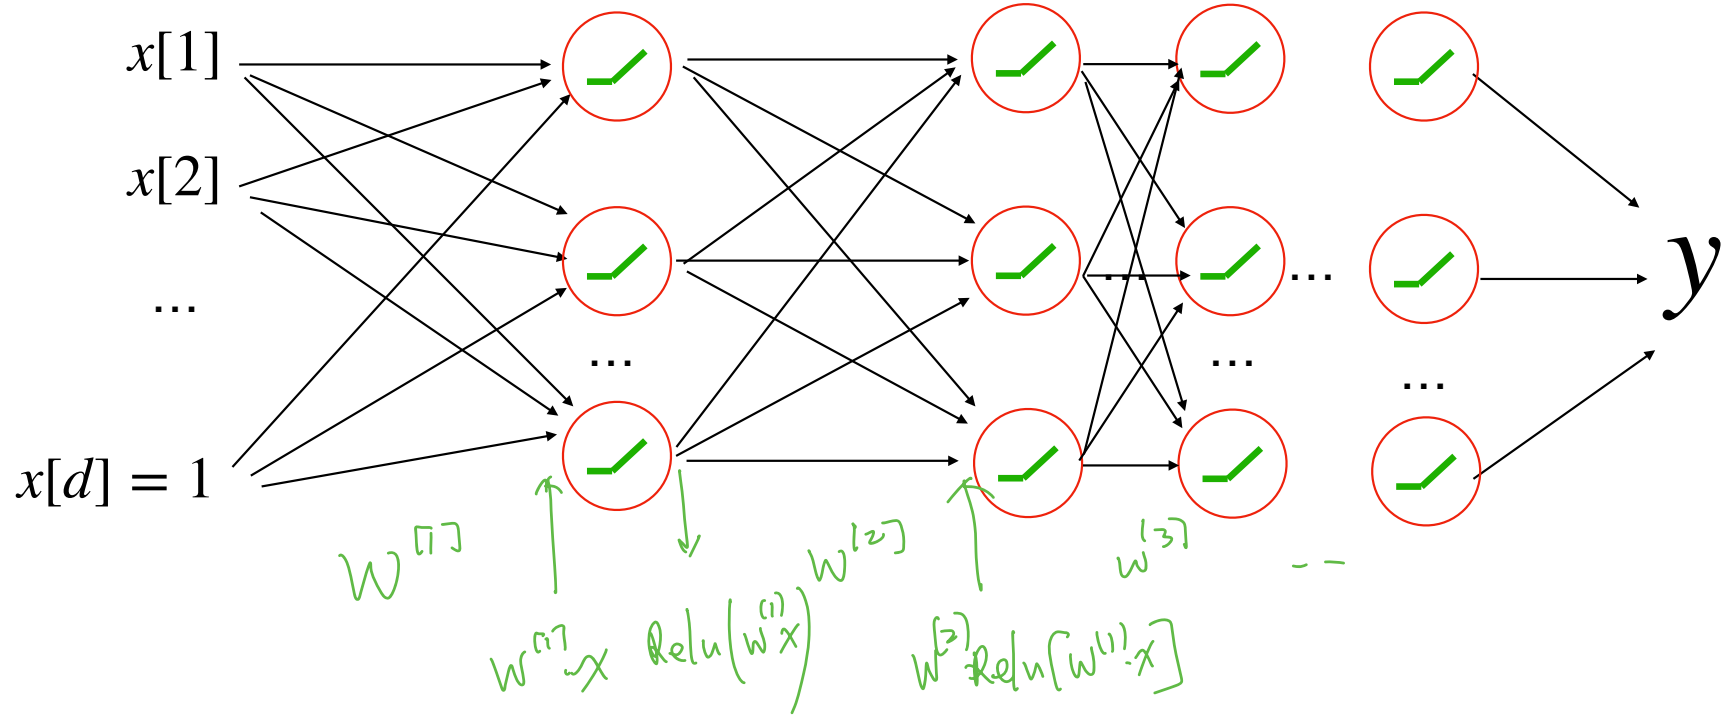


Convolutional Neural Network


Announcements

1. Releasing past semesters' final exams
2. Email for makeup final should be out today

Recap on fully connected NN



Recap on fully connected NN





$\underline{W^{[1]}} \in \mathbb{R}^{d \times d}$

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



Recap on fully connected NN

Also not translation invariant



Recap on fully connected NN

Also not translation invariant



$$W = \begin{bmatrix} | & & | \\ w'_1 & \dots & w'_d \\ | & & | \end{bmatrix}$$

$$X = \begin{bmatrix} 1 \\ 1 \\ 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix}$$

