Convolutional Neural Network

Announcements

1. Releasing past semesters’ final exams

2. Email for makeup final should be out today

Recap on fully connected NN

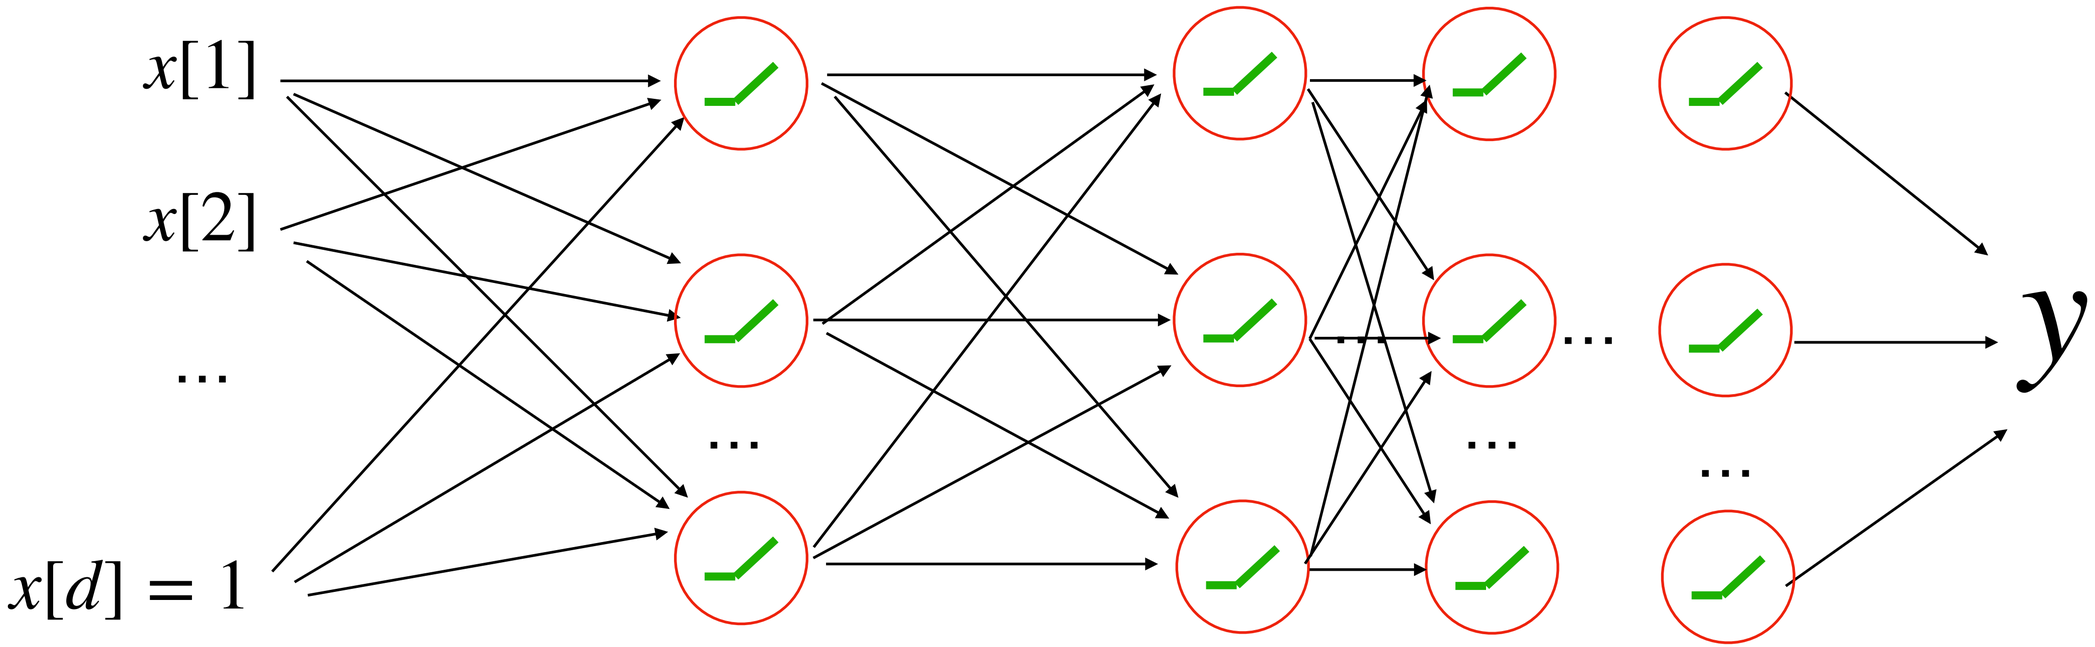

Not very scalable when d is large (e.g., when d is a million)

