

Cornell Bowers C-IS

College of Computing and Information Science

Convolutional Neural Networks

CS4782: Intro to Deep Learning

Thanks to:

Varsha Kishore
Justin Lovelace
Anissa Dallmann
Stephanie Ginting
Alexander Scotte

Logistics

- **HW1** has been released
 - Due next Thursday (February 13)
 - Homework clarifications are listed as pinned posts under HW1 on Ed
- CS 5782 - **Quiz 1** will be released today
 - 20 min duration - make sure to start well before it's due
 - Submission Due: Thursday 11:59 PM
- **Coding Assignment 1** to be released this week.
- Office hours are listed on the course website
- Post questions on Ed

So far...

- MLPs learn complex decision boundaries
- Optimization algorithms use the gradient of the loss to find network parameters
- Different training strategies like regularization, early stopping and normalization can improve training and generalization




Image Classification



input image



“dog”




input image



“cat”

Applications in Medicine




Why not use a Multi-Layer Perceptron?



flatten →

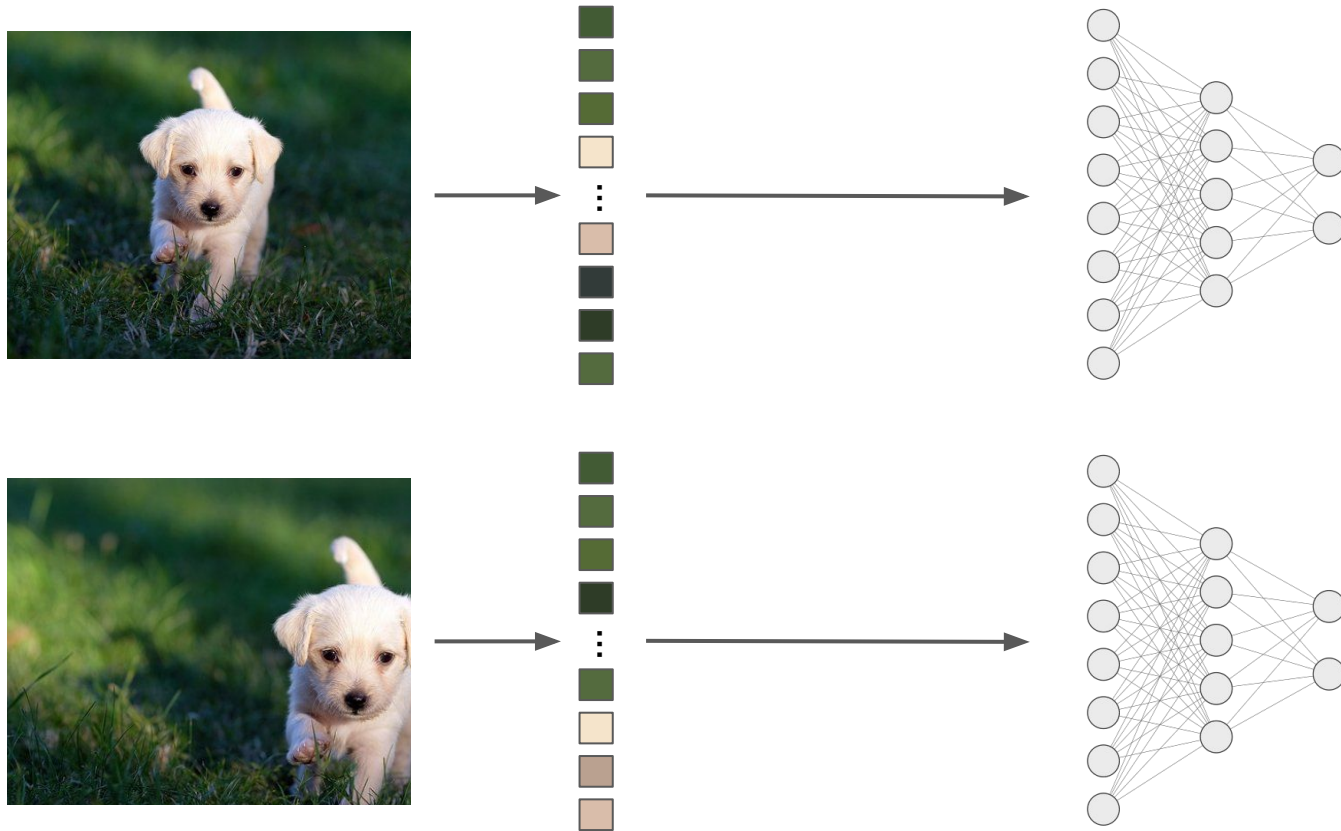


→

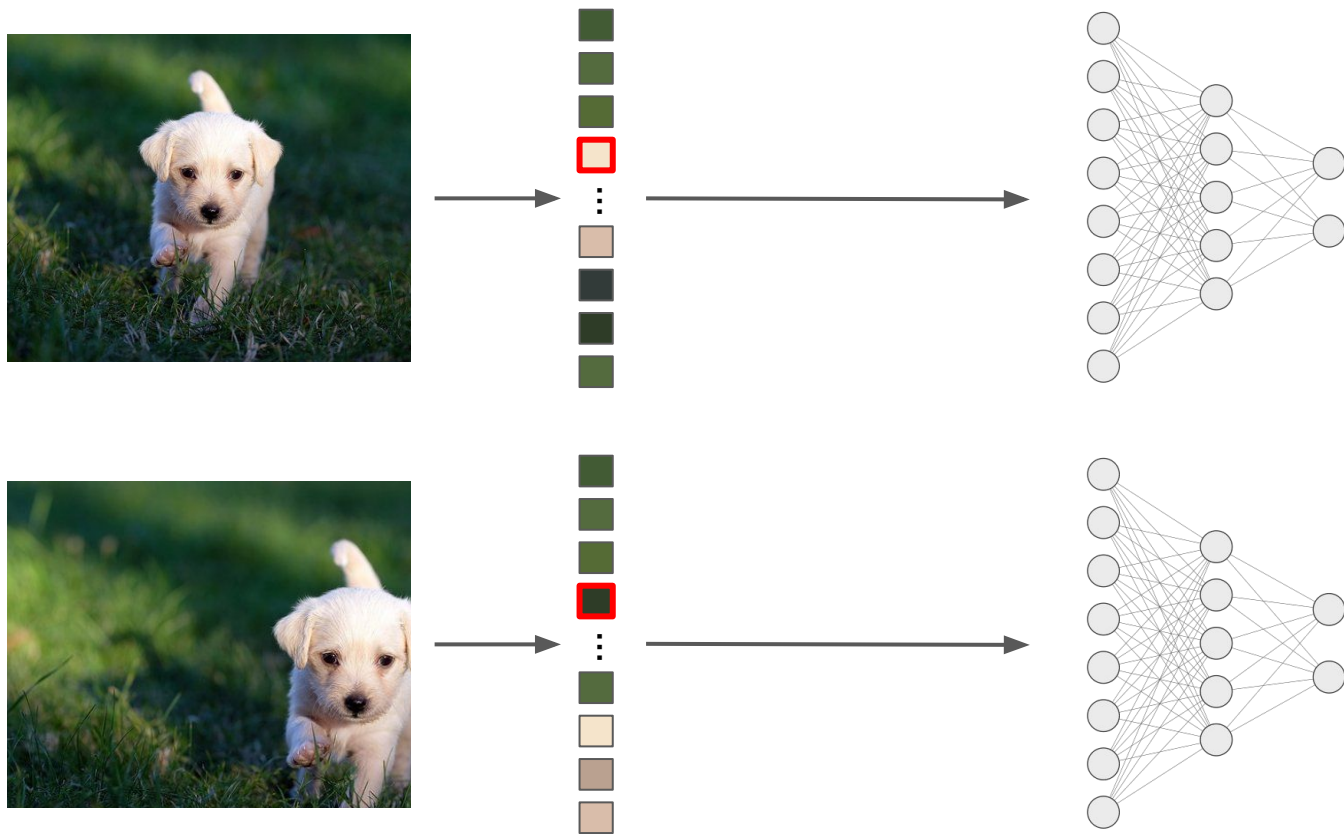


Which pixels were next to each other?

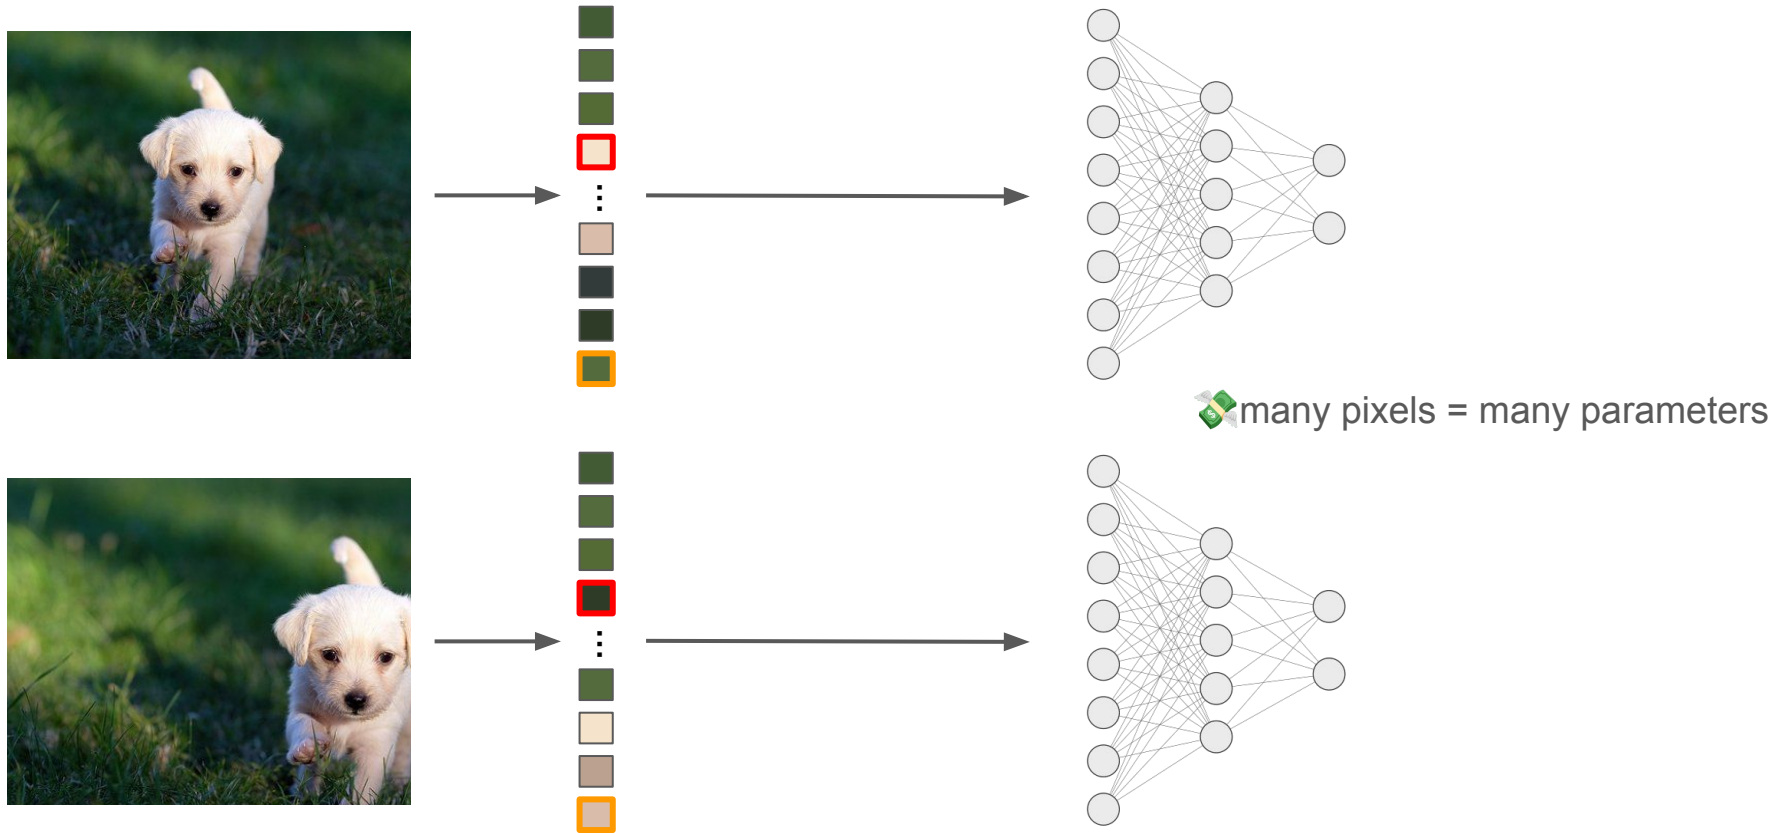
Why not use a Multi-Layer Perceptron?



