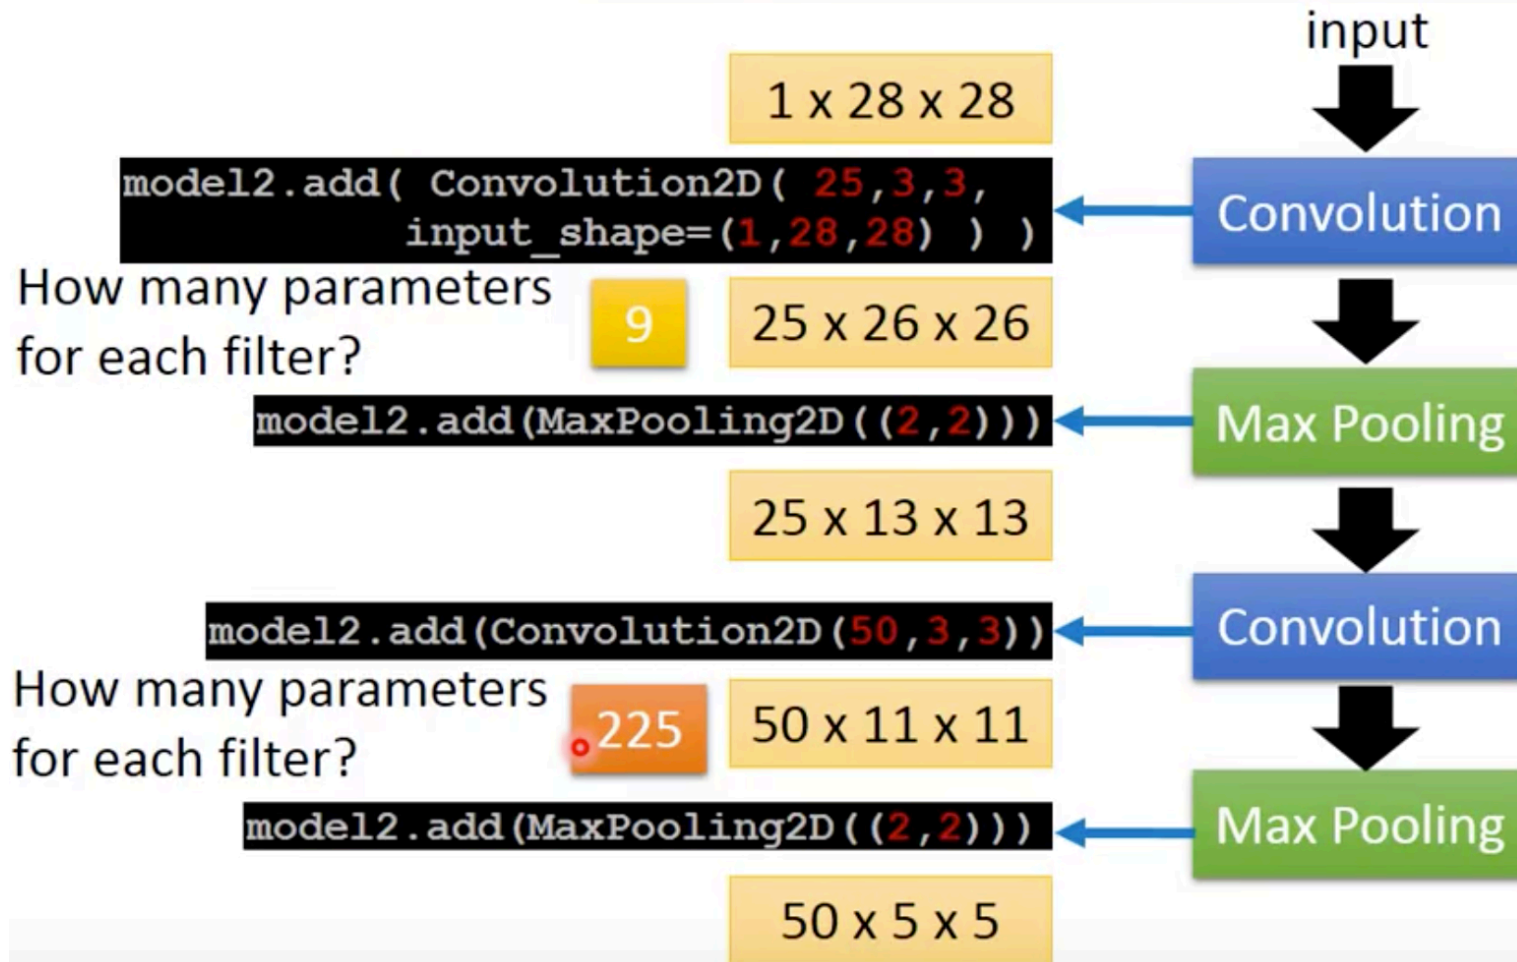
## CNN in Keras

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



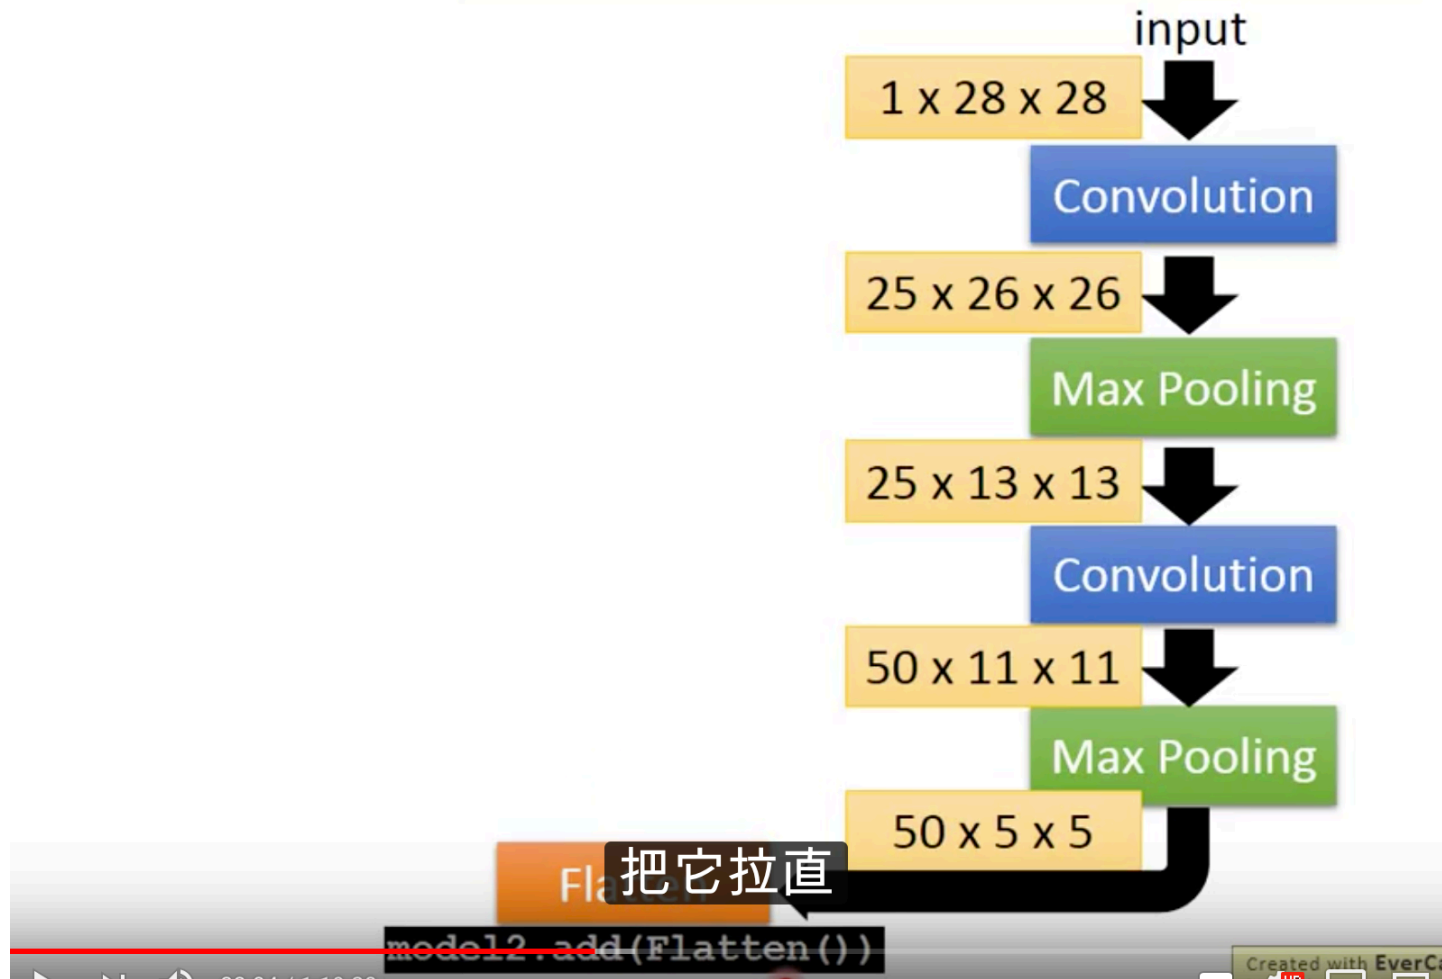
## CNN in Keras