Recap on fully connected NN

Also not translation invariant

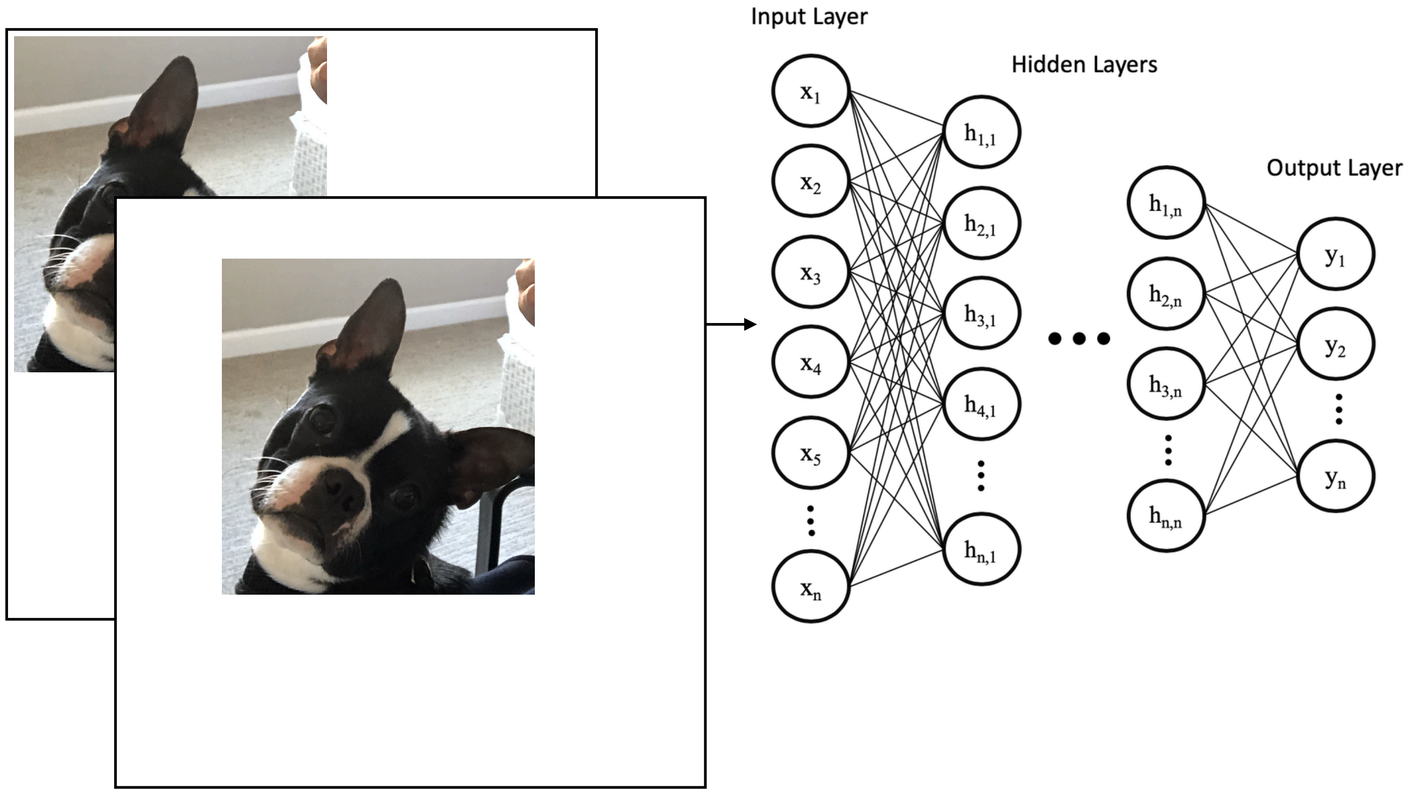

Objective today

Understand the convolution operator and the Convolution network
(designed for dealing w/ image inputs)