Why not use a Multi-Layer Perceptron?



Why not use a Multi-Layer Perceptron?



Convolutional Filters

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

“image”

*

0	1	0
1	0	1
0	1	0

convolutional filter

Convolutional Filters

1 x0	1 x1	1 x0	0	0
0 x1	0 x0	1 x1	1	0
0 x0	1 x1	0 x0	1	1
1	1	0	0	0
1	0	0	1	1

“image”

*

0	1	0
1	0	1
0	1	0

convolutional filter

=

3		

Convolutional Filters

1	1 _{x0}	1 _{x1}	0 _{x0}	0
0	0 _{x1}	1 _{x0}	1 _{x1}	0
0	1 _{x0}	0 _{x1}	1 _{x0}	1
1	1	0	0	0
1	0	0	1	1

“image”

*

0	1	0
1	0	1
0	1	0

convolutional filter

=

3	2	

Convolutional Filters

1	1	1 _{x0}	0 _{x1}	0 _{x0}
0	0	1 _{x1}	1 _{x0}	0 _{x1}
0	1	0 _{x0}	1 _{x1}	1 _{x0}
1	1	0	0	0
1	0	0	1	1

“image”

*

0	1	0
1	0	1
0	1	0

convolutional filter

=

3	2	2

Convolutional Filters

1	1	1	0	0
0 _{x0}	0 _{x1}	1 _{x0}	1	0
0 _{x1}	1 _{x0}	0 _{x1}	1	1
1 _{x0}	1 _{x1}	0 _{x0}	0	0
1	0	0	1	1

“image”

*

0	1	0
1	0	1
0	1	0

convolutional filter

=

3	2	2
1		

Convolutional Filters

1	1	1	0	0
0	0 _{x0}	1 _{x1}	1 _{x0}	0
0	1 _{x1}	0 _{x0}	1 _{x1}	1
1	1 _{x0}	0 _{x1}	0 _{x0}	0
1	0	0	1	1

“image”

*

0	1	0
1	0	1
0	1	0

convolutional filter

=

3	2	2
1	3	

Convolutional Filters

1	1	1	0	0
0	0	1 _{x0}	1 _{x1}	0 _{x0}
0	1	0 _{x1}	1 _{x0}	1 _{x1}
1	1	0 _{x0}	0 _{x1}	0 _{x0}
1	0	0	1	1

“image”

*

0	1	0
1	0	1
0	1	0

convolutional filter

=

3	2	2
1	3	2

Convolutional Filters

1	1	1	0	0
0	0	1	1	0
0 _{x0}	1 _{x1}	0 _{x0}	1	1
1 _{x1}	1 _{x0}	0 _{x1}	0	0
1 _{x0}	0 _{x1}	0 _{x0}	1	1

“image”

*

0	1	0
1	0	1
0	1	0

convolutional filter

=

3	2	2
1	3	2
2		

Convolutional Filters

1	1	1	0	0
0	0	1	1	0
0	1 _{x0}	0 _{x1}	1 _{x0}	1
1	1 _{x1}	0 _{x0}	0 _{x1}	0
1	0 _{x0}	0 _{x1}	1 _{x0}	1

“image”

*

0	1	0
1	0	1
0	1	0

convolutional filter

=

3	2	2
1	3	2
2	1	

Convolutional Filters

1	1	1	0	0
0	0	1	1	0
0	1	0 _{x0}	1 _{x1}	1 _{x0}
1	1	0 _{x1}	0 _{x0}	0 _{x1}
1	0	0 _{x0}	1 _{x1}	1 _{x0}

"image"

*

0	1	0
1	0	1
0	1	0

convolutional filter

=

3	2	2
1	3	2
2	1	2

Convolutional Filters

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

"image"

*

0	1	0
1	0	1
0	1	0

convolutional filter

=


3	2	2
1	3	2
2	1	2



can learn this!

Convolutional Filters

- ❖ Aggregates information from local window around pixel
- ❖ Translational invariance
- ❖ Reduce number of parameters needed to be learned



Discuss with your Neighbor!

