

# HW2 Recitation

10-417/617

# Statistics and Probability

## Conditional Probability

$$P(A|B)P(B) = P(A, B)$$

## Bayes' Theorem

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)}$$



Marginalization

$$P(A) = \int_B P(A, B)$$

$$P(A) = \sum_B P(A, B)$$

# Graphical Models

- joint  $p(x, y)$  or conditional  $p(y | x)$  probability distribution
- represented as  $G=(V,E)$
- Enables us to encode relationships between a set of random variables
- There are two types of graphical models : Directed Graphical Model and Undirected Graphical Model
- Directed edges gives a causality relationships
- Undirected edges give correlations between variable




# Direct Graphical Model

- Assume a directed, acyclic graphical model  $G=(V,E)$  and  $E \subset V \times V$

$$p(\mathbf{Y} = \mathbf{y}) = \prod_{i \in \mathcal{V}} p(y_i \mid \mathbf{y}_{\text{pa}_G(i)})$$

where  $y_{\text{pa}_G(i)}$  is conditional probability distribution on the parents of node  $i$



# Undirected Graphical Model


- An undirected graphical model  $G=(V,E)$  is called Markov Random Field (MRF) if two nodes are conditionally independent whenever they are not connected.

$$p(Y_i | Y_{V \setminus \{i\}}) = p(Y_i | Y_{N(i)})$$

where  $N(i)$  is the neighbor of node  $i$

$$Y_i \perp\!\!\!\perp Y_{V \setminus \text{cl}(i)} | Y_{N(i)},$$

where  $\text{cl}(i)=N(i) \cup \{i\}$  is the closed neighborhood of  $i$ .






# Convolutions


Some figures adapted from Simon Lucey's 16-720b slides and  
Introduction to Deep Learning 11-785

# What is a (Discrete) Convolution?



Concatenation of inner products of filter and receptive fields of signal.

Unlike an arbitrary affine transform, convolutions preserve locality.



\*

X

H

$$(x * k)_{ij} = \sum_{pq} x_{i+p, j+q} k_{r-p, r-q}$$

These would be + in cross-correlation

=



Y

# Why do Convolutions Work?




$\mathbf{X}$

Natural signals are locally smooth!

In natural images, neighboring pixels tend to be highly correlated.


Convolutions exploit this local correlation to generalize better than fully connected networks.

Convolutional layers DO NOT learn random noise better than fully connected layers!




# Computer Vision B.C. (Before Conv-Nets)

Filters were hand-designed to extract features for the intended task!



How do we recover edges?


Canny Algorithm



# Computer Vision A.D. (After Deep Learning)



Now, optimal filters are *learned* through back-propagation!





Feature visualization of convolutional net trained on ImageNet from [Zeiler & Fergus 2013]

How do we do convolutions?

# What is an image?

~~A thousand words.~~ A Matrix  $\mathbf{I}$  of dimensions  $(\mathbf{M}, \mathbf{N})$  with  $\mathbf{I}[\mathbf{i}][\mathbf{j}] = \text{intensity}(\text{pixel}(\mathbf{i}, \mathbf{j}))$



157	153	174	168	150	162	129	161	172	161	165	164
155	182	163	74	75	62	53	17	110	210	180	154
180	180	50	14	54	6	10	93	48	106	159	181
206	109	5	124	131	111	120	204	166	15	56	180
194	68	137	251	237	239	239	228	227	87	71	201
172	105	207	233	233	214	220	239	228	98	74	206
188	88	179	209	185	215	211	158	139	75	20	169
189	97	165	84	10	168	134	11	31	62	22	148
199	168	191	193	158	227	178	143	182	106	36	190
205	174	155	252	236	231	149	178	228	43	95	234
190	216	116	149	236	187	85	150	79	38	218	241
190	224	147	108	227	210	127	102	36	101	255	224
190	214	173	66	103	143	96	50	2	109	249	215
187	196	235	75	1	81	47	0	6	217	255	211
183	202	237	145	0	0	12	108	200	138	243	236
195	206	123	207	177	121	123	200	175	13	96	218

157	153	174	168	150	162	129	161	172	161	165	164
155	182	163	74	75	62	53	17	110	210	180	154
180	180	50	14	54	6	10	93	48	106	159	181
206	109	5	124	131	111	120	204	166	15	56	180
194	68	137	251	237	239	239	228	227	87	71	201
172	105	207	233	233	214	220	239	228	98	74	206
188	88	179	209	185	215	211	158	139	75	20	169
189	97	165	84	10	168	134	11	31	62	22	148
199	168	191	193	158	227	178	143	182	106	36	190
205	174	155	252	236	231	149	178	228	43	95	234
190	216	116	149	236	187	85	150	79	38	218	241
190	224	147	108	227	210	127	102	36	101	255	224
190	214	173	66	103	143	96	50	2	109	249	215
187	196	235	75	1	81	47	0	6	217	255	211
183	202	237	145	0	0	12	108	200	138	243	236
195	206	123	207	177	121	123	200	175	13	96	218

# Components of a CNN

$A_{1,1}$	$A_{1,2}$	$A_{1,3}$	$A_{1,4}$
$A_{2,1}$	$A_{2,2}$	$A_{2,3}$	$A_{2,4}$
$A_{3,1}$	$A_{3,2}$	$A_{3,3}$	$A_{3,4}$
$A_{4,1}$	$A_{4,2}$	$A_{4,3}$	$A_{4,4}$

Input -  $\mathbf{A}$

$W_{1,1}$	$W_{1,2}$
$W_{2,1}$	$W_{2,2}$

Kernel -  $\mathbf{W}$

$B_{1,1}$
-----------

Bias -  $\mathbf{B}$

$Z_{1,1}$	$Z_{1,2}$	$Z_{1,3}$
$Z_{2,1}$	$Z_{2,2}$	$Z_{2,3}$
$Z_{3,1}$	$Z_{3,2}$	$Z_{3,3}$

Output -  $\mathbf{Z}$

$$\mathbf{Z} = (\mathbf{A} \otimes \mathbf{W}) + \mathbf{B}$$



# Convolution steps

Essentially element-wise (Hadamard) multiplications and summations

# Convolution steps

Essentially element-wise (Hadamard) multiplications and summations

$A_{1,1}$	$A_{1,2}$	$A_{1,3}$	$A_{1,4}$
$A_{2,1}$	$A_{2,2}$	$A_{2,3}$	$A_{2,4}$
$A_{3,1}$	$A_{3,2}$	$A_{3,3}$	$A_{3,4}$
$A_{4,1}$	$A_{4,2}$	$A_{4,3}$	$A_{4,4}$

\*

$W_{1,1}$	$W_{1,2}$
$W_{2,1}$	$W_{2,2}$

+

$B_{1,1}$
-----------

=

$z_{1,1}$
-----------

$$z_{1,1} = (A_{1,1} * W_{1,1}) + (A_{1,2} * W_{1,2}) + (A_{2,1} * W_{2,1}) + (A_{2,2} * W_{2,2}) + B$$

# Convolution steps

Essentially element-wise (Hadamard) multiplications and summations

The diagram illustrates a 1D convolution operation. On the left, a 4x4 input grid  $A$  is shown with elements  $A_{1,1}$  through  $A_{4,4}$ . The second and third columns are highlighted in yellow. This grid is multiplied element-wise ( $*$ ) by a 2x2 kernel  $W$  (yellow grid) with elements  $W_{1,1}$ ,  $W_{1,2}$ ,  $W_{2,1}$ , and  $W_{2,2}$ . The result is added ( $+$ ) to a bias  $B_{1,1}$  (orange square). The final result is a 2x2 output grid  $Z$  (green grid) with elements  $Z_{1,1}$  and  $Z_{1,2}$ .

$$Z_{1,2} = (A_{1,2} * W_{1,1}) + (A_{1,3} * W_{1,2}) + (A_{2,2} * W_{2,1}) + (A_{2,3} * W_{2,2}) + B$$

# Convolution steps

Essentially element-wise (Hadamard) multiplications and summations

The diagram illustrates a 1D convolution operation. It shows the following components:

- Input Matrix A:** A 4x4 grid with elements  $A_{1,1}$  through  $A_{4,4}$ . The elements  $A_{1,3}$  and  $A_{2,3}$  are highlighted in yellow.
- Kernel Matrix W:** A 2x2 grid with elements  $W_{1,1}$  and  $W_{1,2}$  in the top row, and  $W_{2,1}$  and  $W_{2,2}$  in the bottom row. The entire kernel is highlighted in yellow.
- Bias B:** A single orange square containing the element  $B_{1,1}$ .
- Output Vector Z:** A 1x3 horizontal vector with elements  $Z_{1,1}$ ,  $Z_{1,2}$ , and  $Z_{1,3}$ , highlighted in green.