$$WX = \begin{bmatrix} w_1 + w_2 \end{bmatrix}$$

$$X' = \begin{bmatrix} 0 \\ 0 \\ 1 \\ 0 \\ \vdots \\ 0 \end{bmatrix}$$

$$WX' = \begin{bmatrix} w_3 + w_4 \end{bmatrix}$$


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



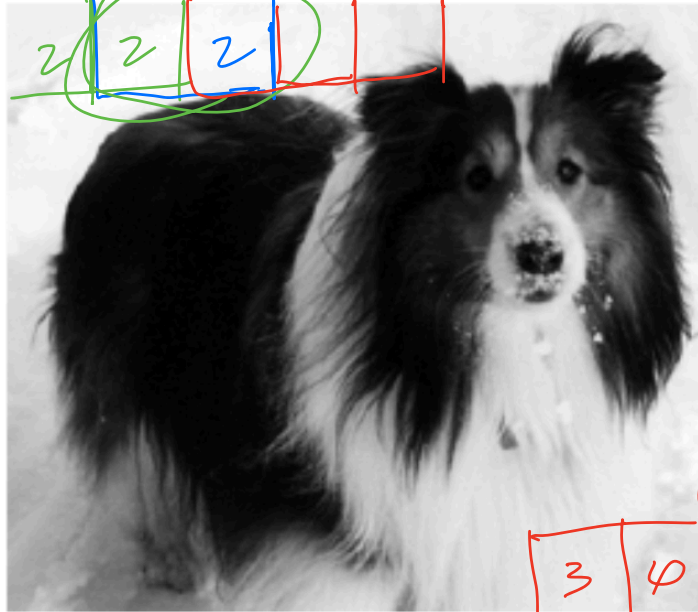
Edge detector



Implementing vertical Edge detector w/ convolution

$$2 \times (-1) + 2 \times 1 = 0$$

$$\rightarrow 2 \times (-1) + 2 \times 1 = 0$$



Kernel 1x2 matrix

| | |
|--------------|--------------|
| -1 | 1 |
| \leftarrow | \leftarrow |

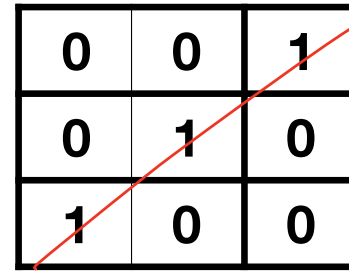
| | |
|---|---|
| 3 | 4 |
|---|---|

$$3 \times (-1) + 4 \times 1 = 1$$


Other Edge detector

Kernel 3x3 matrix

| | | |
|----------|----------|----------|
| 0 | 0 | 1 |
| 0 | 1 | 0 |
| 1 | 0 | 0 |

A 3x3 grid of cells with black borders. The cells contain the following values: top row [0, 0, 1], middle row [0, 1, 0], bottom row [1, 0, 0]. A red diagonal line starts from the bottom-left corner of the grid and extends upwards and to the right, passing through the top-right corner.


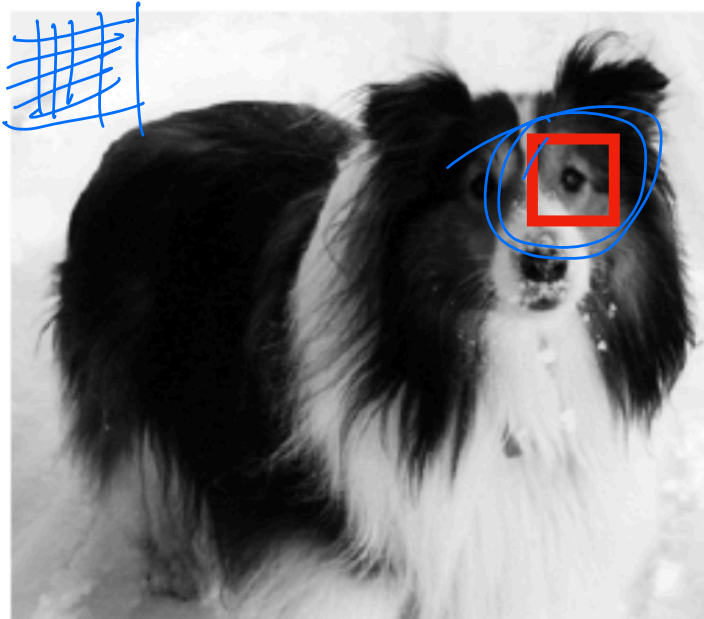
Other Edge detector






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 |




| | | |
|-----------------------|----------------------|----------------------|
| $aw+bx$ $+ ey+ fz$ | $bw+cx$ $+ fy+gz$ | $cw+dx$ $+ gy+hz$ |
| $we+xf$ $+ iy+jz$ | | |