Match the following convolutional filters with the output they produce.





input image



-1	-1	-1
0	0	0
1	1	1

-1	0	1
-1	0	1
-1	0	1




$1/9$	$1/9$	$1/9$
$1/9$	$1/9$	$1/9$
$1/9$	$1/9$	$1/9$



Dilated Convolutions

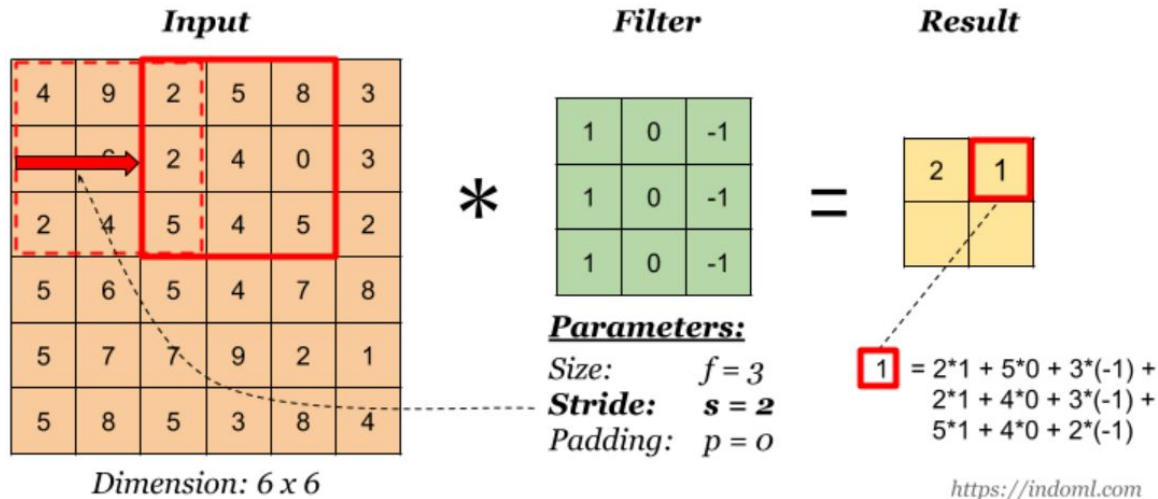


1D and 3D Convolutions



CNNs - Stride

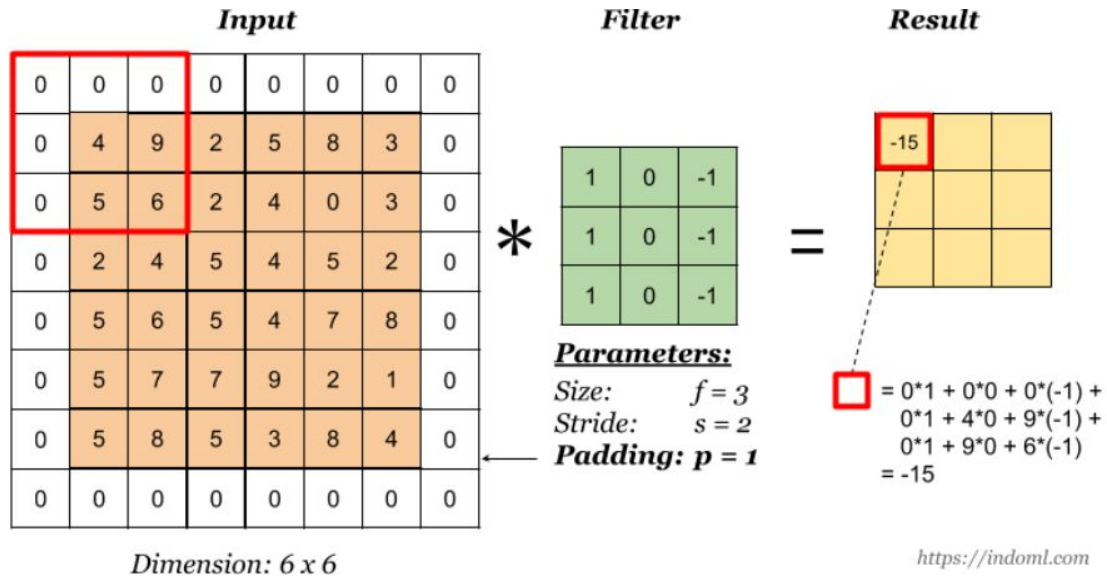
- ❖ Stride controls how many units the filter / the receptive field shift at a time
- ❖ The size of the output image shrinks more as the stride becomes larger
- ❖ The receptive fields overlap less as the stride becomes larger




Filter with stride (s) = 2

CNNs - Padding

- ❖ Padding adds layers of zeros (or other number) around image border
- ❖ Prevents image shrinking and loss of information from image boundary