The operation is represented by the equation:

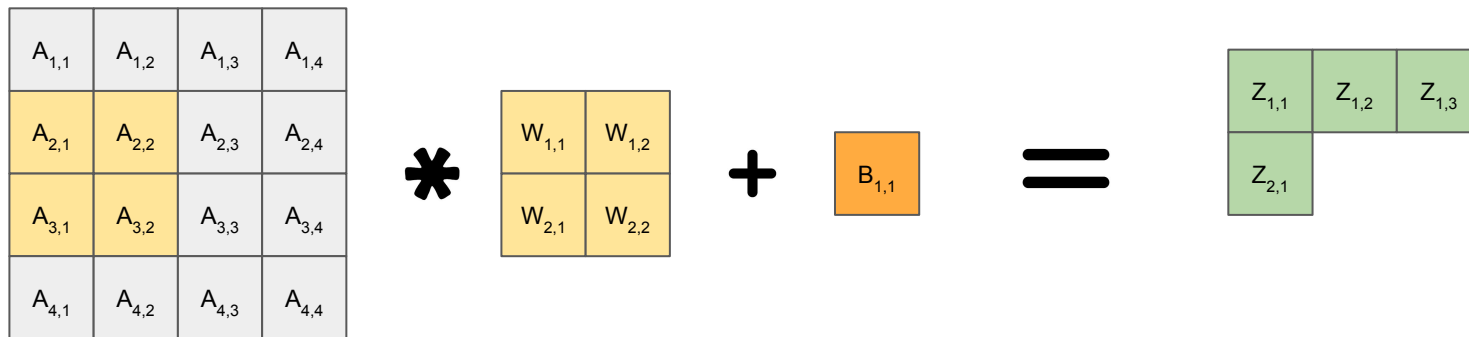
$$\begin{bmatrix} A_{1,1} & A_{1,2} & A_{1,3} & A_{1,4} \\ A_{2,1} & A_{2,2} & A_{2,3} & A_{2,4} \\ A_{3,1} & A_{3,2} & A_{3,3} & A_{3,4} \\ A_{4,1} & A_{4,2} & A_{4,3} & A_{4,4} \end{bmatrix} * \begin{bmatrix} W_{1,1} & W_{1,2} \\ W_{2,1} & W_{2,2} \end{bmatrix} + \begin{bmatrix} B_{1,1} \end{bmatrix} = \begin{bmatrix} Z_{1,1} & Z_{1,2} & Z_{1,3} \end{bmatrix}$$

The third element of the output vector,  $Z_{1,3}$ , is expanded as follows:

$$Z_{1,3} = (A_{1,3} * W_{1,1}) + (A_{1,4} * W_{1,2}) + (A_{2,3} * W_{2,1}) + (A_{2,4} * W_{2,2}) + B$$

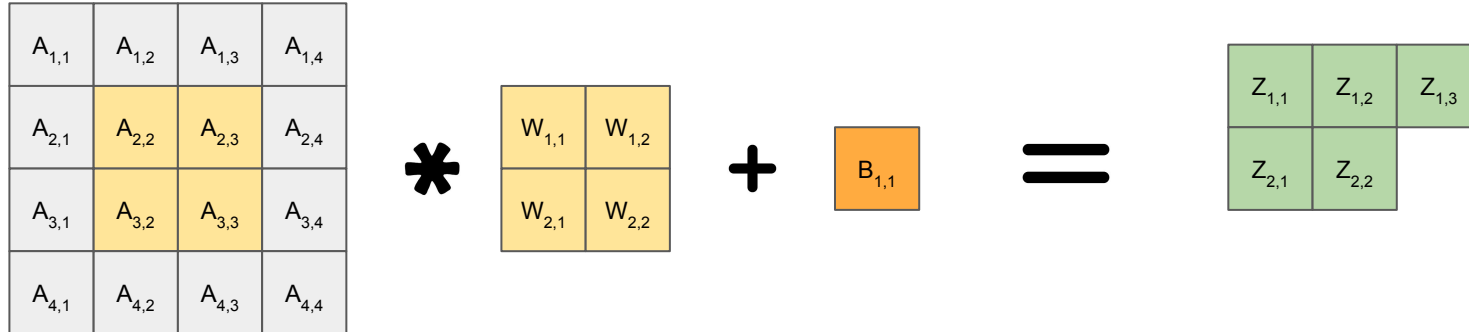
# Convolution steps

Essentially element-wise (Hadamard) multiplications and summations



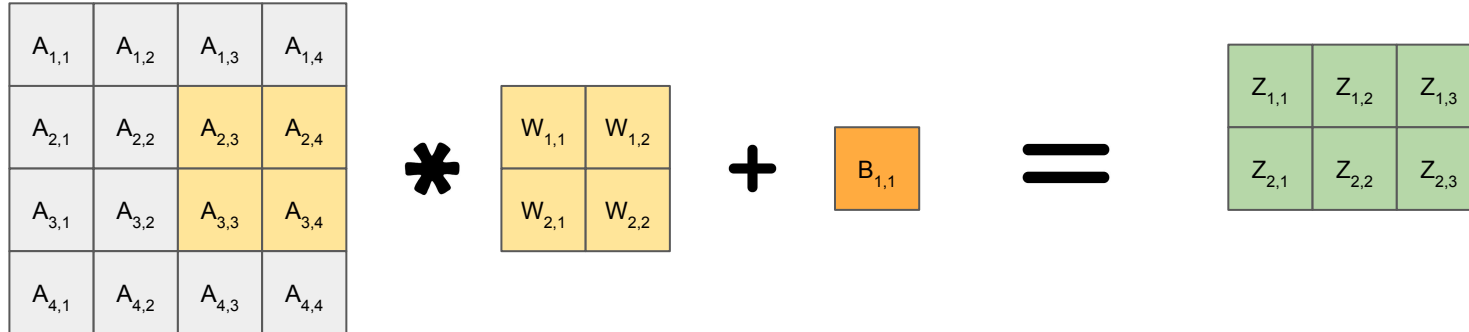
# Convolution steps

Essentially element-wise (Hadamard) multiplications and summations



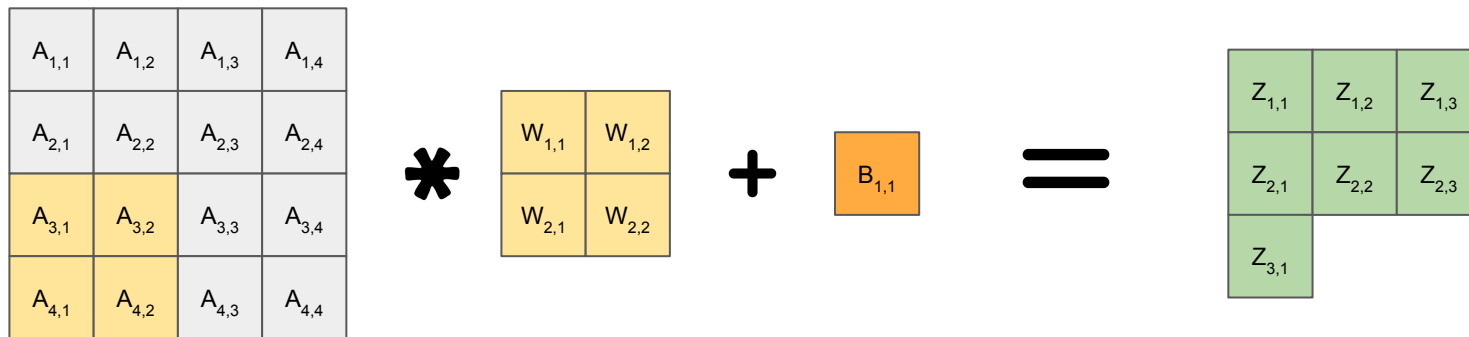
# Convolution steps

Essentially element-wise (Hadamard) multiplications and summations



# Convolution steps

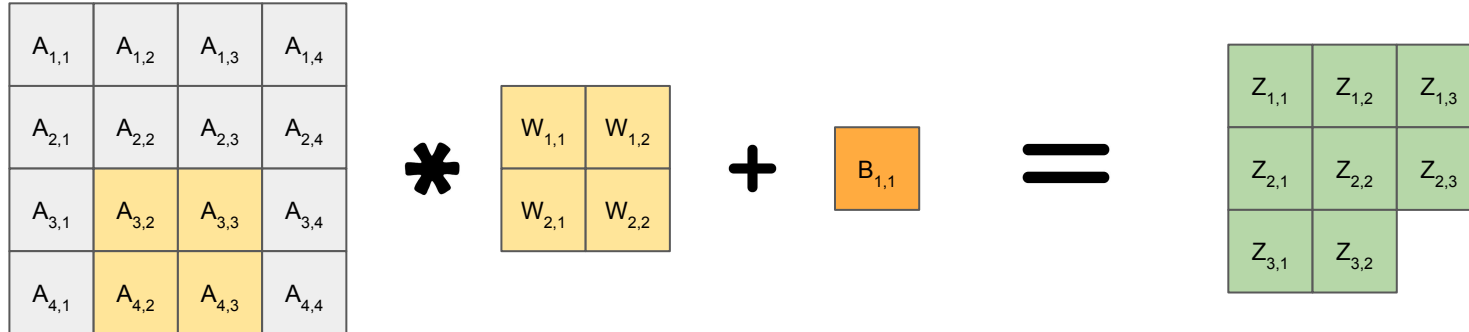
Essentially element-wise (Hadamard) multiplications and summations





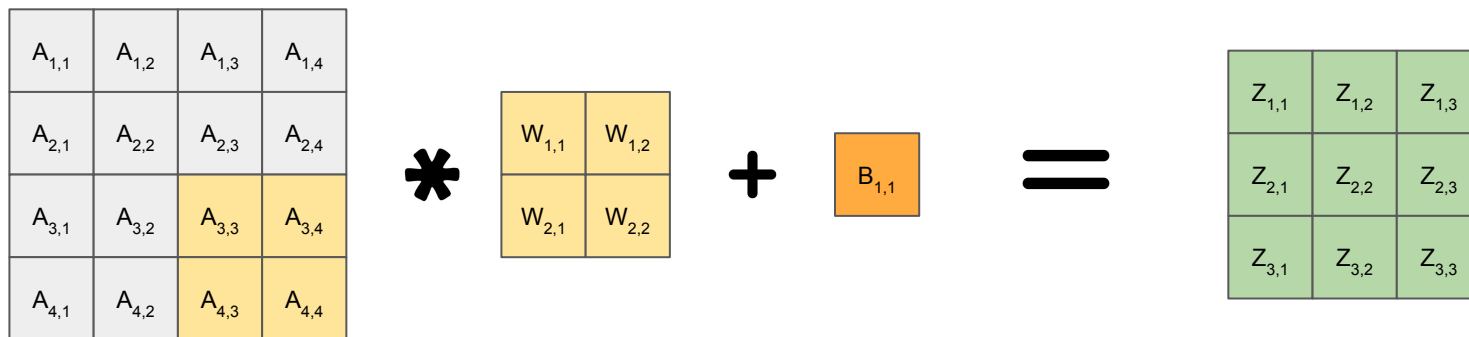
# Convolution steps

Essentially element-wise (Hadamard) multiplications and summations



# Convolution steps

Essentially element-wise (Hadamard) multiplications and summations



# Output Size

$A_{1,1}$	$A_{1,2}$	$A_{1,3}$	$A_{1,4}$
$A_{2,1}$	$A_{2,2}$	$A_{2,3}$	$A_{2,4}$
$A_{3,1}$	$A_{3,2}$	$A_{3,3}$	$A_{3,4}$
$A_{4,1}$	$A_{4,2}$	$A_{4,3}$	$A_{4,4}$



$Z_{1,1}$	$Z_{1,2}$	$Z_{1,3}$
$Z_{2,1}$	$Z_{2,2}$	$Z_{2,3}$
$Z_{3,1}$	$Z_{3,2}$	$Z_{3,3}$

# Output Size

$A_{1,1}$	$A_{1,2}$	$A_{1,3}$	$A_{1,4}$
$A_{2,1}$	$A_{2,2}$	$A_{2,3}$	$A_{2,4}$
$A_{3,1}$	$A_{3,2}$	$A_{3,3}$	$A_{3,4}$
$A_{4,1}$	$A_{4,2}$	$A_{4,3}$	$A_{4,4}$



$Z_{1,1}$	$Z_{1,2}$	$Z_{1,3}$
$Z_{2,1}$	$Z_{2,2}$	$Z_{2,3}$
$Z_{3,1}$	$Z_{3,2}$	$Z_{3,3}$

$$\text{Output Width} = \left[ \frac{(W_{\text{in}} - W_k + 2P)}{S} \right] + 1$$

$$\text{Output Height} = \left[ \frac{(H_{\text{in}} - H_k + 2P)}{S} \right] + 1$$

# Output Size

$A_{1,1}$	$A_{1,2}$	$A_{1,3}$	$A_{1,4}$
$A_{2,1}$	$A_{2,2}$	$A_{2,3}$	$A_{2,4}$
$A_{3,1}$	$A_{3,2}$	$A_{3,3}$	$A_{3,4}$
$A_{4,1}$	$A_{4,2}$	$A_{4,3}$	$A_{4,4}$



$Z_{1,1}$	$Z_{1,2}$	$Z_{1,3}$
$Z_{2,1}$	$Z_{2,2}$	$Z_{2,3}$
$Z_{3,1}$	$Z_{3,2}$	$Z_{3,3}$

$$\text{Output Width} = \left[ \frac{(W_{\text{in}} - W_k + 2P)}{S} \right] + 1$$

# Output Size

$A_{1,1}$	$A_{1,2}$	$A_{1,3}$	$A_{1,4}$
$A_{2,1}$	$A_{2,2}$	$A_{2,3}$	$A_{2,4}$
$A_{3,1}$	$A_{3,2}$	$A_{3,3}$	$A_{3,4}$
$A_{4,1}$	$A_{4,2}$	$A_{4,3}$	$A_{4,4}$



$Z_{1,1}$	$Z_{1,2}$	$Z_{1,3}$
$Z_{2,1}$	$Z_{2,2}$	$Z_{2,3}$
$Z_{3,1}$	$Z_{3,2}$	$Z_{3,3}$

$$\text{Output Width} = \left[ \frac{(W_{\text{in}} - W_k + 2\mathbf{P})}{\mathbf{S}} \right] + 1$$


**P**: Padding (here - 0)

**S**: Stride (here - 1)

# Padding

- Attaching zeros (usually) around inputs.
- Images can be padded to the left, right, top, and bottom.