Outline today

1. Edge detector and convolution

2. Convolution layer and a pooling layer

3. Case study on LeNet (ResNet)

Edge detector

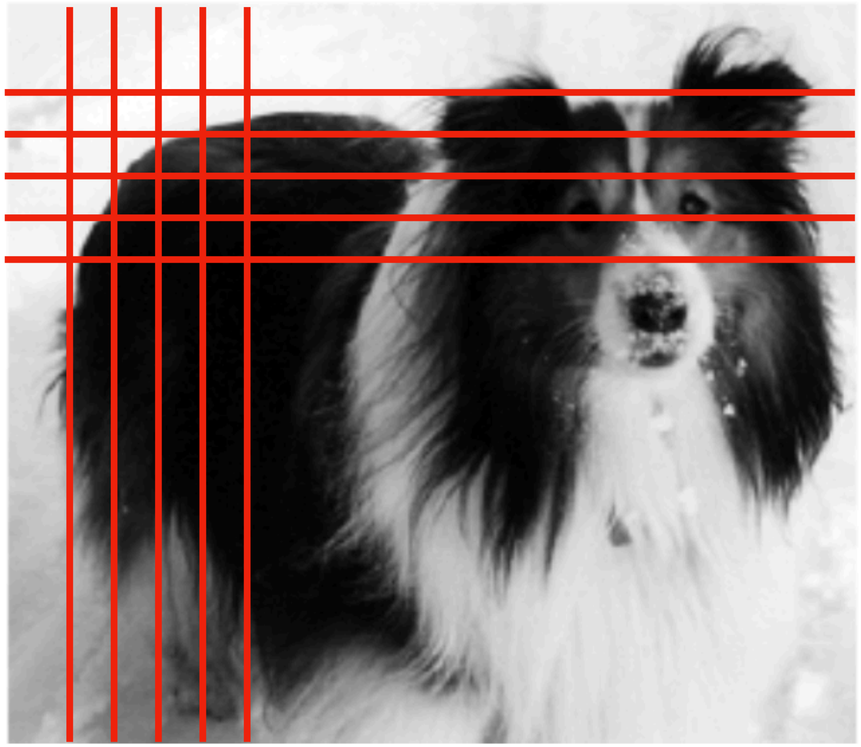

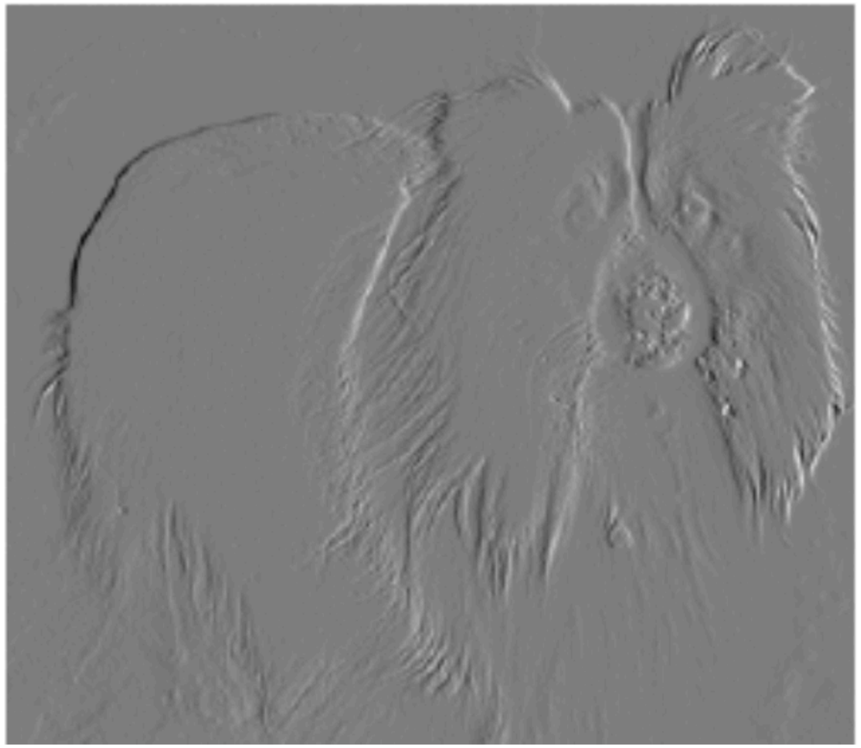

Implementing vertical Edge detector w/ convolution

Kernel 1x2 matrix

-1	**1**