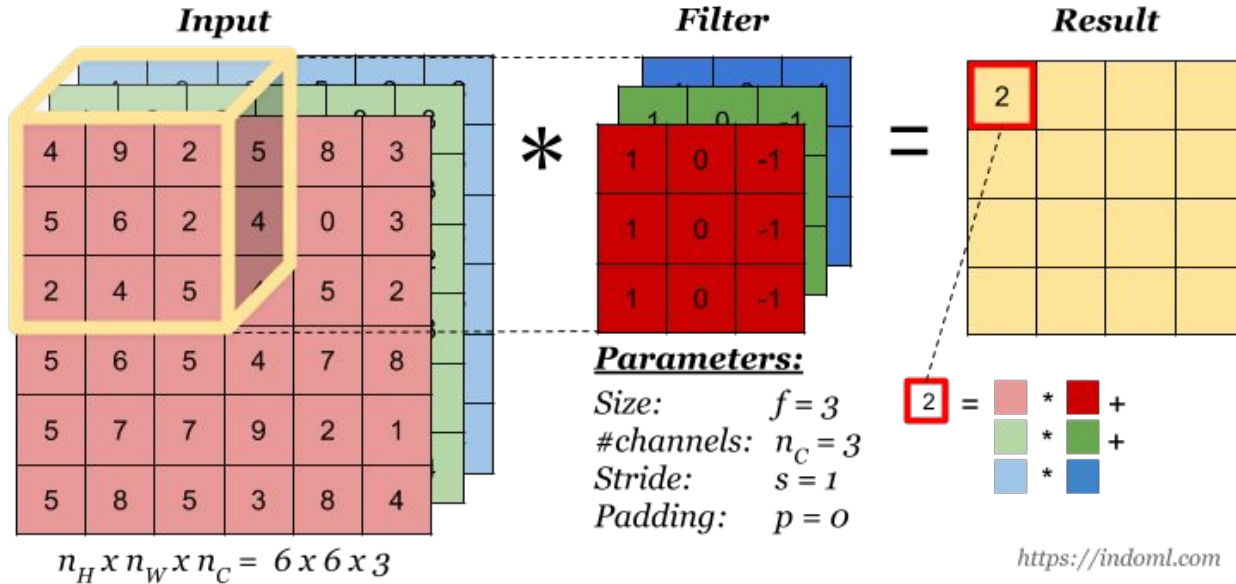
Stacking Convolutions



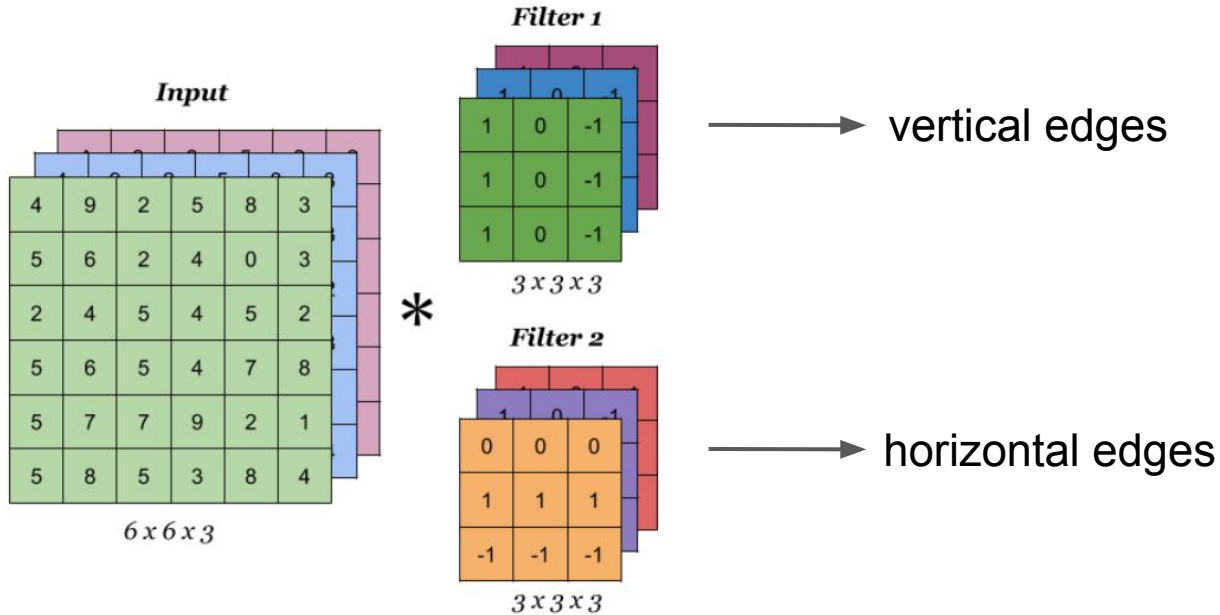
- ❖ Size of receptive field increases with each layer
- ❖ Capture more complex features

Convolution Over Volumes


What if our input image has more than one channel?




Convolution Operation with Multiple Filters



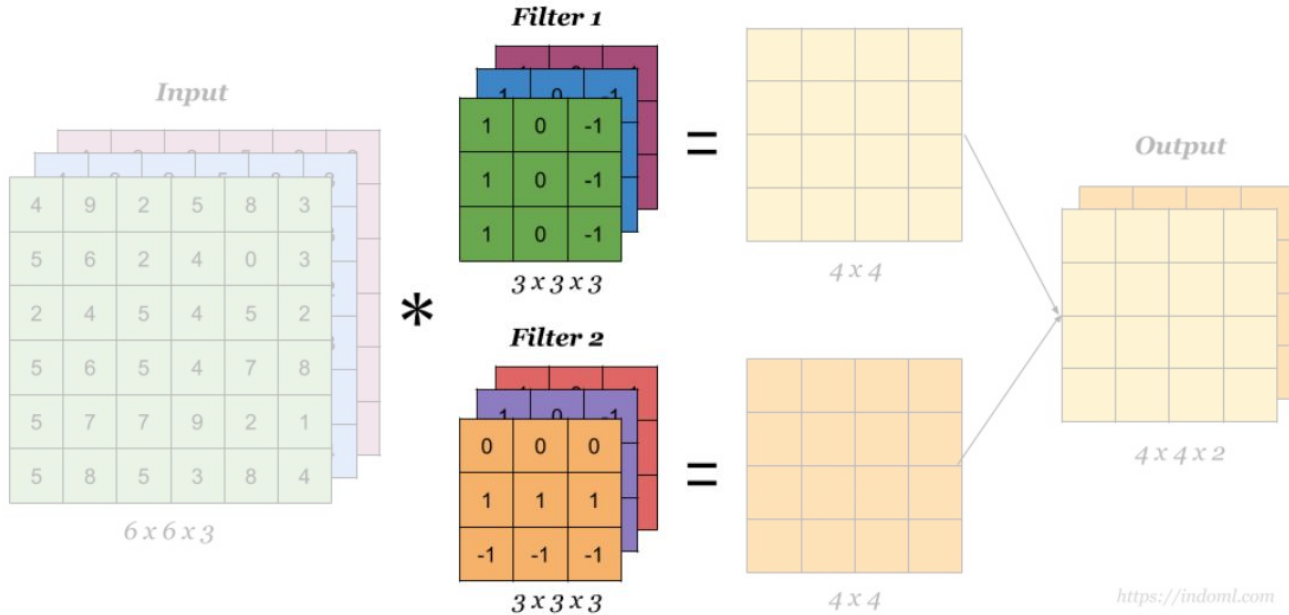
Convolution Operation with Multiple Filters