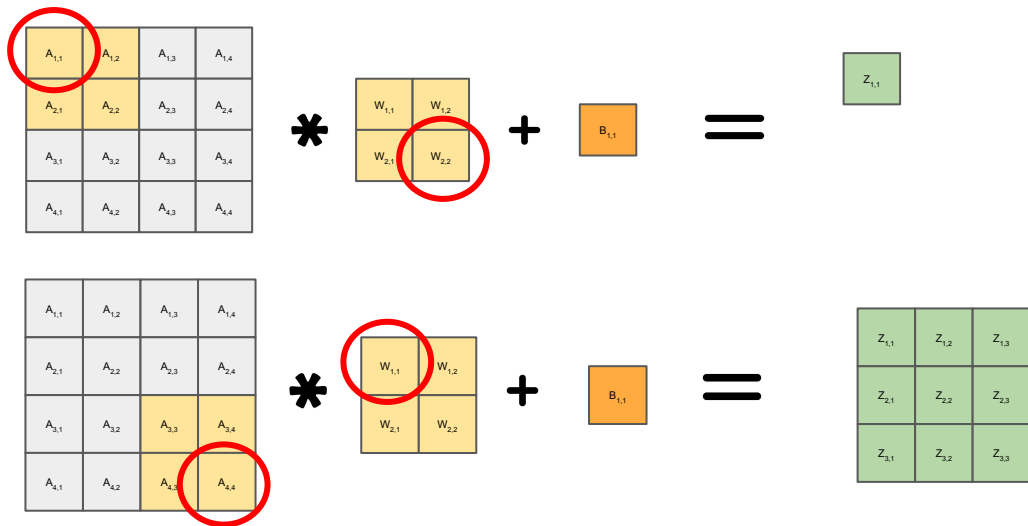
# Padding





$$z_{1,1} = (A_{1,1} * W_{1,1}) + (A_{1,2} * W_{1,2}) + (A_{2,1} * W_{2,1}) + (A_{2,2} * W_{2,2}) + B$$



# Padding



# Padding



Never Meet...

# Padding

Increase output size

Preserve input size

**More Kernel Interactions!**

# Padding

$A_{1,1}$	$A_{1,2}$	$A_{1,3}$	$A_{1,4}$
$A_{2,1}$	$A_{2,2}$	$A_{2,3}$	$A_{2,4}$
$A_{3,1}$	$A_{3,2}$	$A_{3,3}$	$A_{3,4}$
$A_{4,1}$	$A_{4,2}$	$A_{4,3}$	$A_{4,4}$

\*

$W_{1,1}$	$W_{1,2}$
$W_{2,1}$	$W_{2,2}$

+

$B_{1,1}$
-----------

=

$Z_{1,1}$	$Z_{1,2}$	$Z_{1,3}$
$Z_{2,1}$	$Z_{2,2}$	$Z_{2,3}$
$Z_{3,1}$	$Z_{3,2}$	$Z_{3,3}$

# Padding

0	0	0	0	0	0
0	$A_{1,1}$	$A_{1,2}$	$A_{1,3}$	$A_{1,4}$	0
0	$A_{2,1}$	$A_{2,2}$	$A_{2,3}$	$A_{2,4}$	0
0	$A_{3,1}$	$A_{3,2}$	$A_{3,3}$	$A_{3,4}$	0
0	$A_{4,1}$	$A_{4,2}$	$A_{4,3}$	$A_{4,4}$	0
0	0	0	0	0	0

\*

$W_{1,1}$	$W_{1,2}$
$W_{2,1}$	$W_{2,2}$

+

$B_{1,1}$
-----------

=

$Z_{1,1}$	$Z_{1,2}$	$Z_{1,3}$	$Z_{1,4}$
$Z_{2,1}$	$Z_{2,2}$	$Z_{2,3}$	$Z_{2,4}$
$Z_{3,1}$	$Z_{3,2}$	$Z_{3,3}$	$Z_{3,4}$
$Z_{4,1}$	$Z_{4,2}$	$Z_{4,3}$	$Z_{4,4}$

# Padding

0	0	0	0	0	0
0	$A_{1,1}$	$A_{1,2}$	$A_{1,3}$	$A_{1,4}$	0
0	$A_{2,1}$	$A_{2,2}$	$A_{2,3}$	$A_{2,4}$	0
0	$A_{3,1}$	$A_{3,2}$	$A_{3,3}$	$A_{3,4}$	0
0	$A_{4,1}$	$A_{4,2}$	$A_{4,3}$	$A_{4,4}$	0
0	0	0	0	0	0

\*

$W_{1,1}$	$W_{1,2}$
$W_{2,1}$	$W_{2,2}$

+

$B_{1,1}$
-----------

=

$Z_{1,1}$	$Z_{1,2}$	$Z_{1,3}$	$Z_{1,4}$
$Z_{2,1}$	$Z_{2,2}$	$Z_{2,3}$	$Z_{2,4}$
$Z_{3,1}$	$Z_{3,2}$	$Z_{3,3}$	$Z_{3,4}$
$Z_{4,1}$	$Z_{4,2}$	$Z_{4,3}$	$Z_{4,4}$

# Padding

0	0	0	0	0	0
0	$A_{1,1}$	$A_{1,2}$	$A_{1,3}$	$A_{1,4}$	0
0	$A_{2,1}$	$A_{2,2}$	$A_{2,3}$	$A_{2,4}$	0
0	$A_{3,1}$	$A_{3,2}$	$A_{3,3}$	$A_{3,4}$	0
0	$A_{4,1}$	$A_{4,2}$	$A_{4,3}$	$A_{4,4}$	0
0	0	0	0	0	0

\*

$W_{1,1}$	$W_{1,2}$
$W_{2,1}$	$W_{2,2}$

+

$B_{1,1}$
-----------

=

$Z_{1,1}$	$Z_{1,2}$	$Z_{1,3}$	$Z_{1,4}$
$Z_{2,1}$	$Z_{2,2}$	$Z_{2,3}$	$Z_{2,4}$
$Z_{3,1}$	$Z_{3,2}$	$Z_{3,3}$	$Z_{3,4}$
$Z_{4,1}$	$Z_{4,2}$	$Z_{4,3}$	$Z_{4,4}$



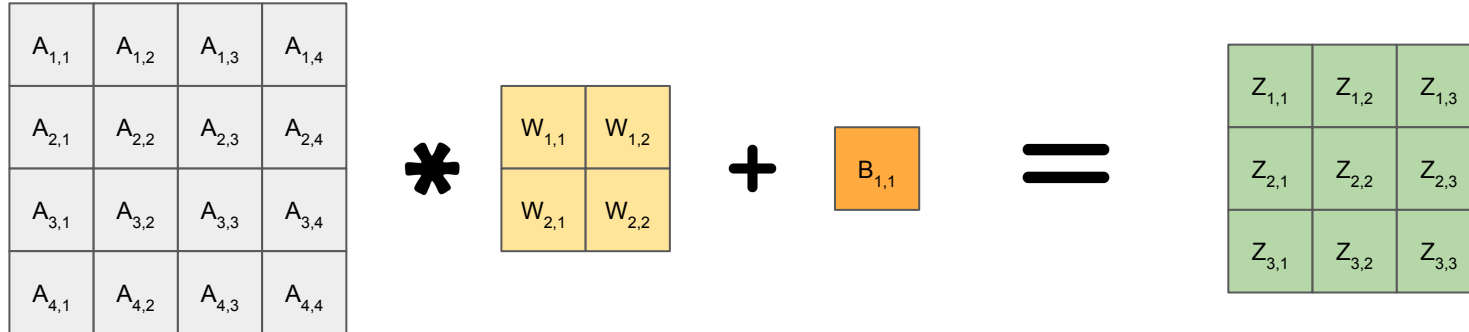
# Stride

Taking bigger steps!



# Stride = 1

What we did before - The kernel "moves" one pixel (or element) at a time.