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



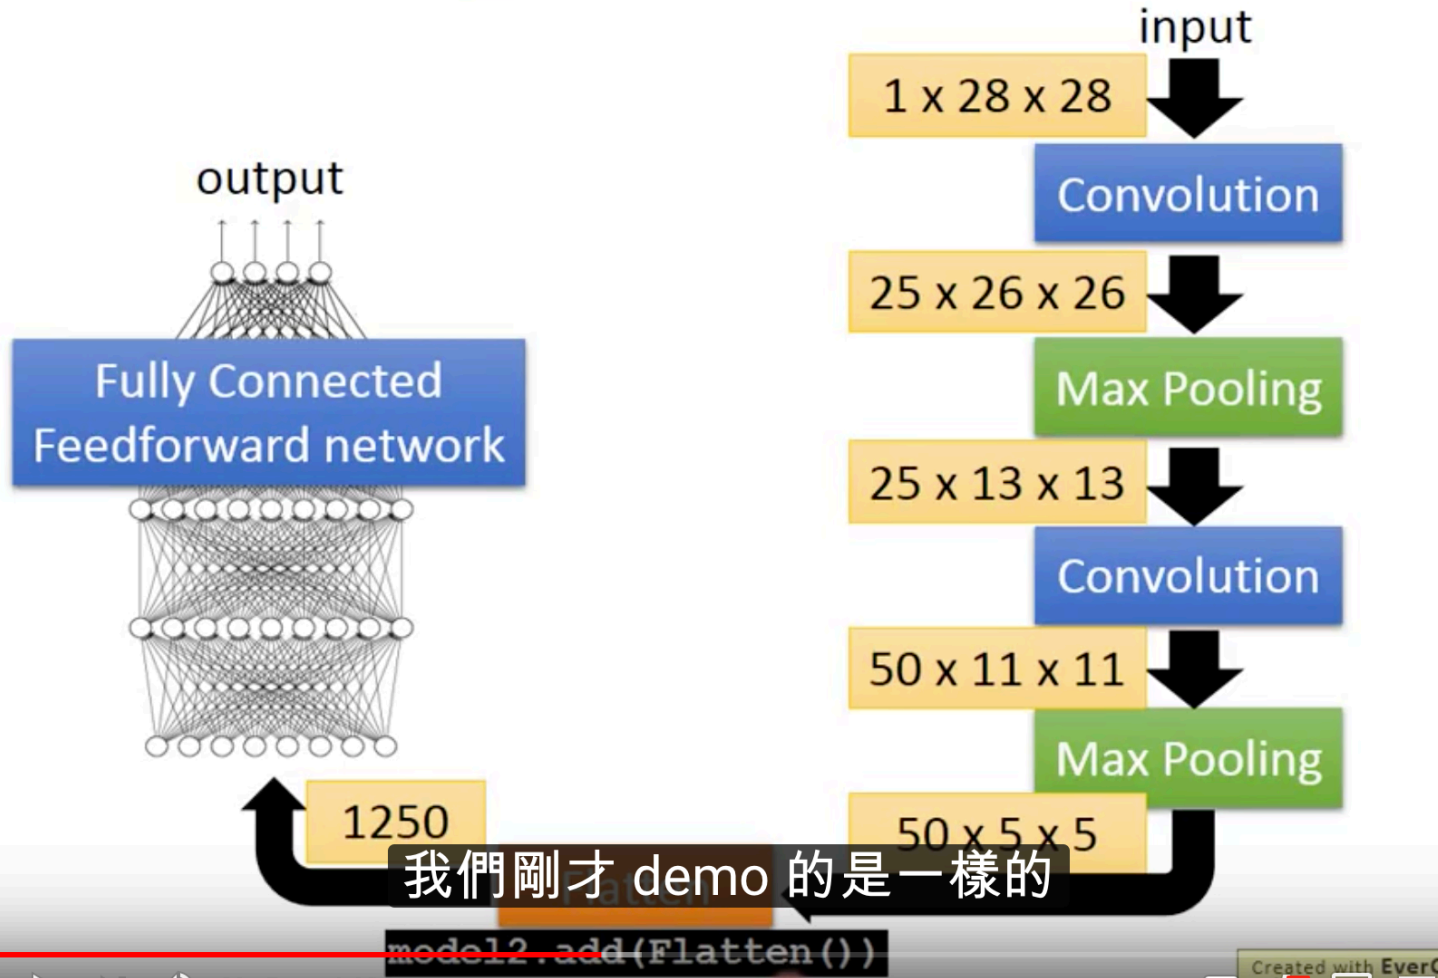
## CNN in Keras

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



# CNN in Keras

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





# Deep Dream



- Given a photo, machine adds what it sees .....



# Deep Dream

- Given a photo, machine adds what it sees .....