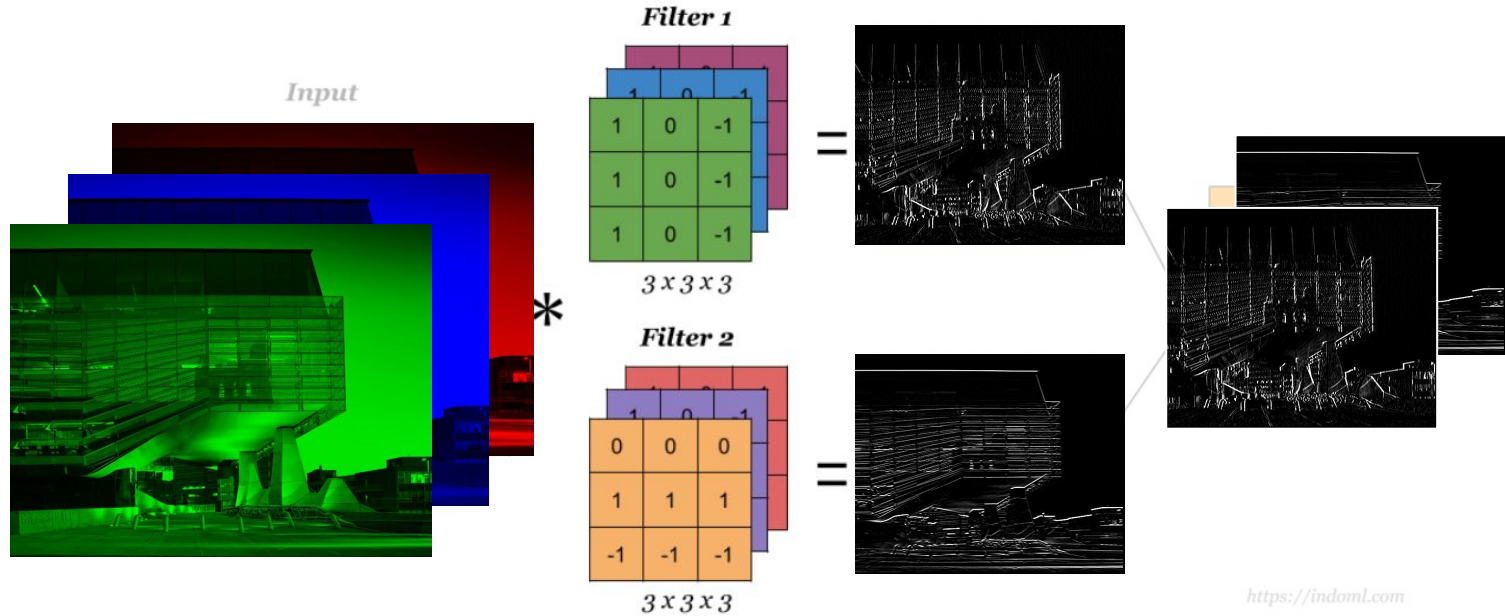
Convolution Operation with Multiple Filters



Convolution Operation with Multiple Filters




Convolution Operation with Multiple Filters




Discuss: 1x1 Convolutions

What is the result of convolving a 64x64x192 dimensional cube with a 1x1 filter?




Convolution Layer

MLP Layer



Input


Convolution Layer



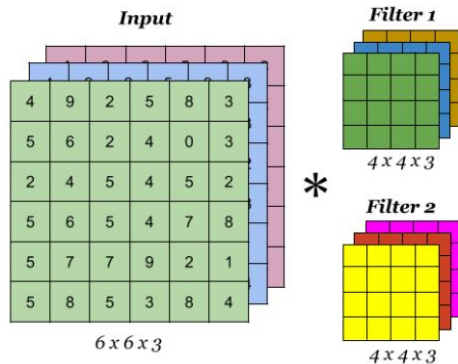
6x6x3

Convolution Layer

MLP Layer




Convolution Layer




Convolution Layer

MLP Layer




Convolution Layer

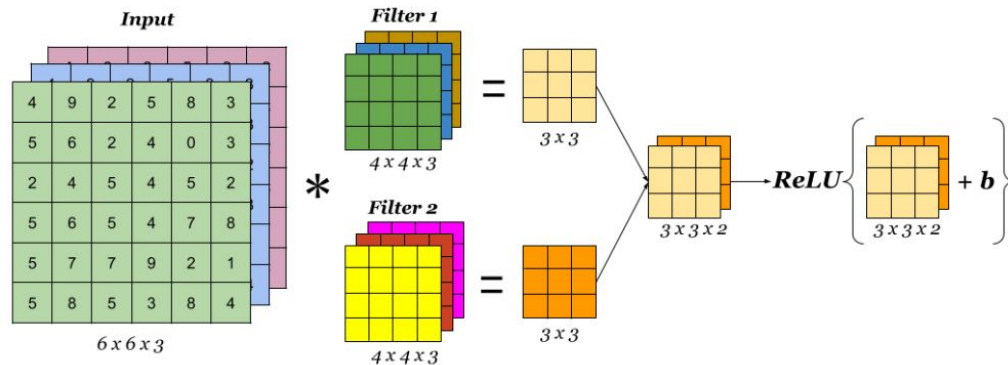


Convolution Layer

MLP Layer




Convolution Layer




Convolution Layer

MLP Layer



Convolution Layer




CNN/MLP Equivalence

Differences in a convolution layer:


- neurons are connected to a local region
- Weights are shared across multiple parameters

CONV layers can be converted to Fully connected layers and vice versa!