# Stride = 2

Start at the same place

$A_{1,1}$	$A_{1,2}$	$A_{1,3}$	$A_{1,4}$
$A_{2,1}$	$A_{2,2}$	$A_{2,3}$	$A_{2,4}$
$A_{3,1}$	$A_{3,2}$	$A_{3,3}$	$A_{3,4}$
$A_{4,1}$	$A_{4,2}$	$A_{4,3}$	$A_{4,4}$

\*

$W_{1,1}$	$W_{1,2}$
$W_{2,1}$	$W_{2,2}$

+

$B_{1,1}$
-----------

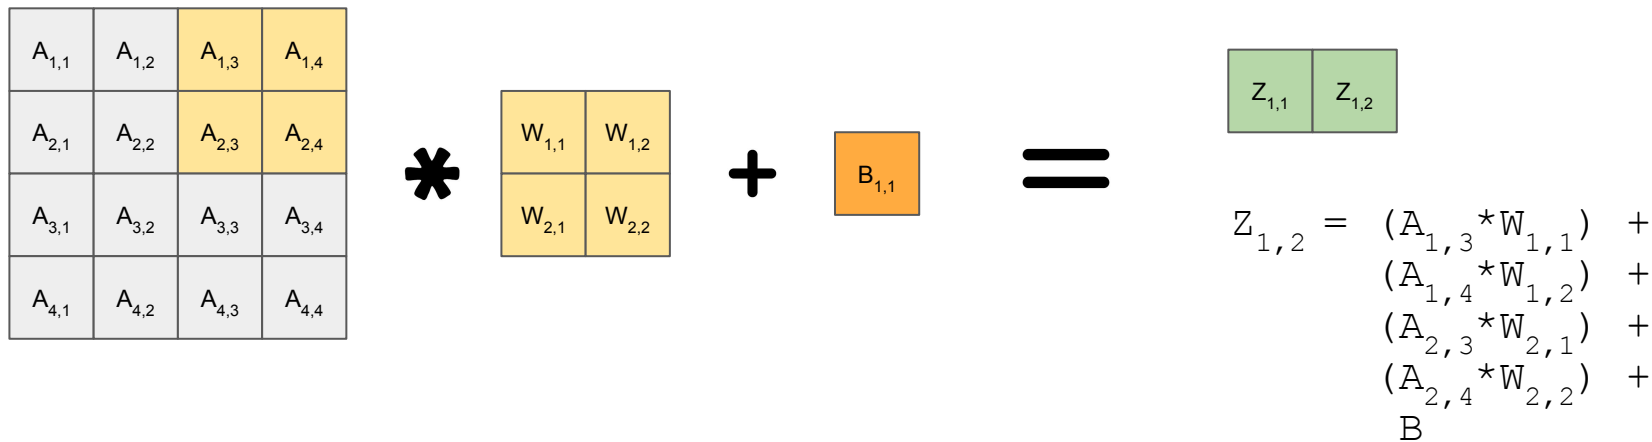
=

$Z_{1,1}$
-----------

$$Z_{1,1} = (A_{1,1} * W_{1,1}) + (A_{1,2} * W_{1,2}) + (A_{2,1} * W_{2,1}) + (A_{2,2} * W_{2,2}) + B$$


# Stride = 2

Move two elements to the right



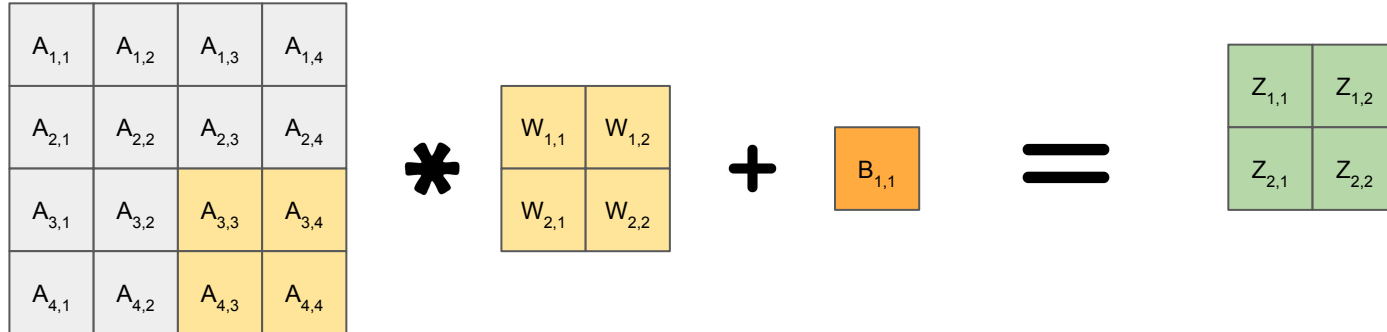
# Stride = 2

Move two elements down.




# Stride = 2

Move two elements to the right.



# Multi-channel CNN



# Multi-channel CNN


- Each kernel (or **filter**) has as many channels as the input does
- Channel **c of the kernel** convolves with channel **c** (corresponding) **of the input**.
- The number of output channels from the convolution = number of **filters**