Other Edge detector

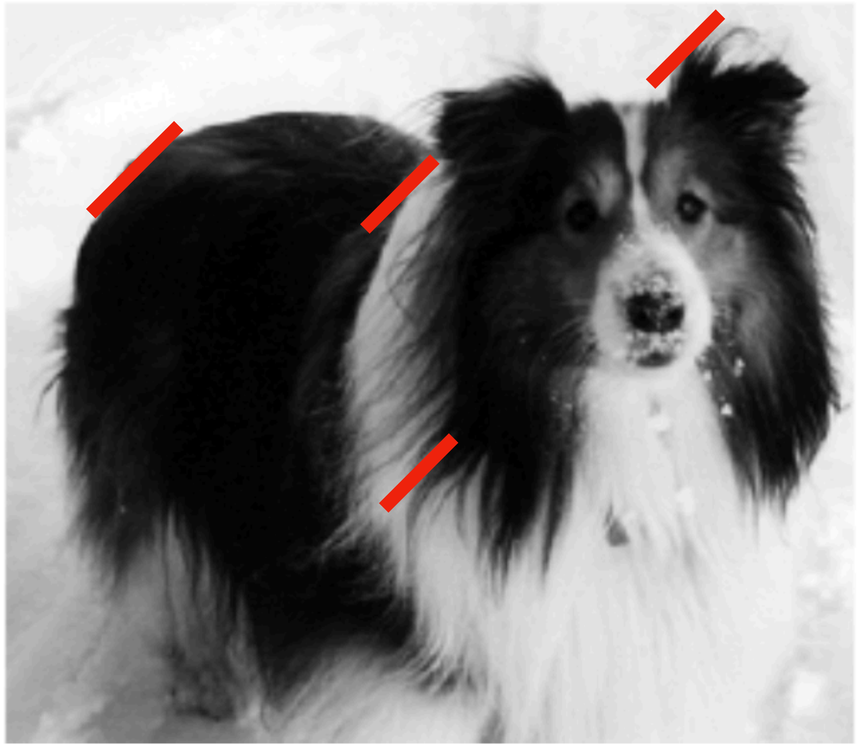

Kernel 3x3 matrix

0	0	1
0	1	0
1	0	0

More examples

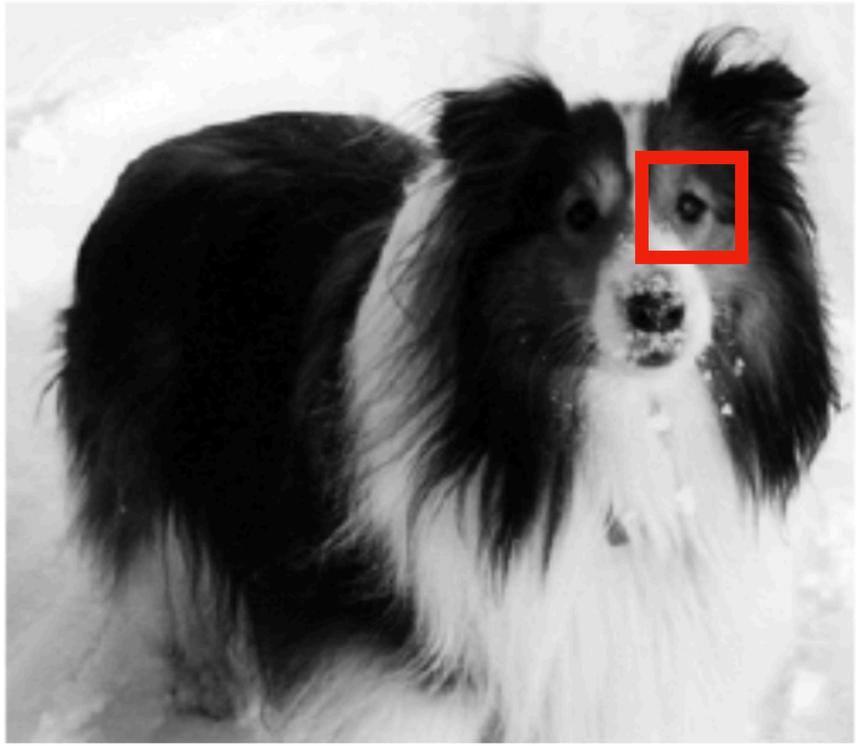

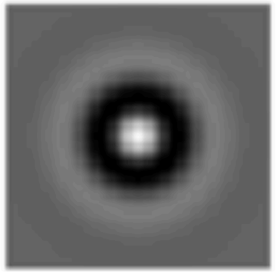

The Filter bank

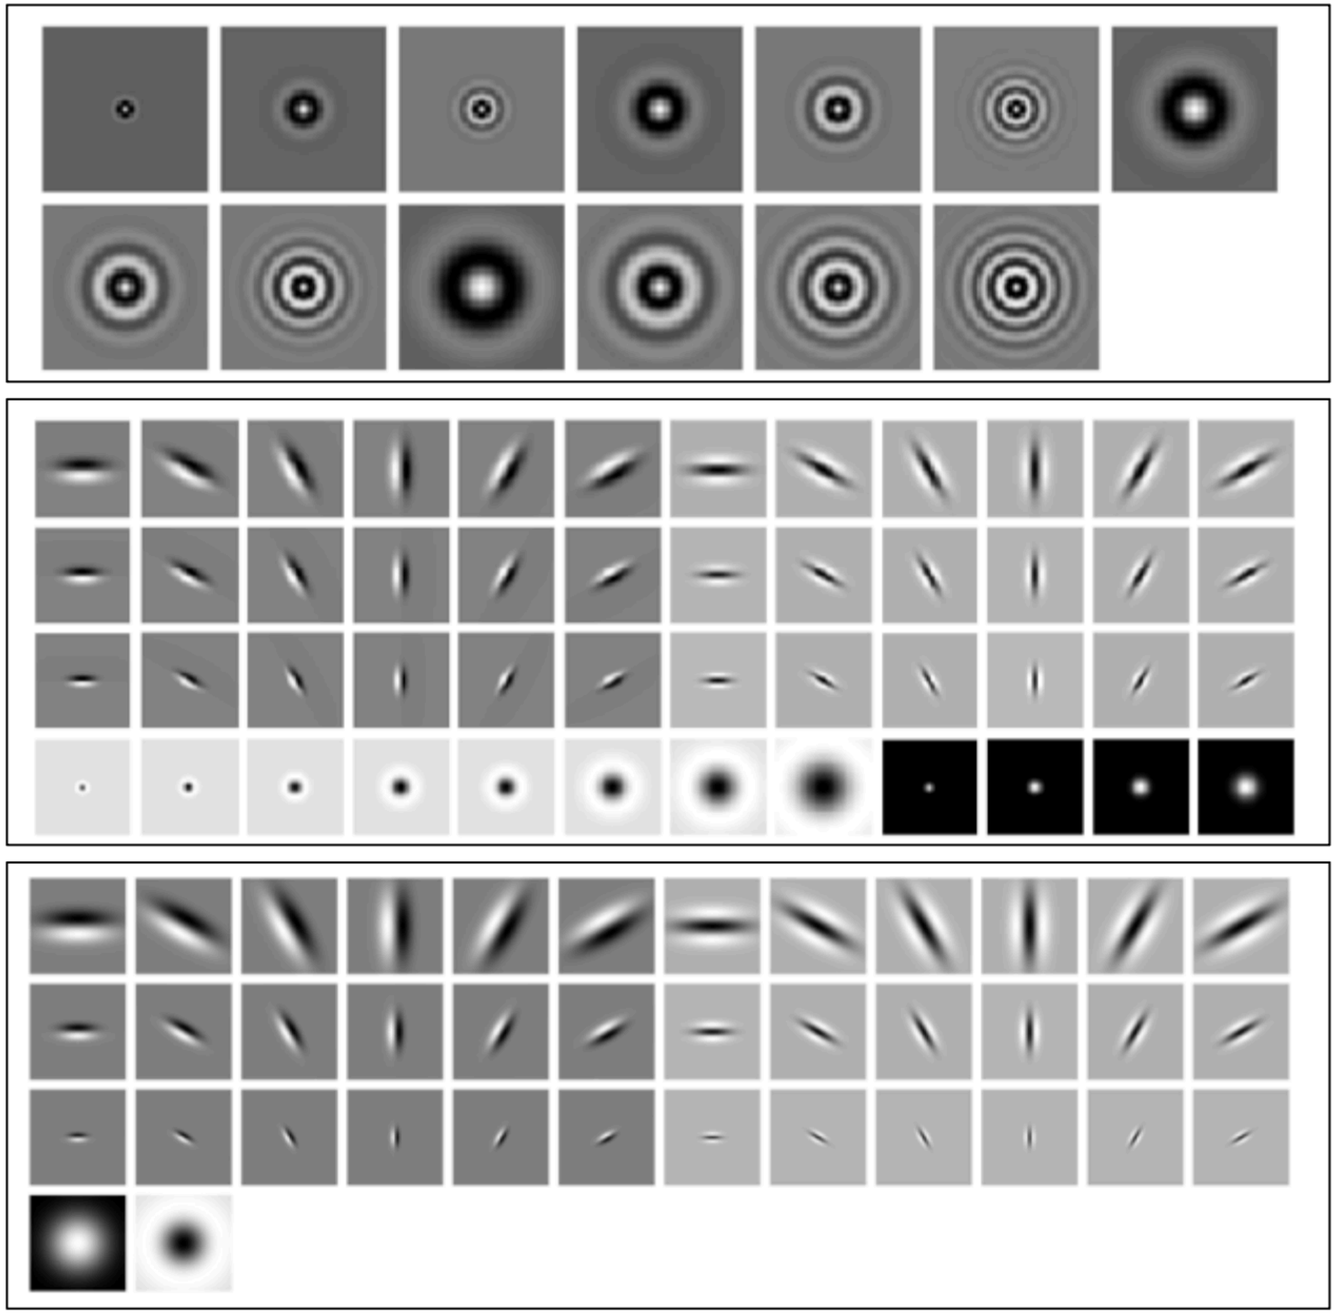

Summary of Convolution operator

Image (2d matrix)