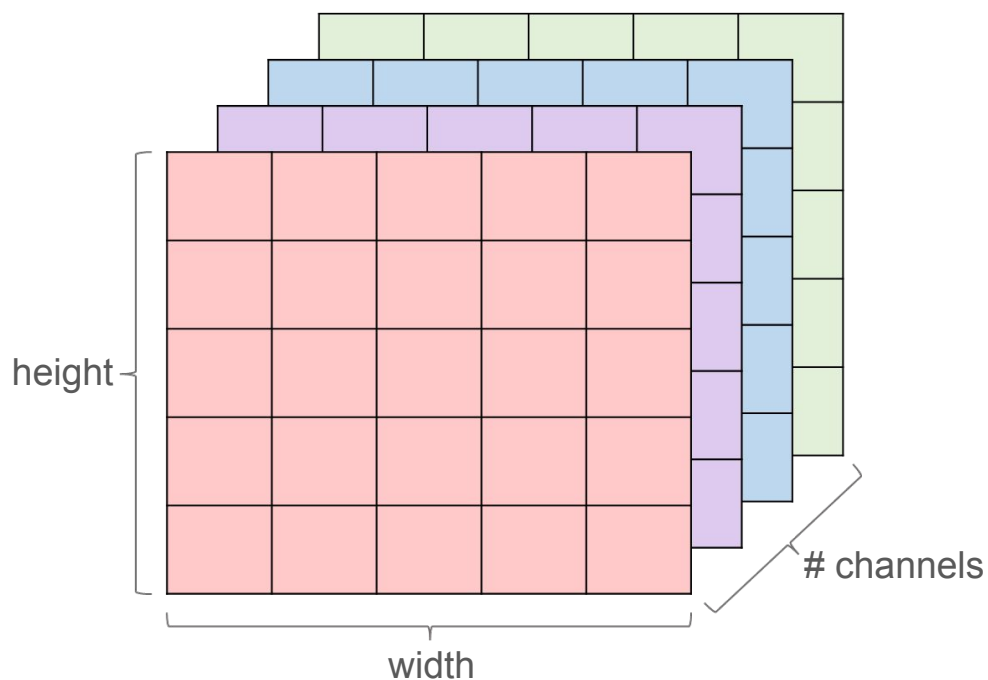


Discuss: Trade-offs between CNNs and MLPs


How would this image change if you used an MLP instead of a 1×1 convolution filter to produce a $(64 \times 64 \times 1)$ feature map? Hint: think about parameter counts and feature interactions.



CNN Layer Output Visualization



=



Convolutional Neural Networks (CNNs)


✓ Convolutions

Maintain spatial relation between pixels

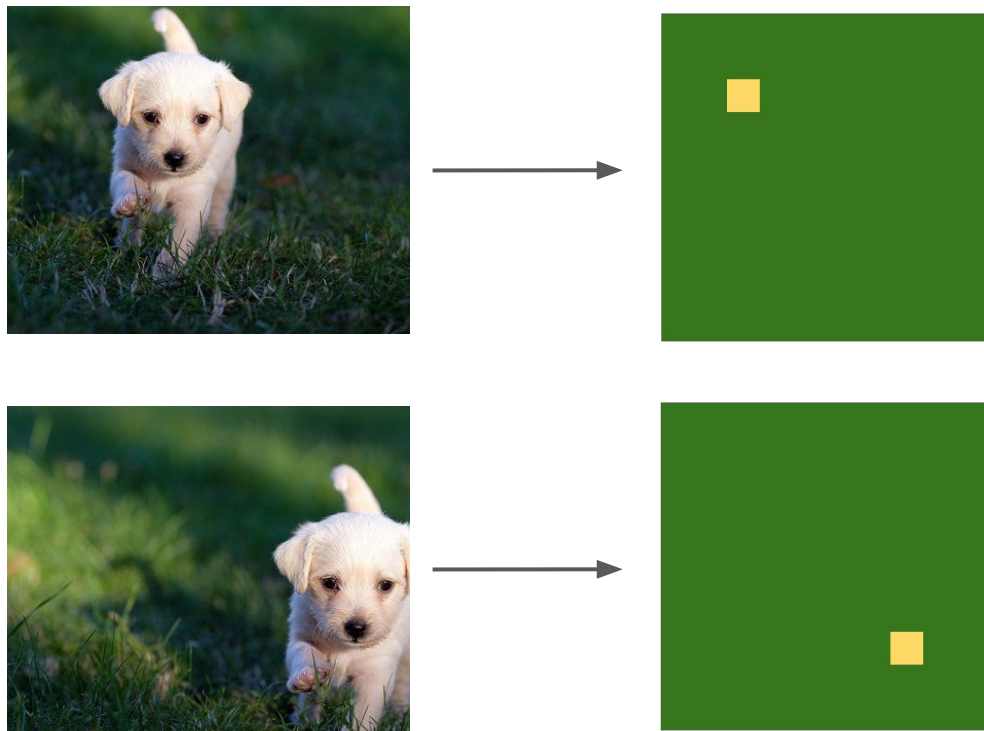
Reduce number of parameters through weight sharing