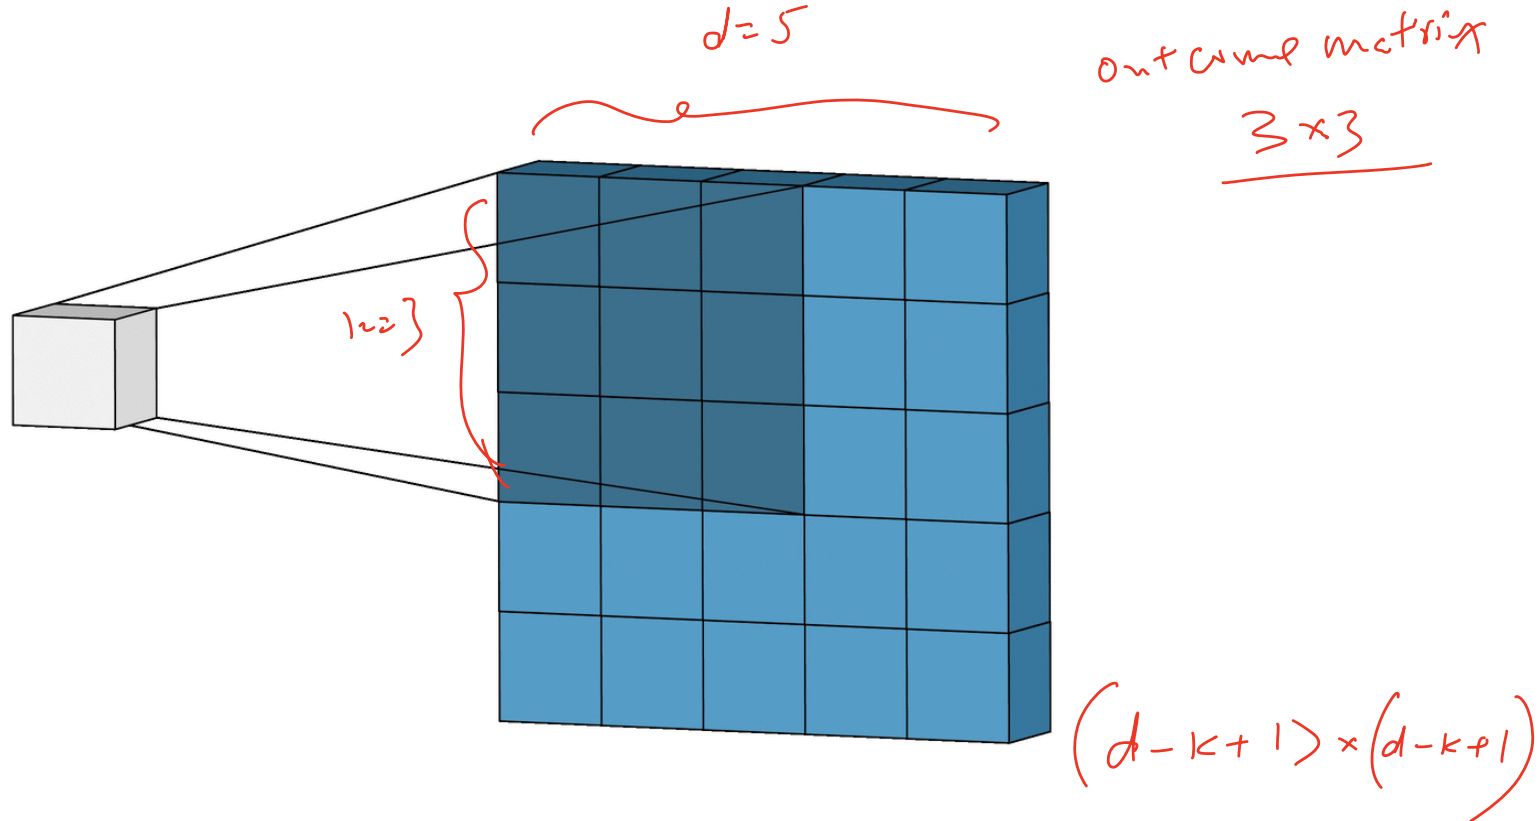
Visualization of convolution



Visualization of convolution



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)

RGB Image



$d \times d \times 3$

\rightarrow # of channels

Convolution over volumes (3d tensor)

RGB Image



$d \times d \times 3$

*

Filter



$k \times k \times 3$

Convolution over volumes (3d tensor)