a	**b**	**c**	**d**
e	**f**	**g**	**h**
i	**j**	**k**	**l**

Kernel 2x2 matrix

W	**X**
Y	**Z**

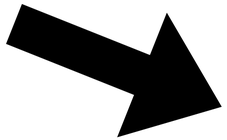

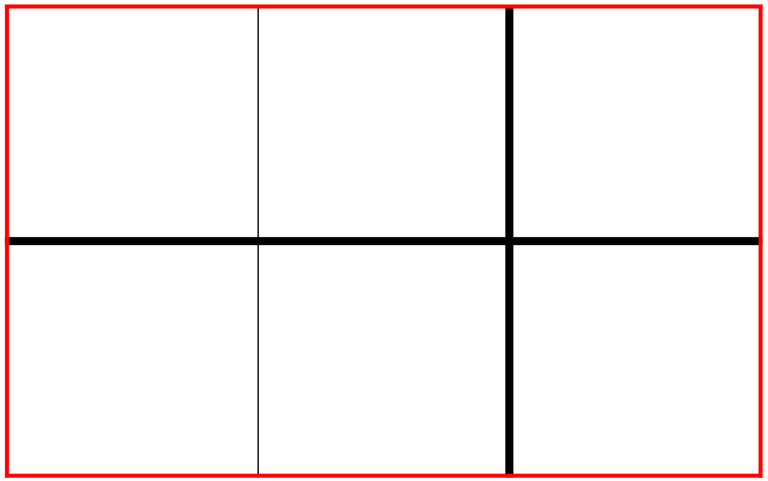

Visualization of convolution

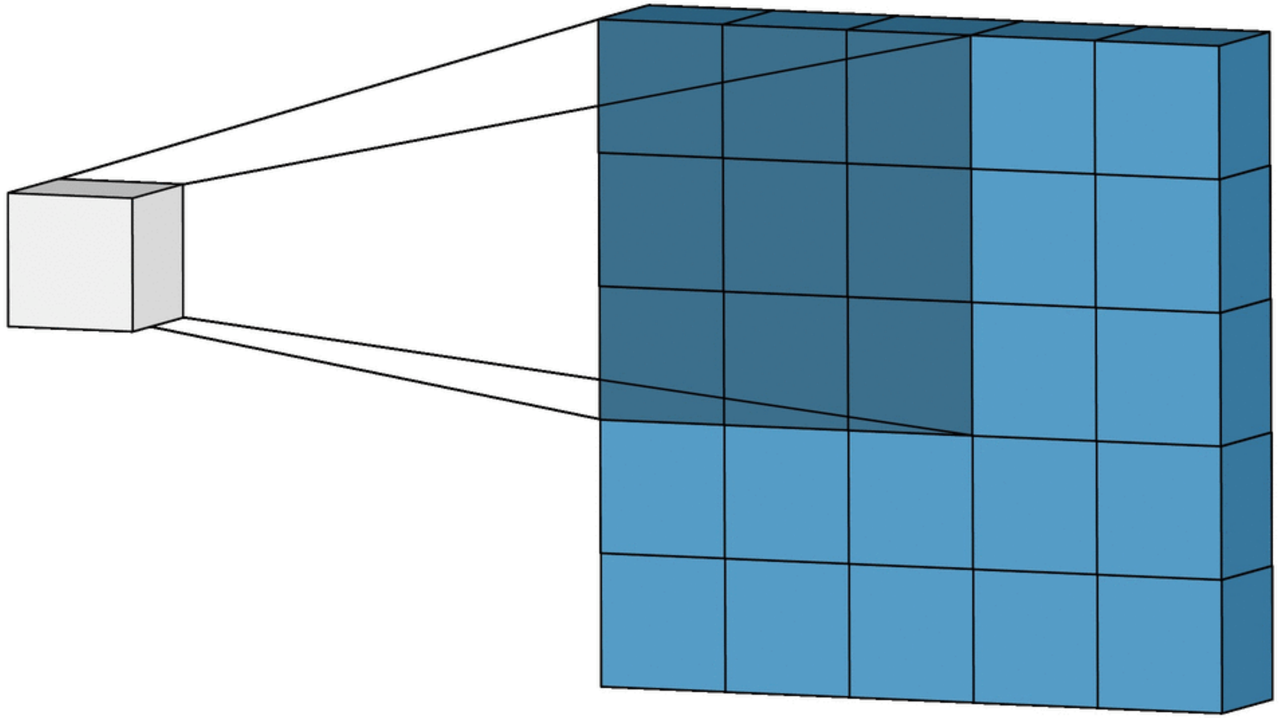

Q: if the image is $d \times d$, and kernel is $k \times k$, what is the dim of the output matrix?

Convolution over volumes (3d tensor)