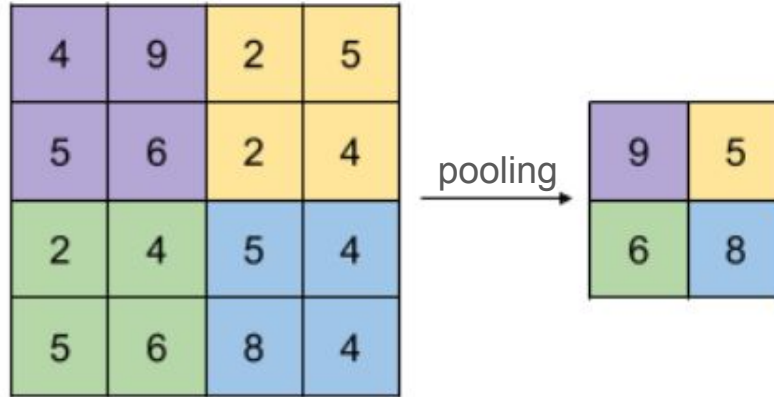
input image




Ensuring translational invariance



Max Pooling




CNNs - Pooling




CNNs - Pooling

- ❖ Down sample feature maps that highlight the most present feature in the patch
- ❖ Improve efficiency by reducing computations with downsampling
- ❖ Increase receptive field size

Max Pooling



Avg Pooling



Convolutional Neural Networks (CNNs)

✓ Convolutions

Maintain spatial relation between pixels

Reduce number of parameters through weight sharing


✓ Pooling

Captures key information from across different areas of the feature maps

Together with convolutions allows for translational invariance




input image



Normalization

- ❖ Normalize channels to mean 0 and variance 1 across each training batch
- ❖ Increases speed of training by enabling the use of larger learning rates
- ❖ Improves stability of training



The Batch Normalization Algorithm

Input: Values of x over a mini-batch: $\mathcal{B} = \{x_{1\dots m}\}$;

Parameters to be learned: γ, β

Output: $\{y_i = \text{BN}_{\gamma,\beta}(x_i)\}$

$$\mu_{\mathcal{B}} \leftarrow \frac{1}{m} \sum_{i=1}^m x_i \quad // \text{ mini-batch mean}$$

$$\sigma_{\mathcal{B}}^2 \leftarrow \frac{1}{m} \sum_{i=1}^m (x_i - \mu_{\mathcal{B}})^2 \quad // \text{ mini-batch variance}$$


$$\hat{x}_i \leftarrow \frac{x_i - \mu_{\mathcal{B}}}{\sqrt{\sigma_{\mathcal{B}}^2 + \epsilon}} \quad // \text{ normalize}$$

$$y_i \leftarrow \gamma \hat{x}_i + \beta \equiv \text{BN}_{\gamma,\beta}(x_i) \quad // \text{ scale and shift}$$

Algorithm 1: Batch Normalizing Transform, applied to activation x over a mini-batch.

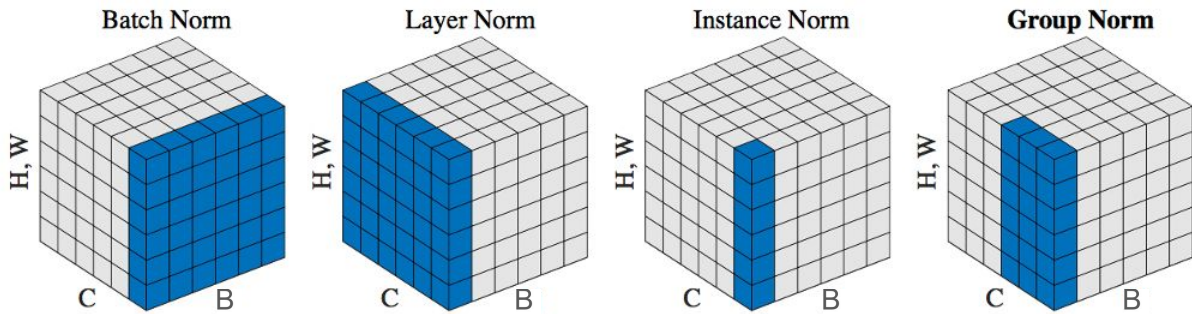
Discuss!

What is the dimension of the mean when you compute the batch norm of a volume of dimension $(b \times c \times h \times w)$?



Normalization Layers

- Normalization layers improve training stability
- Can train with larger learning rates
 - Faster training
- A large learning rate acts as an implicit regularizer
 - Better generalization
- Normalization can also be applied across different dimensions for different use cases



Convolutional Neural Networks (CNNs)

✓ Convolutions

Maintain spatial relation between pixels
Reduce number of parameters through weight sharing

✓ Pooling


Captures key information from across different areas of the feature maps
Together with convolutions allows for translational invariance

✓ BatchNorm


Increases speed and stability of training




input image





Convolutional Neural Networks (CNNs)



Convolutional Neural Networks (CNNs)



Convolutional Neural Networks (CNNs)



Convolutional Neural Networks (CNNs)

✓ Convolutions

Maintain spatial relation between pixels
Reduce number of parameters through weight sharing

✓ Pooling

Captures key information from across different areas of the feature maps
Together with convolutions allows for translational invariance

✓ BatchNorm

Increases speed and stability of training



input image





Image Classification



Practical Guide

- Input image dimensions is divisible by 2
- Small conv filters (3x3 or 5x5)
- Zero padding is used to maintain spatial resolution
- Max pooling for downsampling
- Pooling layers have a receptive field of 2 and stride of 2

Summary

- CNNs are primarily designed to process and analyze visual data, such as images and videos.
- Key components: convolution layers, pooling layers, activation functions, normalization layers
- Advantages:
 - Translational Invariance
 - Parameter sharing
 - Feature learning
- Can be trained with backprop
- Used for tasks such as segmentation, classification, object detection, etc.