RGB Image

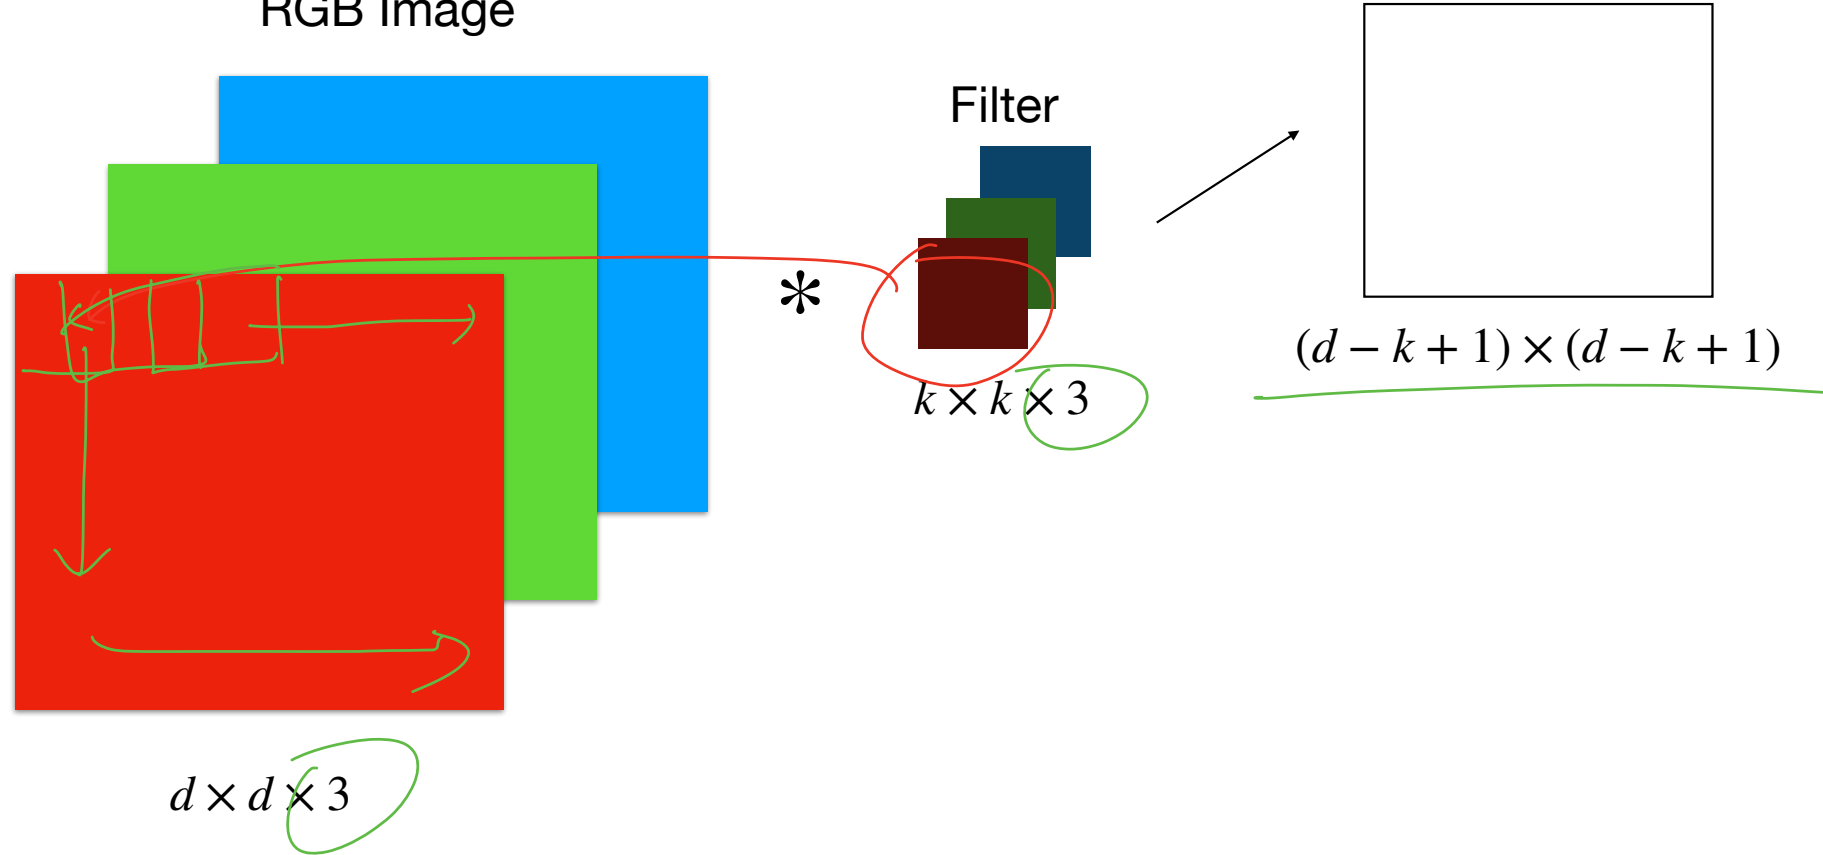
Filter

*

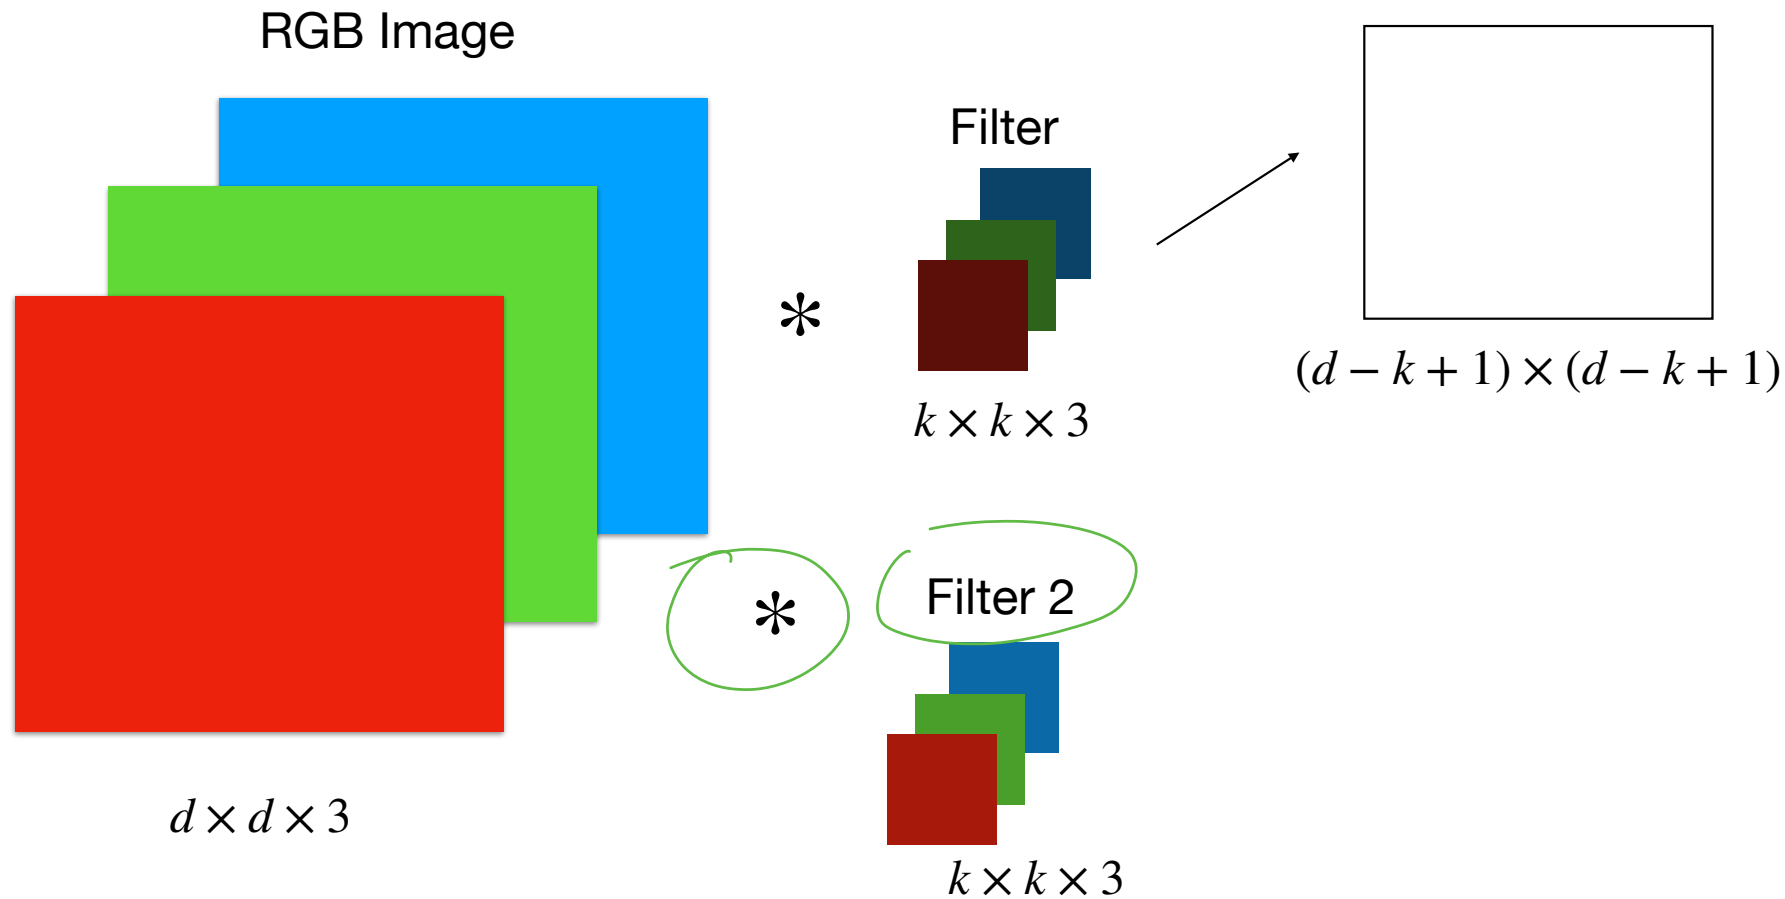
$k \times k \times 3$

$(d - k + 1) \times (d - k + 1)$

$d \times d \times 3$



Convolution over volumes (3d tensor)



Convolution over volumes (3d tensor)

RGB Image



$d \times d \times 3$

*

Filter





$k \times k \times 3$

*

Filter 2





$k \times k \times 3$