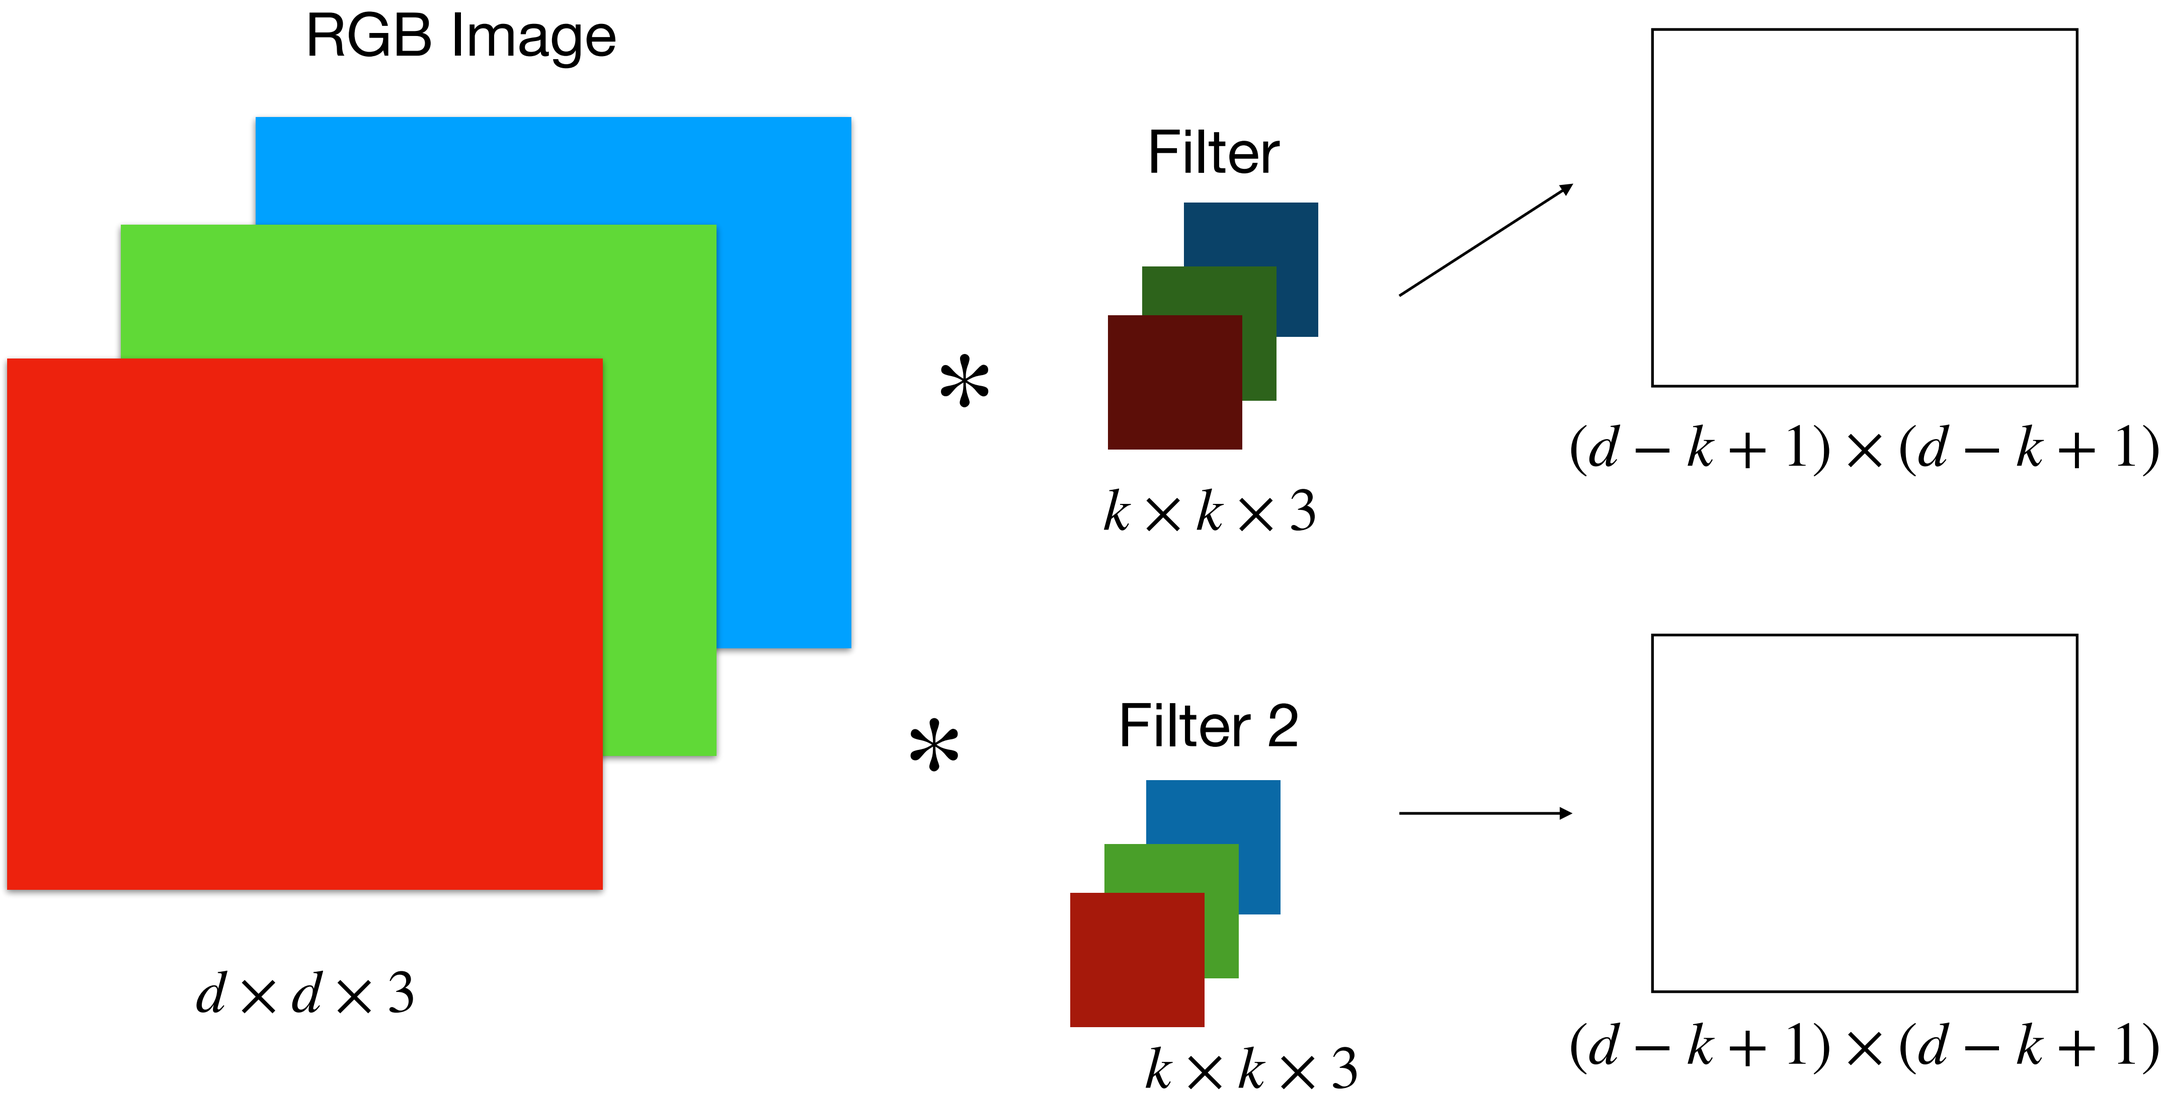

Key question

Can we learn these detectors / filters in an end-to-end fashion?