$C_{in}$  = Input channels

$C_{kernel}$  = Kernel channels =  $C_{in}$


$K$  = Number of Kernels =  $C_{out}$  = Number of output channels

# 1 Filter with 3-channel input




1 channel input

$\otimes$




=



1 almost-output  
map




# 1 Filter with 3-channel input




2 channel input

$\otimes$




Kernel

=




2 almost-output maps

# 1 Filter with 3-channel input





3 channel input

$\otimes$




Kernel

=




3 almost-output maps

# 1 Filter with 3-channel input




3 channel input

$\otimes$



Kernel


=



3 almost-output maps


→

Add  
through  
channels

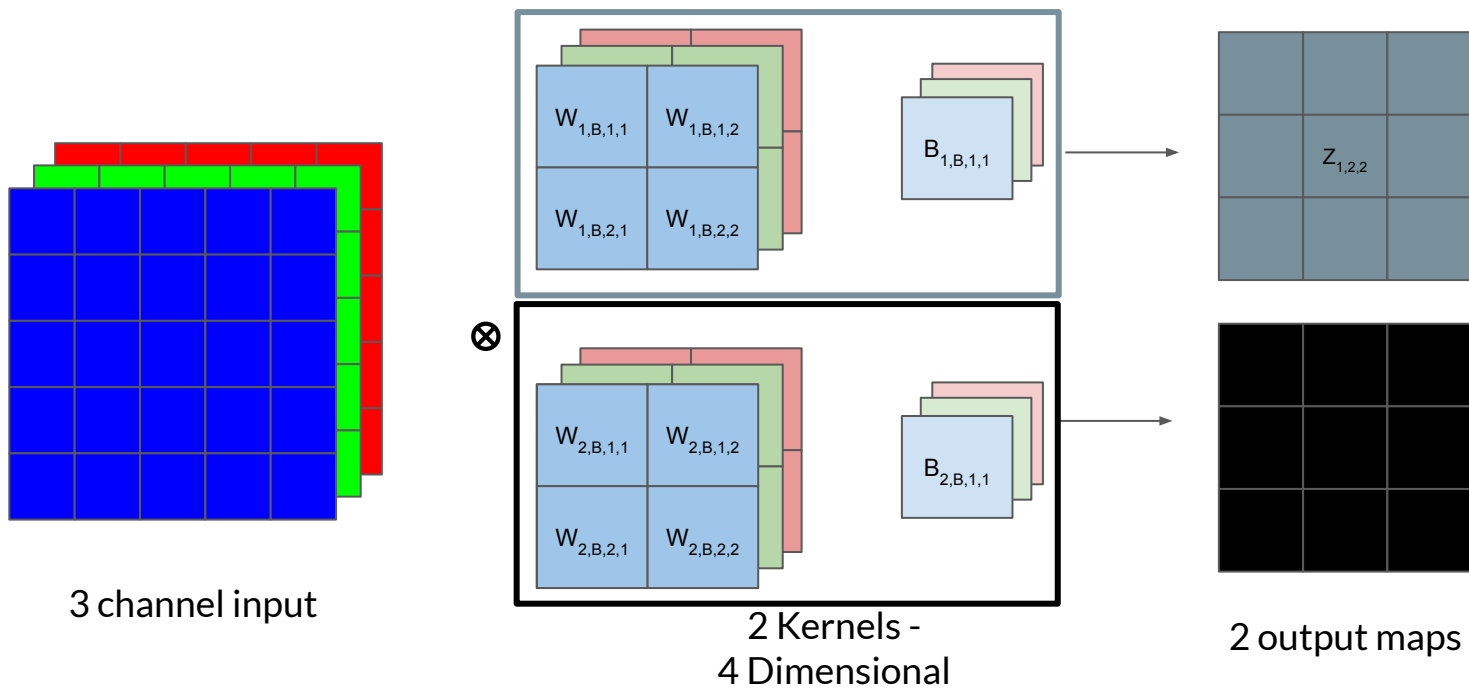


1 output map


# 2 Filters with 3-channel input



# 2 Filters with 3-channel input




# 2 Filters with 3-channel input



# Pooling


- Usually follows convolutions
- Introduces Jitter Invariance
- Reduces feature-map size
- **Max, Mean, Min**

# Pooling





# Pooling



# Convolutional Layer Implementation

- Earlier we said that a discrete convolution is a concatenation of inner products between the filter and receptive fields. This involves a lot of matrix multiplications!
- Parallelized matrix operations make this much faster than using a sliding filter.
  - `im2col`, `im2col_bw` are needed

**THIS IS THE HARDEST PART OF THE PROGRAMMING. START EARLY!**

# Forward propagation

Here are the following steps

1. Transform our input image into a matrix (im2col)
2. Reshape our kernel (flatten)
3. Perform matrix multiplication between reshaped input image and kernel

# Forward propagation

Here are the following steps

1. Transform our input image into a matrix (im2col)
2. Reshape our kernel (flatten)
3. Perform matrix multiplication between reshaped input image and kernel

# Step 1: Transform our input image into a matrix (im2col)


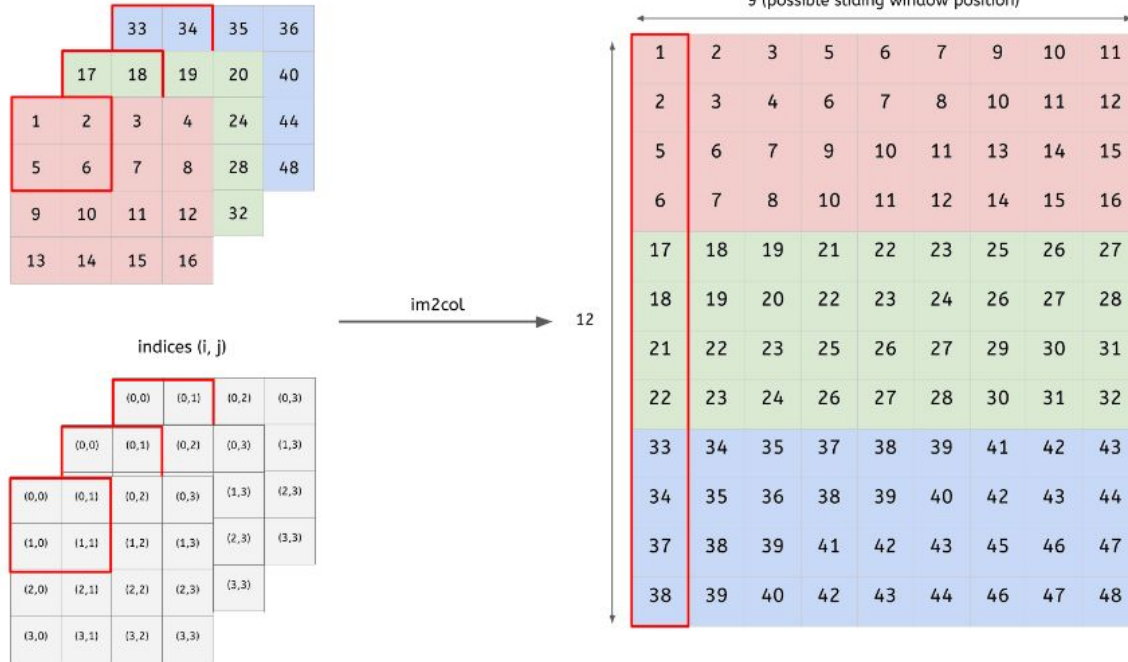


Figure 2: Illustration of im2col.

# Example

We will perform a convolution between an (1,3,4,4) input image and kernels of shape (2,3,2,2)



# Forward propagation

Here are the following steps

1. Transform our input image into a matrix (im2col)
2. Reshape our kernel (flatten)
3. Perform matrix multiplication between reshaped input image and kernel

# Forward propagation

Here are the following steps

1. Transform our input image into a matrix (im2col)
2. Reshape our kernel (flatten)
3. Perform matrix multiplication between reshaped input image and kernel



# Step 2: Reshape our kernel (flatten)




Figure 2: Reshaped version of the 2 kernels

As you can see, each filter is flattened and then stacked together. Thus, for  $X$  filter, we will flatten and stack  $X$  filters together.

# Forward propagation

Here are the following steps

1. Transform our input image into a matrix (im2col)
2. Reshape our kernel (flatten)
3. Perform matrix multiplication between reshaped input image and kernel

# Forward propagation

Here are the following steps

1. Transform our input image into a matrix (im2col)
2. Reshape our kernel (flatten)
3. Perform matrix multiplication between reshaped input image and kernel




## Step 3: Matrix multiplication between reshaped input and kernel

- At the end, we need to reshape our matrix back to a feature map
- Be aware the `np.reshape()` method doesn't return the expected result here
  - Elements in the wrong order
  - A little bit of numpy *\*magic\** solves the problem

# Pooling Layer

- We can make the average pooling operation faster by using **im2col** method.
- Be aware that the `np.reshape()` method doesn't return the expected result here (elements in wrong order). A little bit of numpy gymnastic solves the problem.



# Backward propagation

## Reminder:

- We performed a convolution between (1,3,4,4) input image and kernels of shape (2,3,2,2) which output an (2,3,3) image.
- During the backward pass, the (2,3,3) image contains the error/gradient (“**dout**”) which needs to be back-propagated to the:
  - (1,3,4,4) input image (layer).
  - (2,3,2,2) kernels.

## ★ Layer gradient: Intuition

---

- The formula to compute the layer gradient is:


$$\frac{\partial L}{\partial I} = Conv(K, \frac{\partial L}{\partial O})$$

- $\frac{\partial L}{\partial I}$ : Input gradient.
  - $K$ : Kernels.
  - $\frac{\partial L}{\partial O}$ : Output gradient.
  - $Conv$ : Convolution operation.
- To do so, we will proceed as follow:
    - **A.** Reshape dout ( $\frac{\partial L}{\partial O}$ ).
    - **B.** Reshape kernels  $w$  into single matrix  $w_{col}$ .
    - **C.** Perform matrix multiplication between reshaped  $dout$  and kernel.
    - **D.** Reshape back to image (**col2im**).
  - We are going to see how it works intuitively and then how to implement it using Numpy.