$(d - k + 1) \times (d - k + 1)$

identify edges



$(d - k + 1) \times (d - k + 1)$

identify circles



$(d - k + 1) \times (d - k + 1) \times 2$

Key question

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

Outline today

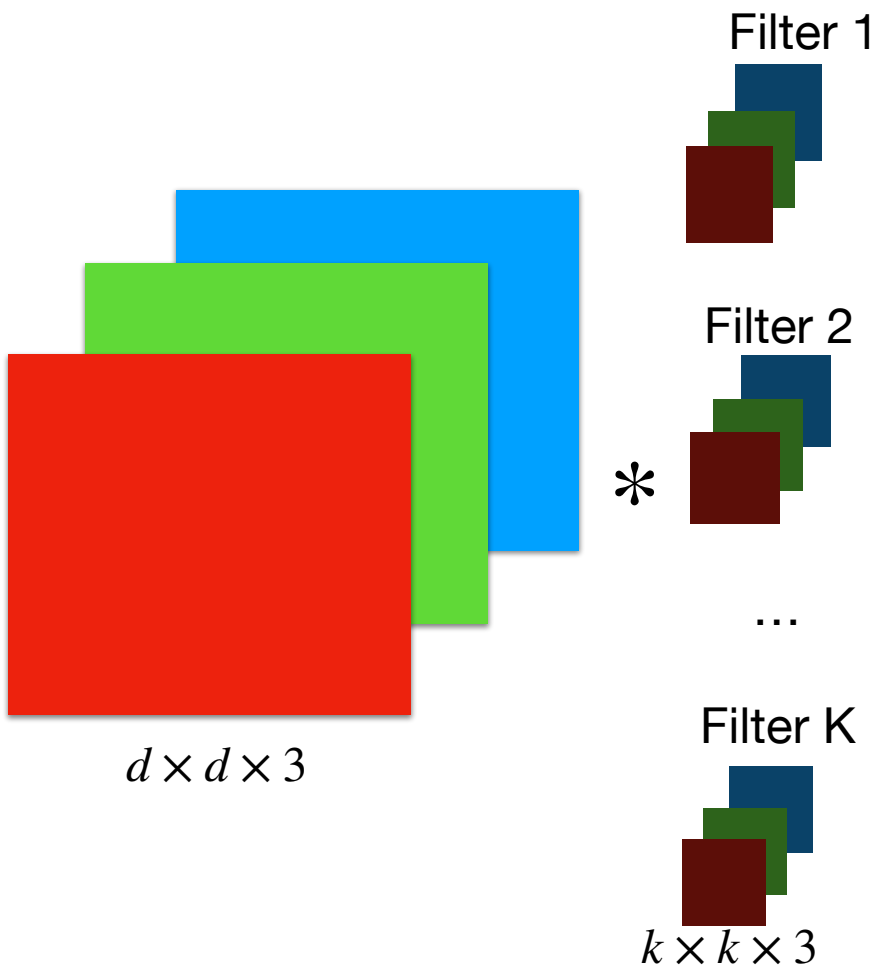
1. Edge detector and convolution
2. Convolution layer and a pooling layer
3. Case study on LeNet (ResNet)