那 Deep Dream 還有一個進階的版本

# Deep Style

- Given a photo, make its style like famous paintings

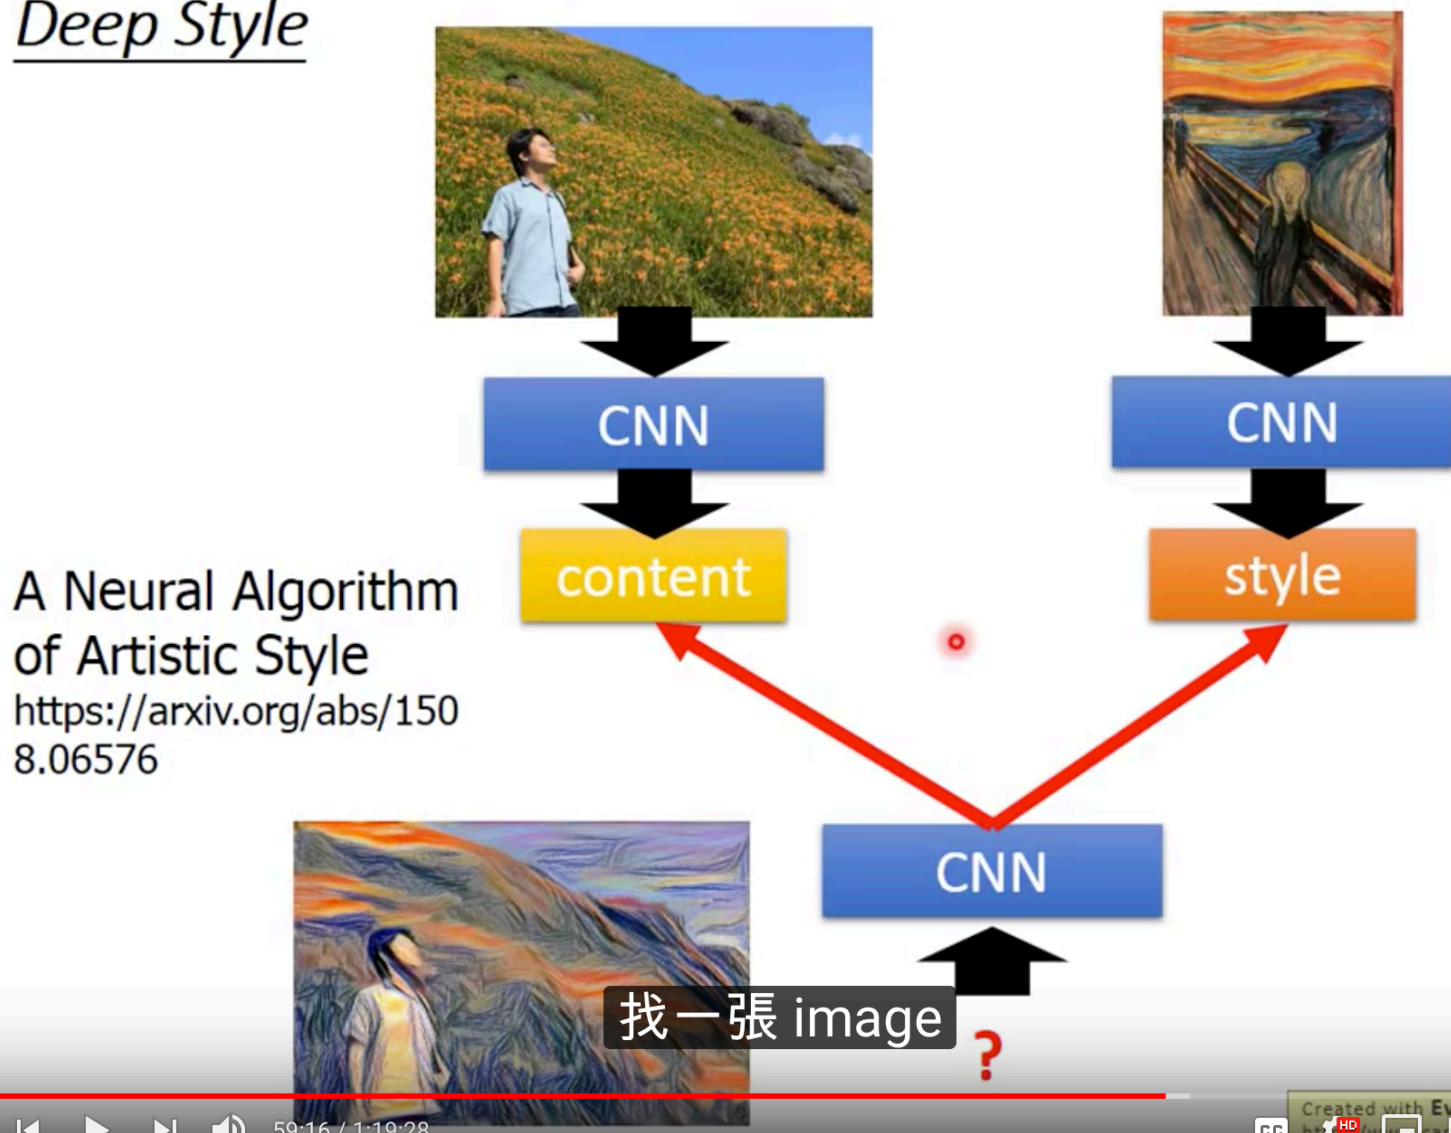


# Deep Style

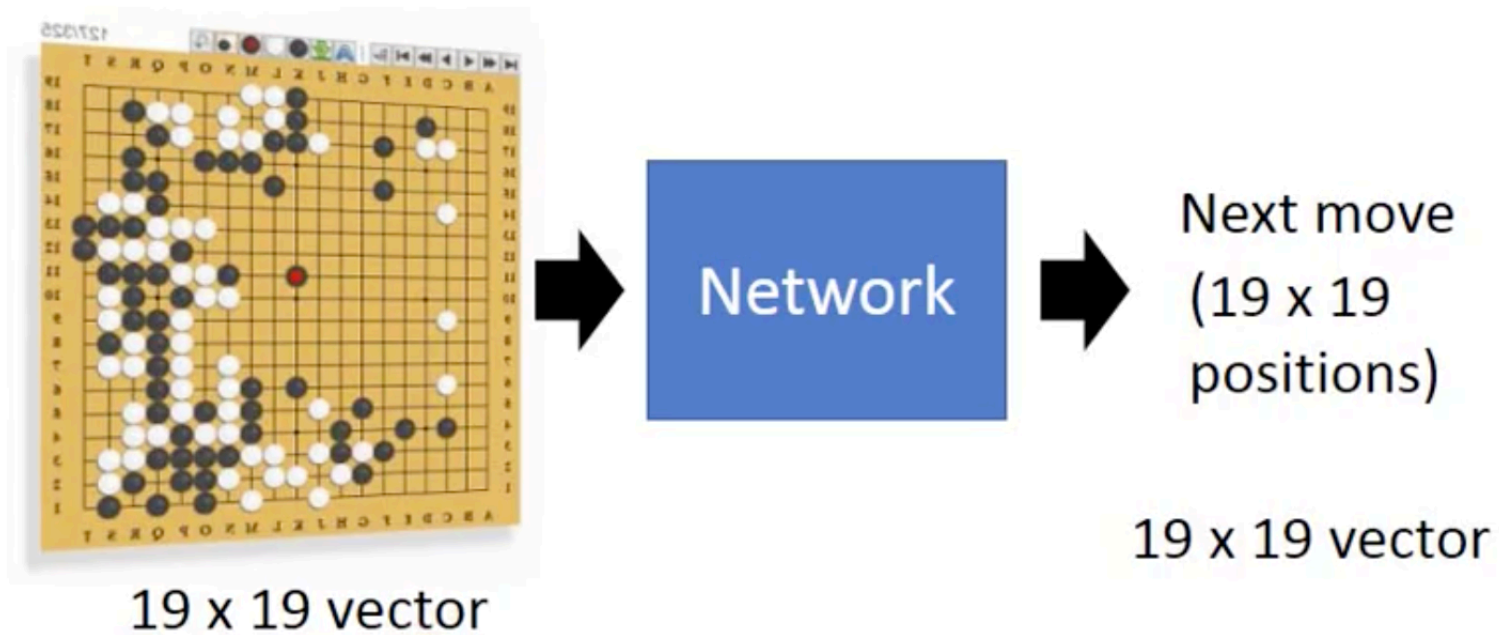


A Neural Algorithm  
of Artistic Style  
<https://arxiv.org/abs/1508.06576>

# Deep Style



# More Application: Playing Go



19 x 19 vector


Black: 1  
white: -1

Fully-connected feedforward network can be used


But CNN performs much better.

# More Application: Playing Go

Training: record of previous plays 黑: 5之五 → 白: 天元 → 黑: 五之5 ...



Target:  
“天元” = 1  
else = 0



Target:  
“五之5” = 1  
else = 0

我們知道說，大家都是

# Why CNN for playing Go?

- Some patterns are much smaller than the whole image



- The same patterns appear in different regions.



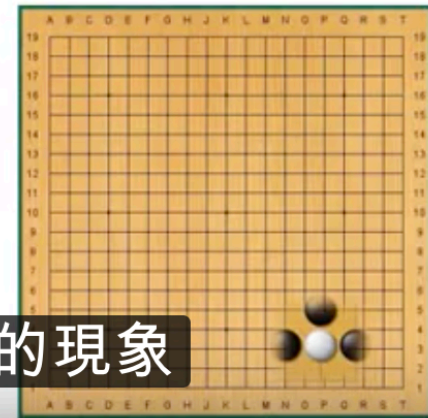

# Why CNN for playing Go?

- Some patterns are much smaller than the whole image

Alpha Go uses 5 x 5 for first layer



- The same patterns appear in different regions.



在圍棋上，也可能有同樣的現象

# Why CNN for playing Go?

- Subsampling the pixels will not change the object