Outline today

1. Edge detector and convolution

2. Convolution layer and a pooling layer

3. Case study on LeNet (ResNet)

Building a convolution layer

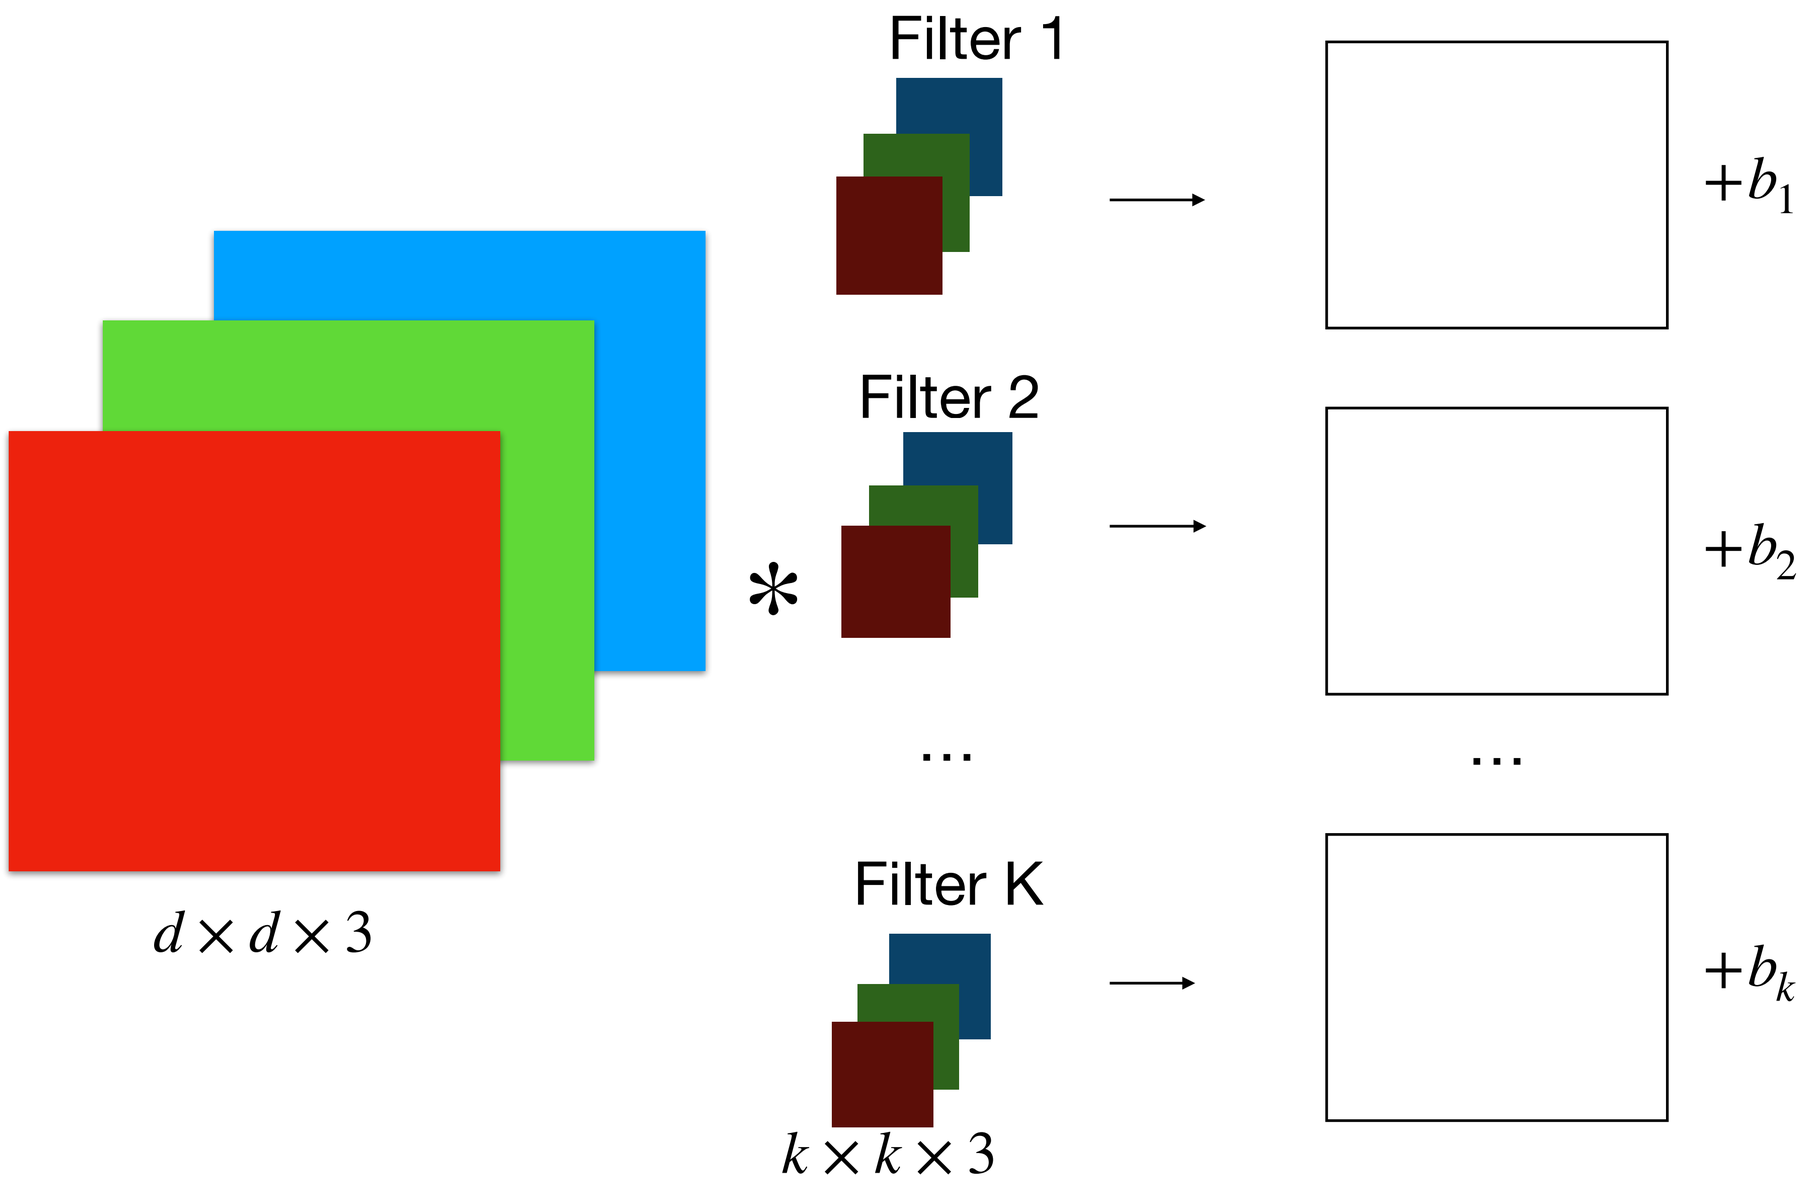

Pooling layer

We use a pooling layer to downsize the inputs

e.g., Max pooling (2x2 filter and stride 2)