Building a convolution layer

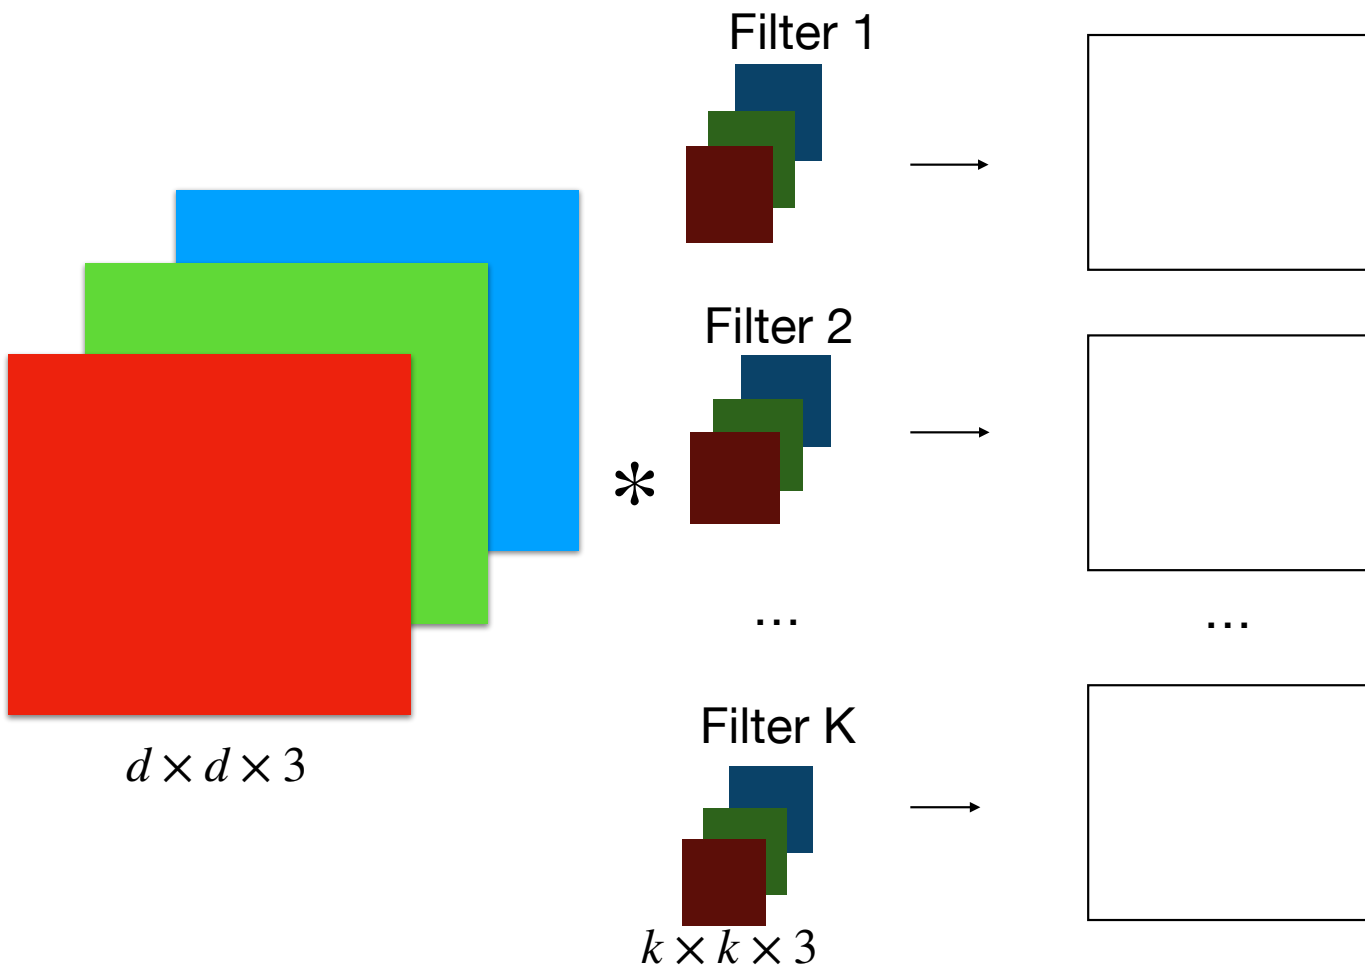


$d \times d \times 3$

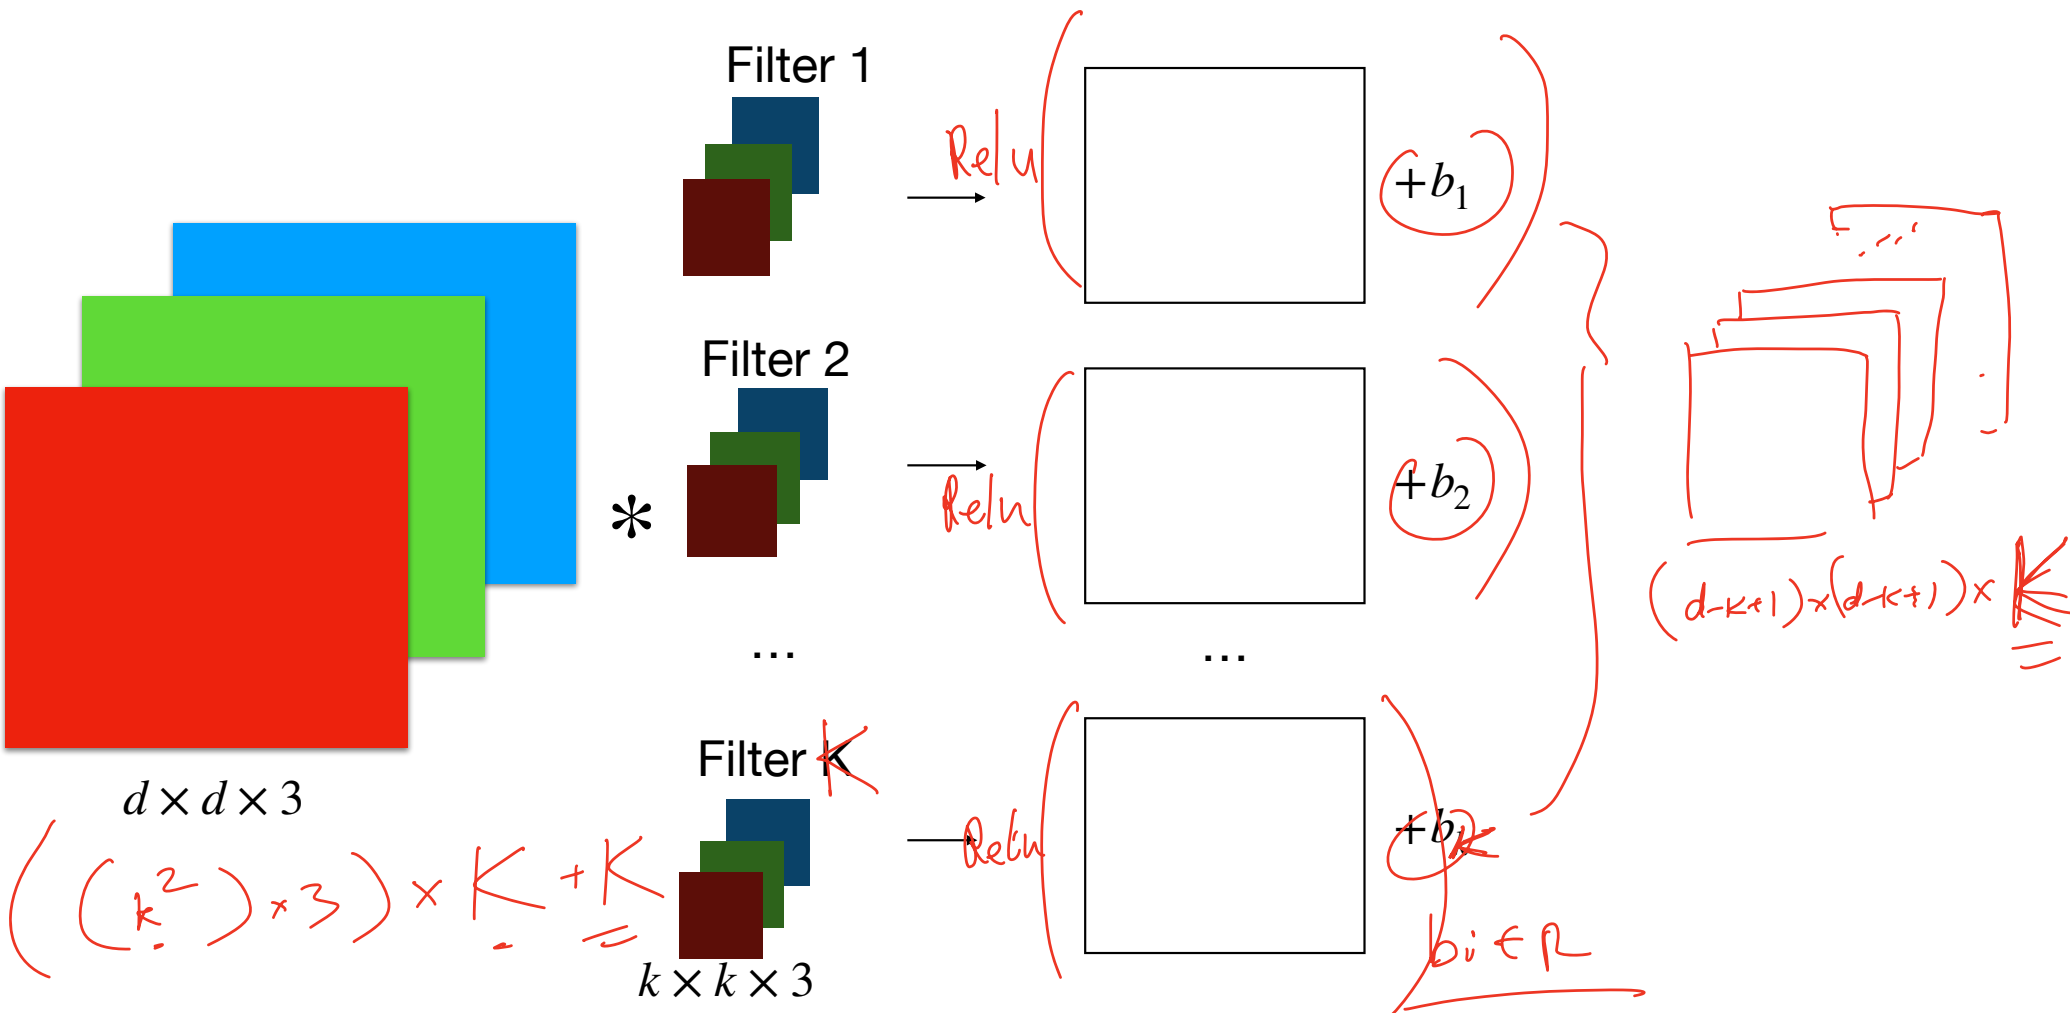
Building a convolution layer



Building a convolution layer




Building a convolution layer



Pooling layer

We use a pooling layer to downsize the inputs



Pooling layer

We use a pooling layer to downsize the inputs