Max Pooling

How to explain this???

# Why CNN for playing Go?

- Subsampling the pixels will not change the object



Max Pooling

How to explain this???

**Neural network architecture.** The input to the policy network is a  $19 \times 19 \times 48$  image stack consisting of 48 feature planes. The first hidden layer zero pads the input into a  $23 \times 23$  image, then convolves  $k$  filters of kernel size  $5 \times 5$  with stride 1 with the input image and applies a rectifier nonlinearity. Each of the subsequent hidden layers 2 to 12 zero pads the respective previous hidden layer into a  $21 \times 21$  image, then convolves  $k$  filters of kernel size  $3 \times 3$  with stride 1, again followed by a rectifier nonlinearity. The final layer convolves 1 filter of kernel size  $1 \times 1$  with stride 1, with a different bias for each position, and applies a softmax function. The match version of AlphaGo used  $k = 192$  filters; Fig. 2b and Extended Data Table 3 additionally show the results of training with  $k = 128, 256$  and 384 filters.

叫吃的狀態呢，等等

# Why CNN for playing Go?

- Subsampling the pixels will not change the object




Max Pooling

How to explain this???

**Neural network architecture.** The input to the policy network is a  $19 \times 19 \times 48$  image stack consisting of 48 feature planes. The first hidden layer zero pads the input into a  $23 \times 23$  image, then convolves  $k$  filters of kernel size  $5 \times 5$  with stride 1 with the input image and applies a rectifier nonlinearity. Each of the subsequent hidden layers 2 to 12 zero pads the respective previous hidden layer into a  $21 \times 21$  image, then convolves  $k$  filters of kernel size  $3 \times 3$  with stride 1, again followed by a rectifier nonlinearity. The final layer convolves 1 filter of kernel size  $1 \times 1$  with stride 1 with a different bias for each position and applies a softmax function. The **Alpha Go does not use Max Pooling .....** Extended Data Table 3 additionally show the results of training with  $k = 128, 256$  and 384 filters.