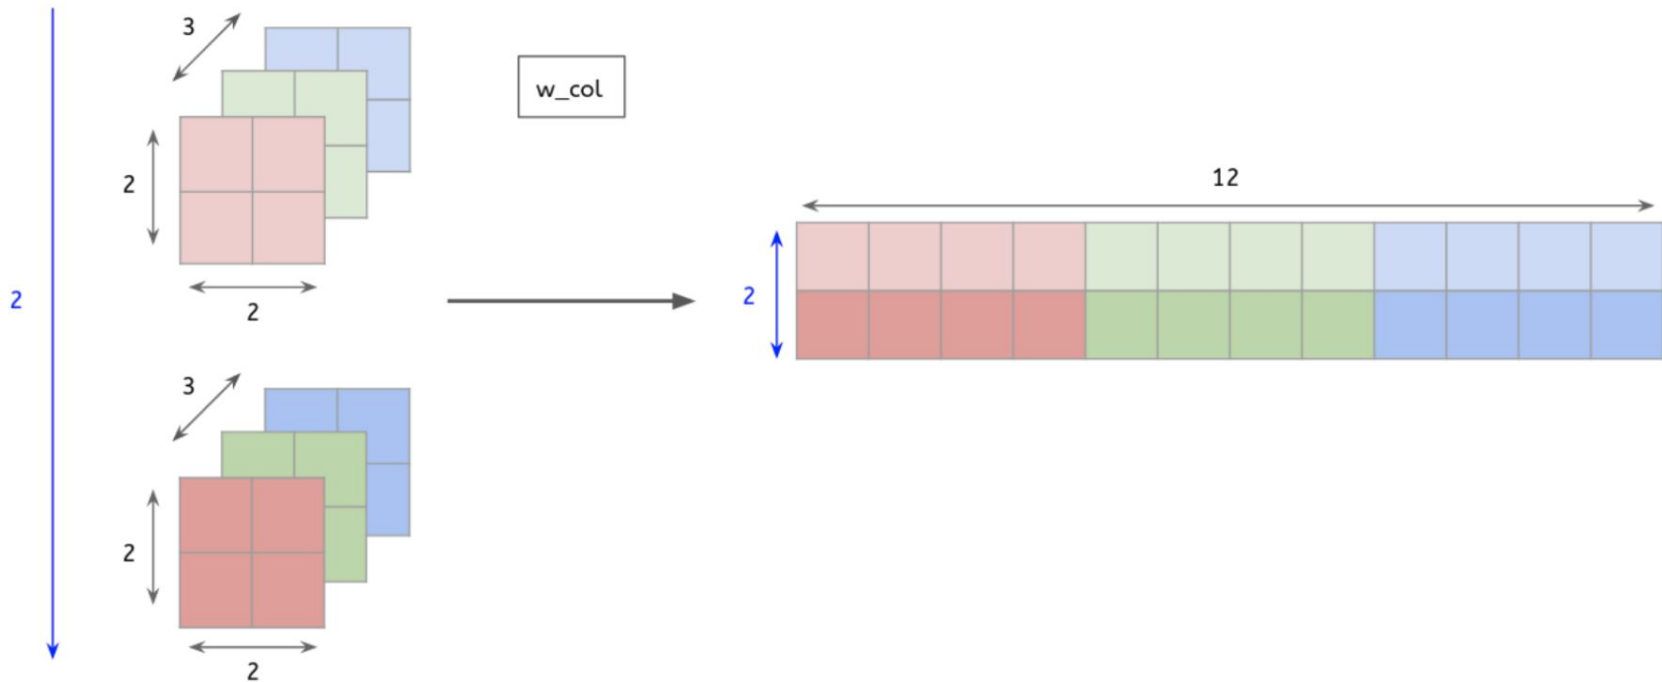


## A) Reshape dout

- During backward propagation, the output of the forward convolution contains the error that needs to be back-propagated.




**B) Reshape w into w\_col**




### C) Perform matrix multiplication between reshaped dout and w\_col

- In order to perform to perform the matrix multiplication, we need to transpose `w_col` .
- We will denoted the output as `dX_col` .



## D) Reshape back to image (col2im)

- Here, **col2im** is more than a simple backward operation of im2col. Indeed, we have to take care of cases where errors will overlap with others.



## ○ Kernel gradient: Intuition

---


- The formula to compute the kernel gradient is:

$$\frac{\partial L}{\partial K} = \text{Conv}(I, \frac{\partial L}{\partial O})$$

- $\frac{\partial L}{\partial K}$ : Kernels gradient.
  - $I$ : Input image.
  - $\frac{\partial L}{\partial O}$ : Output gradient.
  - $\text{Conv}$ : Convolution operation.
- To do so, we will:
    - **A.** Reshape `dout` ( $\frac{\partial L}{\partial O}$ ).
    - **B.** Apply `im2col` on `x` to get `x_col`.
    - **C.** Perform matrix multiplication between reshaped `dout` and `x_col` to get `dw_col`.
    - **D.** Reshape `dw_col` back to `dw`.
  - We are going to see how it works intuitively and then how to implement it using Numpy.

## A) Reshape `dout`

- Be aware that the `np.reshape()` method doesn't return the expected result here (elements in wrong order). A little bit of numpy gymnastic solves the problem.



## B) Apply im2col on X to get X col

		33	34	35	36
	17	18	19	20	40
1	2	3	4	24	44
5	6	7	8	28	48
9	10	11	12	32	
13	14	15	16		

indices (i, j)

		(0,0)	(0,1)	(0,2)	(0,3)		
		(0,0)	(0,1)	(0,2)	(0,3)	(1,3)	
		(0,0)	(0,1)	(0,2)	(0,3)	(1,3)	(2,3)
	(1,0)	(1,1)	(1,2)	(1,3)	(2,3)	(3,3)	
(2,0)	(2,1)	(2,2)	(2,3)	(3,3)			
(3,0)	(3,1)	(3,2)	(3,3)				

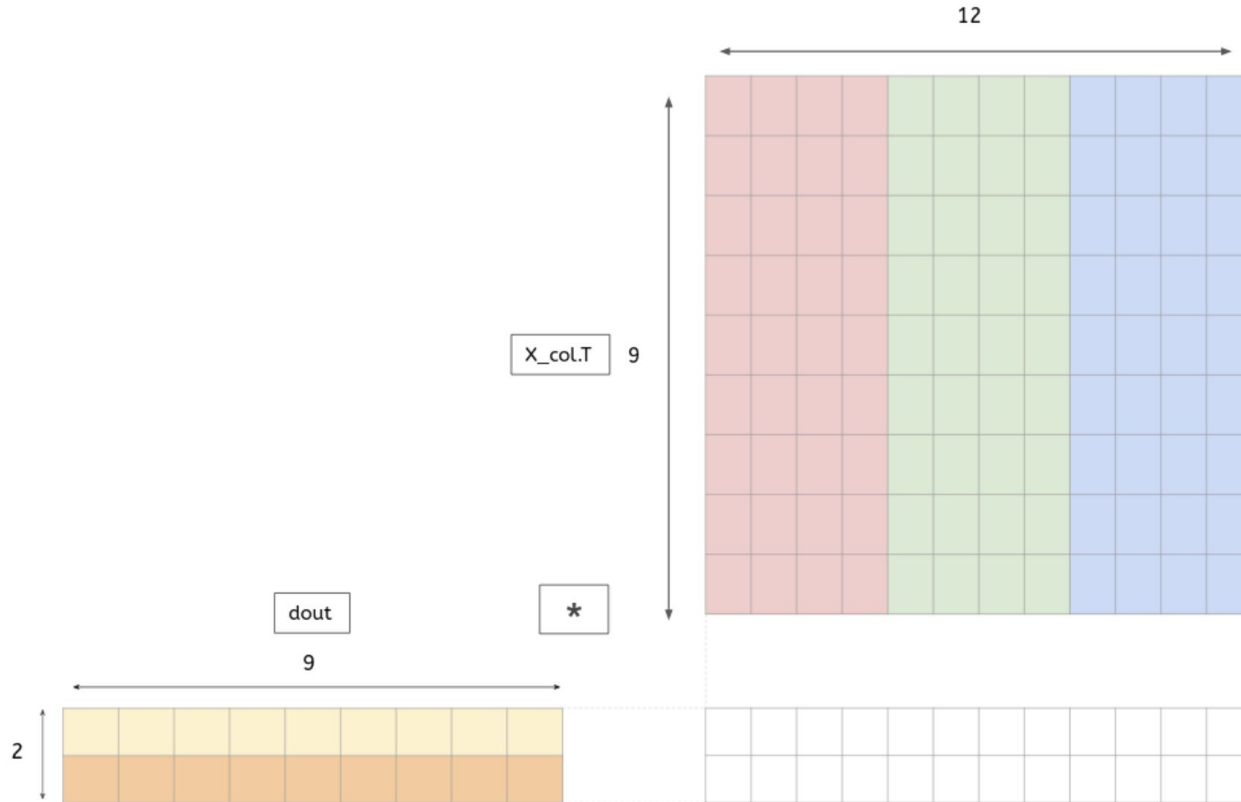
im2col → 12

9 (possible sliding window position)