e.g.,some Output of a convolution layer

1	**3**	**0**	**6**
5	**4**	**12**	**2**
7	**1**	**9**	**0**
4	**3**	**1**	**8**

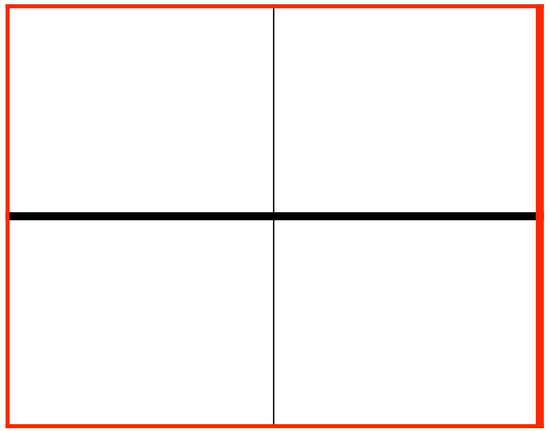

Outline today

1. Edge detector and convolution

2. Convolution layer and a pooling layer

3. Case study on LeNet

LeNet 5

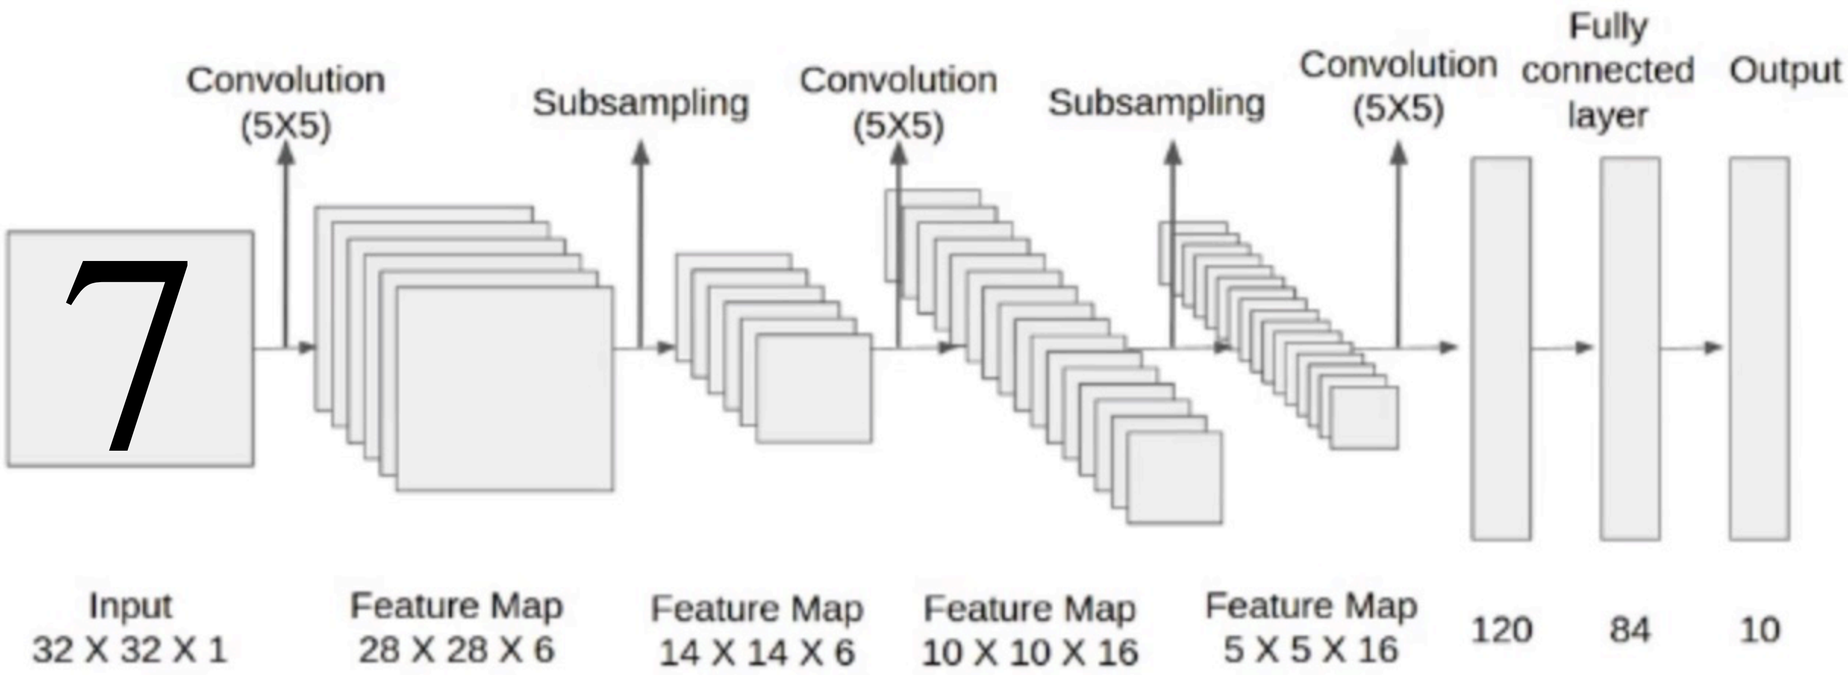
Convolution
(5X5)
Subsampling
Convolution
(5X5)
Subsampling
Convolution
(5X5)
Fully
connected
layer
Output
7
Input
32 X 32 X 1
Feature Map
28 X 28 X 6
Feature Map
14 X 14 X 6
Feature Map
10 X 10 X 16
Feature Map
5 X 5 X 16
120
84
10

ResNet

(The last reading quiz in on the classic ResNet paper!)

Summary for today

Convolutional neural network works well for images where pixels have strong local spatial correlations

Limitations:
convolution cannot capture global information (correlation among very distant pixels);
Fine-grained details may lost during pooling.