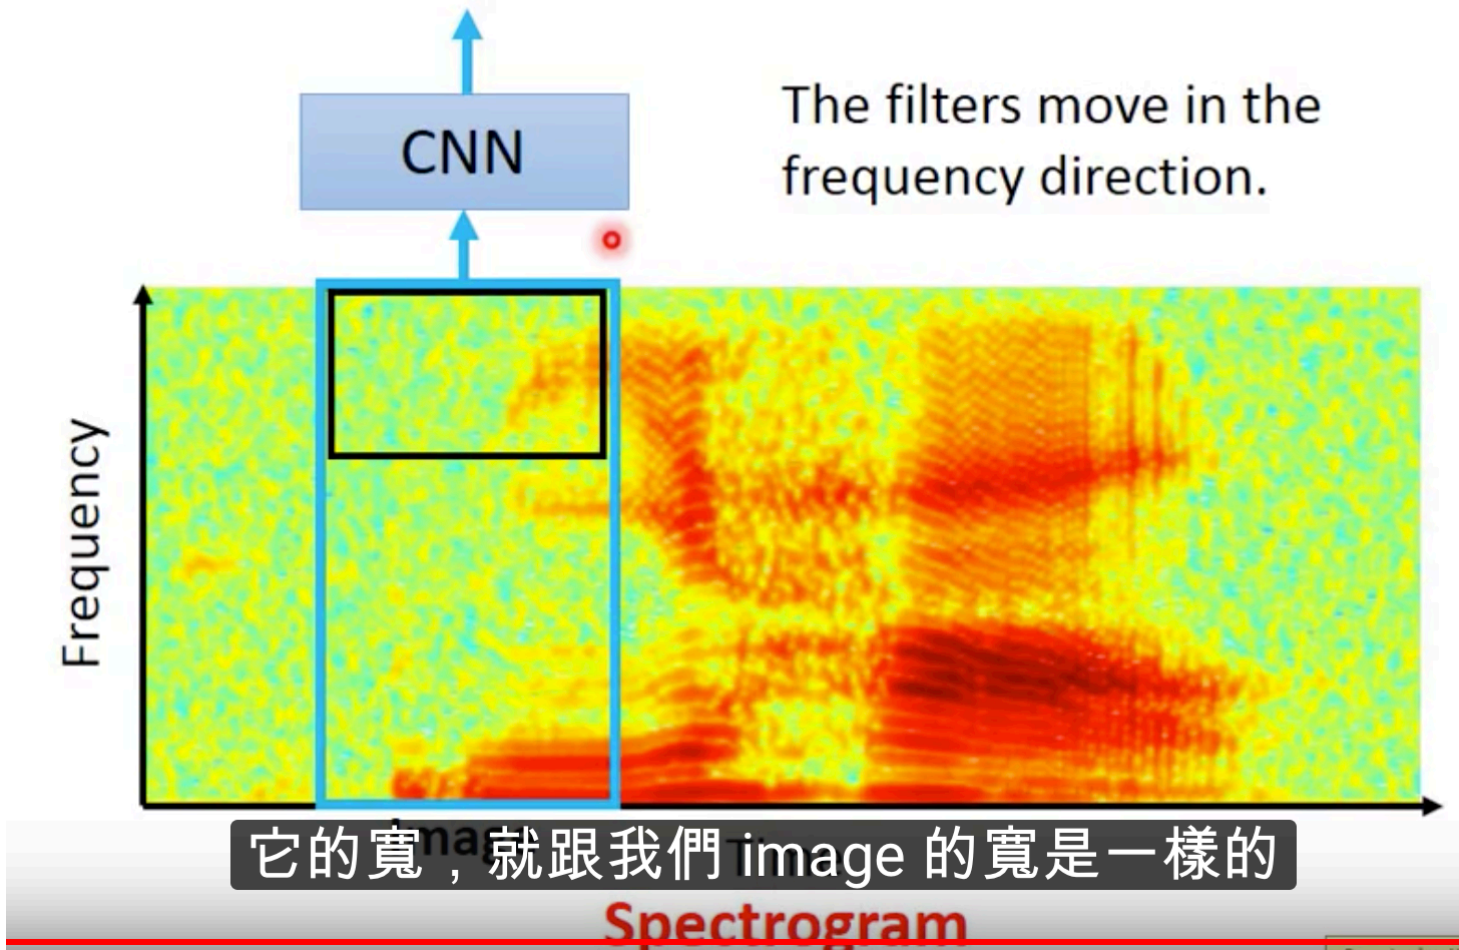
它是沒有用 Max Pooling 的

# More Application: Speech

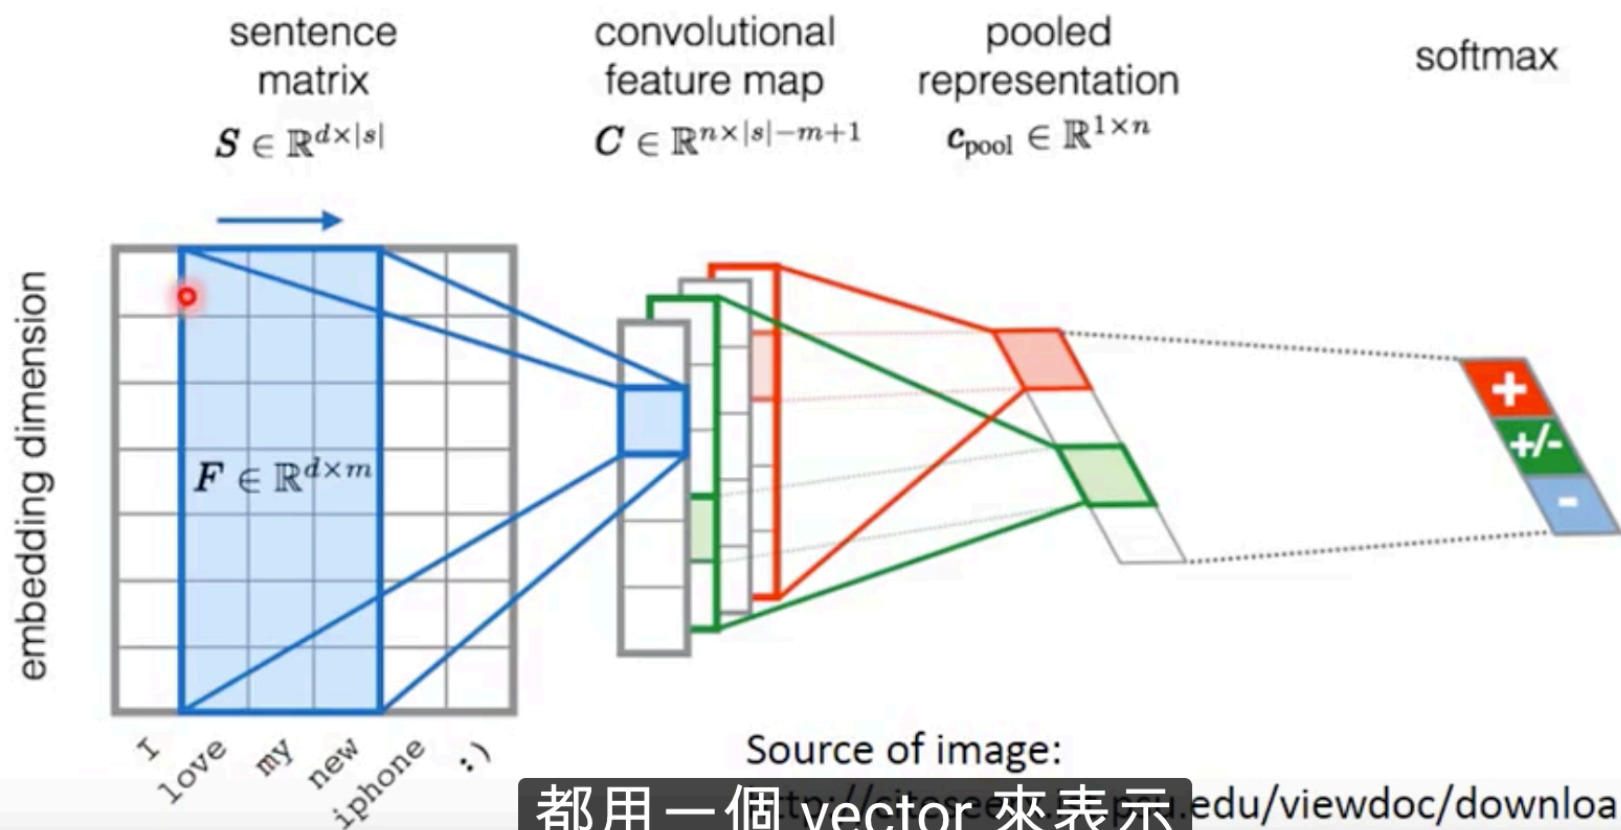


Time  
**Spectrogram**

# More Application: Speech



# More Application: Text



Source of image:

都用一個 vector 來表示

<http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.703.6858&rep=rep1&type=pdf>

## To learn more .....

- The methods of visualization in these slides
  - <https://blog.keras.io/how-convolutional-neural-networks-see-the-world.html>
- More about visualization
  - <http://cs231n.github.io/understanding-cnn/>
- Very cool CNN visualization toolkit
  - <http://yosinski.com/deepvis>
  - <http://scs.ryerson.ca/~aharley/vis/conv/>