1	2	3	5	6	7	9	10	11
2	3	4	6	7	8	10	11	12
5	6	7	9	10	11	13	14	15
6	7	8	10	11	12	14	15	16
17	18	19	21	22	23	25	26	27
18	19	20	22	23	24	26	27	28
21	22	23	25	26	27	29	30	31
22	23	24	26	27	28	30	31	32
33	34	35	37	38	39	41	42	43
34	35	36	38	39	40	42	43	44
37	38	39	41	42	43	45	46	47
38	39	40	42	43	44	46	47	48

### C) Perform matrix multiplication between reshaped dout and X\_col to get dw\_col


- In order to perform the matrix multiplication, we need to transpose `X_col`.





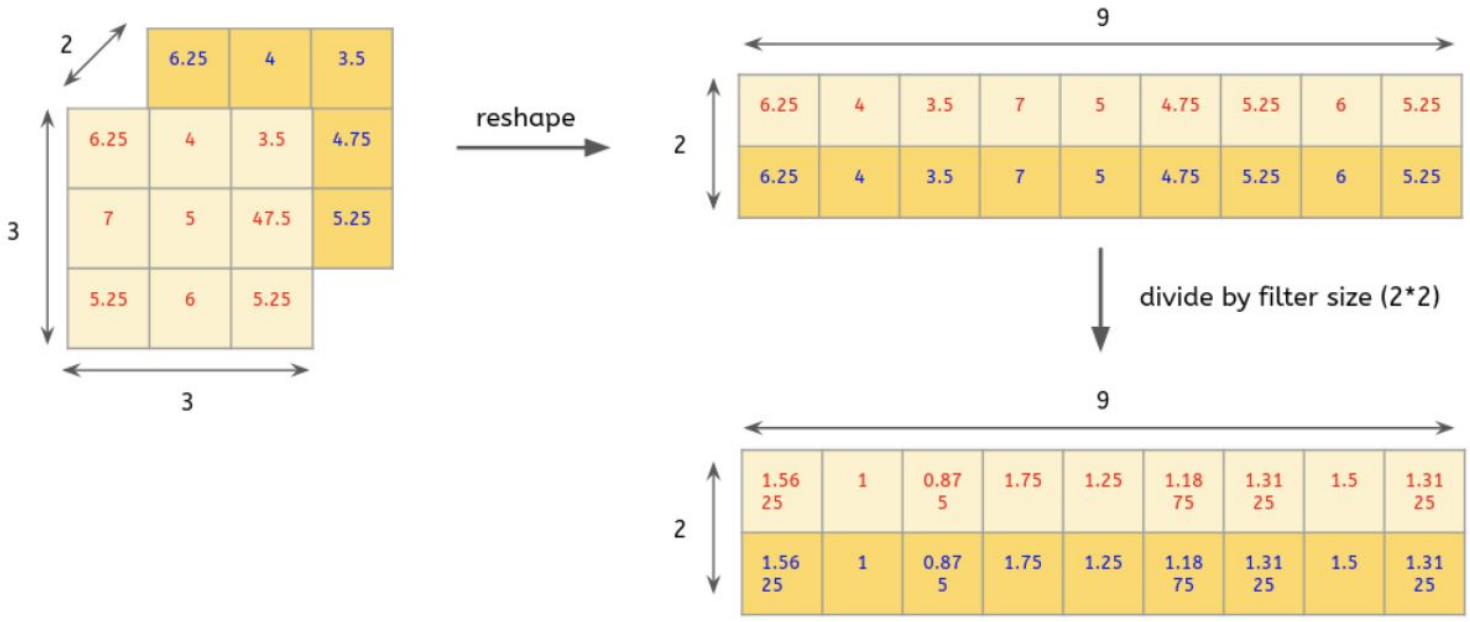
## D) Reshape dw\_col back to dw

- We simply need to reshape `dw_col` back to its original kernel shape.



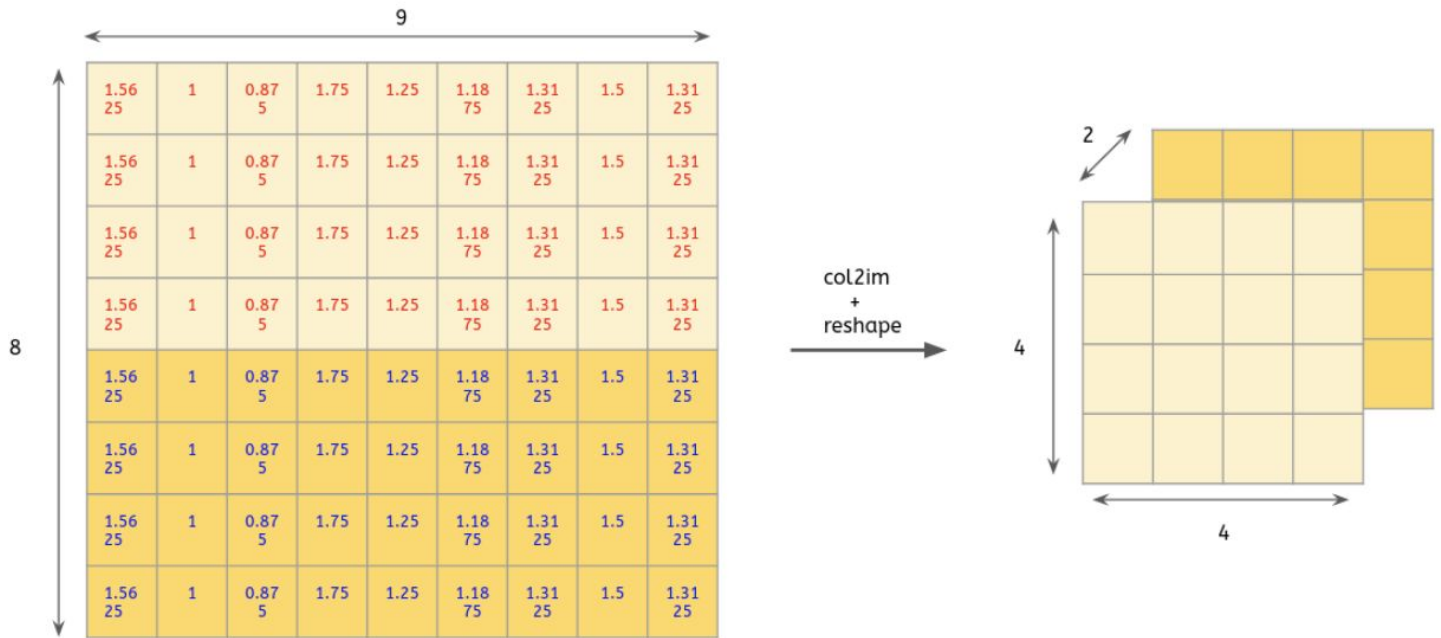
# 2) Pooling layer

- We first have to reshape our filters and divide by the filter size.





- Finally, we apply `col2im`.
- Be aware that the `np.reshape()` method doesn't return the expected result here (elements in wrong order). A little bit of numpy gymnastic solves the problem.



# Some suggestions

Please **start early**, this assignment takes time!!!

- Im2col usually takes the most effort, try derive the example we did above using your hand to better understand it
- Make sure to test it correctly with small hand-derived example.