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



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

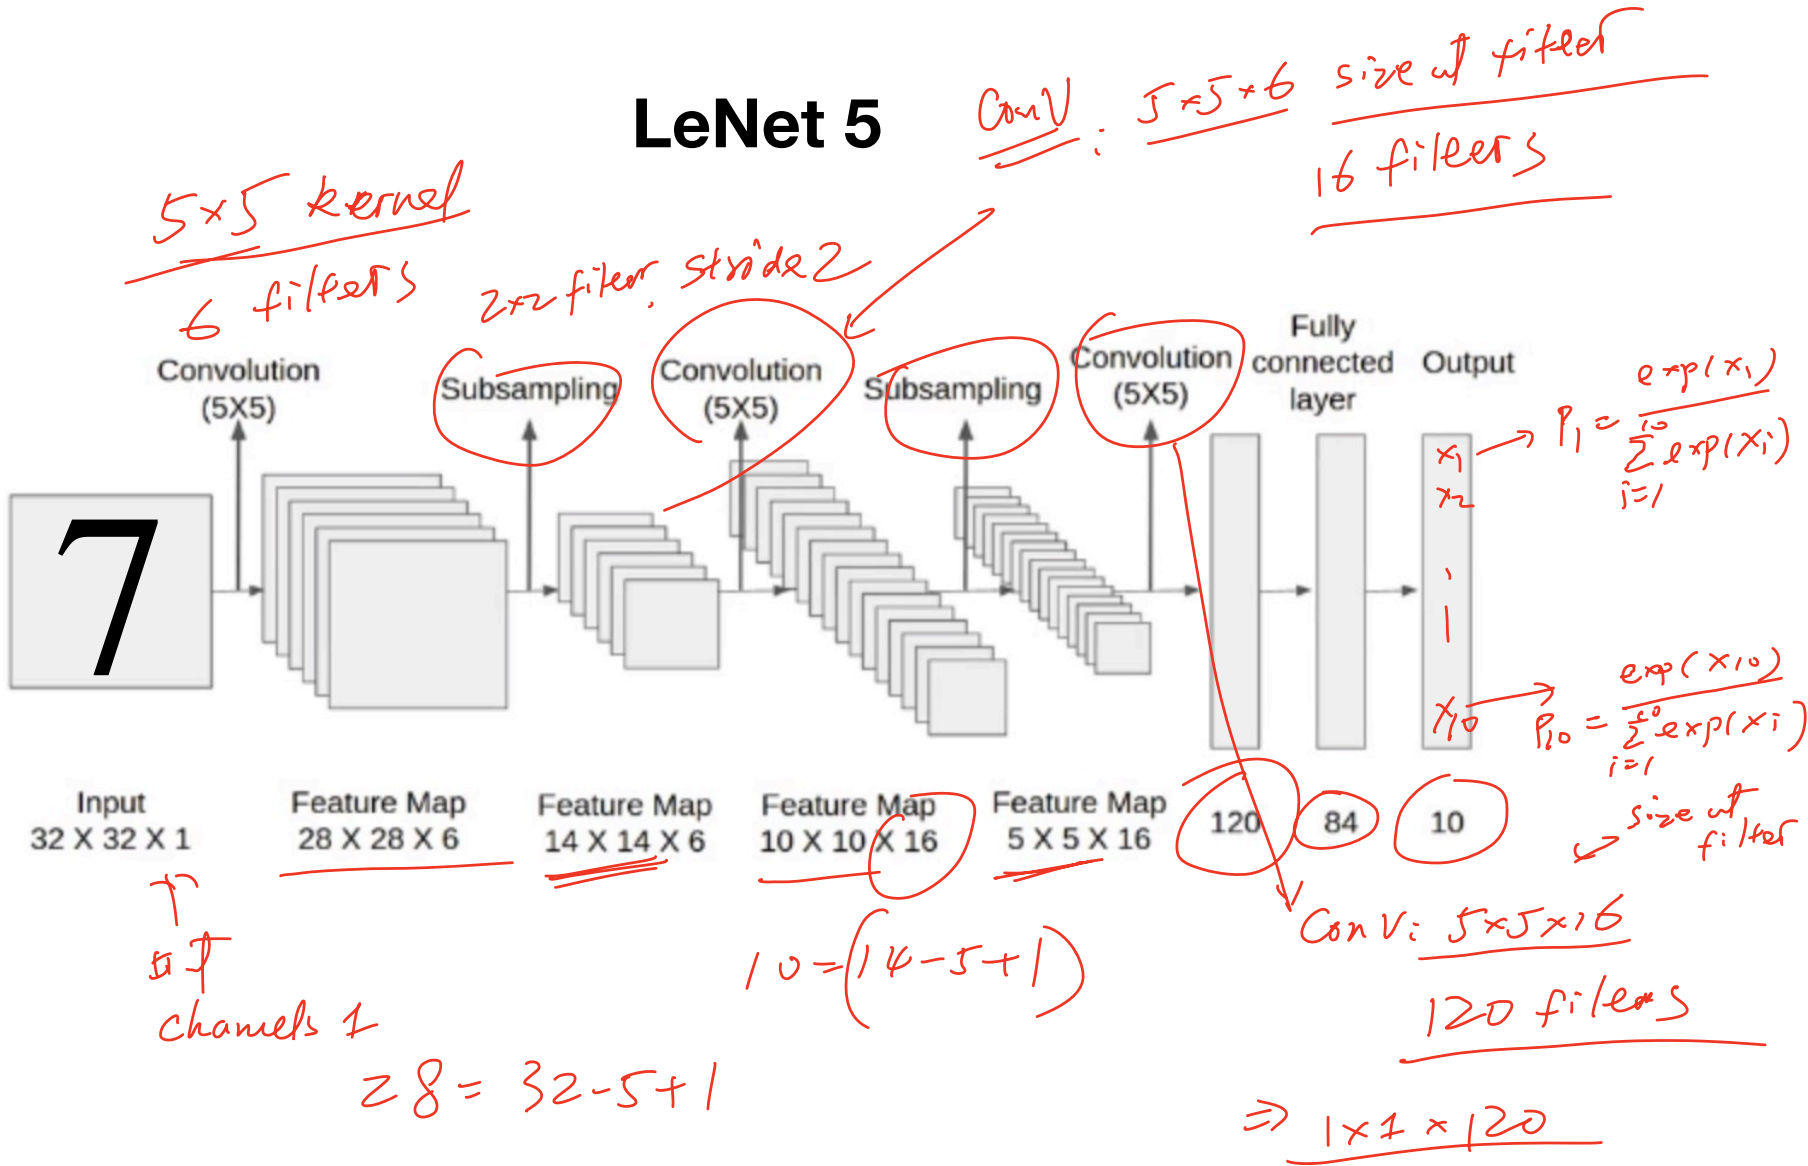
| | | | |
|---|---|----|---|
| 1 | 3 | 0 | 6 |
| 5 | 4 | 12 | 2 |
| 7 | 1 | 9 | 0 |
| 4 | 3 | 1 | 8 |

| | |
|---|----|
| 5 | 12 |
| 7 | 9 |


Outline today

1. Edge detector and convolution
2. Convolution layer and a pooling layer
3. Case study on LeNet

LeNet 5



ResNet



$$\frac{\partial l}{\partial w_0} = \frac{\partial l}{\partial z_{i,n}} \frac{\partial z_{i,n}}{\partial z_{i,n-1}} \dots \frac{\partial z_1}{\partial w_0} \leftarrow \text{Vanishing Gradient}$$

$< 1 \quad < 1 \quad < 1$

(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

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 be lost during pooling.