





# CS 445

## Introduction to Machine Learning

# Convolution Neural Networks

Instructor: Dr. Kevin Molloy



# Announcements

## PA 3

- Due on November 13<sup>th</sup> at 5:00 pm
- Good goal to complete dense models by the end of the weekend

## Review of Last Class

- Softmax and One-Hot Encoding (multiclass classification with ANN)
- Using Convolutions on Images
  - Filters/kernels enable the NN to learn the features that are important in the image for classifying
  - These kernels provide "position invariance" to the features (the pattern can be detected anywhere in the picture)

# Learning Objectives

- Reducing the Number of Parameters
- Preventing Overfitting in NNs (and CNNs)
- Data Normalization
- Experimental Design for NNs

# Strides and Padding

10	10	10	0	0	0
10	10	10	0	0	0
10	10	10	0	0	0
10	10	10	0	0	0
10	10	10	0	0	0
10	10	10	0	0	0
10	10	10	0	0	0

	10	10	10	0	0	0	
	10	10	10	0	0	0	
	10	10	10	0	0	0	
	10	10	10	0	0	0	
	10	10	10	0	0	0	
	10	10	10	0	0	0	
	10	10	10	0	0	0	

- Corners only appear once. Solution is to pad the image (see left)

Example of stride = 2

10	10	10	0	0	0
10	10	10	0	0	0
10	10	10	0	0	0
10	10	10	0	0	0
10	10	10	0	0	0
10	10	10	0	0	0

10	10	10	0	0	0
10	10	10	0	0	0
10	10	10	0	0	0
10	10	10	0	0	0
10	10	10	0	0	0
10	10	10	0	0	0

Recall  $(W - F + 2P/S) + 1$

# Basic CNN

4	3	3	...	2
255	42	125	...	211
123	98	111	...	48
...	...	...	...	...
54	192	185	...	6

$a^{[0]}$

100 x 100 x 1

$w_1$	$w_2$	$w_3$
$w_4$	$w_5$	$w_6$
$w_7$	$w_8$	$w_9$

$w_1$	$w_2$	$w_3$
$w_4$	$w_5$	$w_6$
$w_7$	$w_8$	$w_9$

$w_1$	$w_2$	$w_3$
$w_4$	$w_5$	$w_6$
$w_7$	$w_8$	$w_9$

$w_1$	$w_2$	$w_3$
$w_4$	$w_5$	$w_6$
$w_7$	$w_8$	$w_9$

4 filters  
Shape of W (3,3,4)

$$\text{ReLU}\left( \begin{array}{cccc} 4 & 3 & \dots & 2 \\ 255 & 42 & \dots & 211 \\ \dots & \dots & \dots & \dots \\ 54 & 192 & \dots & 6 \end{array} + b_1^{[1]} \right) =$$

		...	
		...	
...	...	...	...
		...	

$$\text{ReLU}\left( \begin{array}{cccc} 4 & 3 & \dots & 2 \\ 255 & 42 & \dots & 211 \\ \dots & \dots & \dots & \dots \\ 54 & 192 & \dots & 6 \end{array} + b_2^{[1]} \right) =$$

		...	
		...	
...	...	...	...
		...	

$$\text{ReLU}\left( \begin{array}{cccc} 4 & 3 & \dots & 2 \\ 255 & 42 & \dots & 211 \\ \dots & \dots & \dots & \dots \\ 54 & 192 & \dots & 6 \end{array} + b_3^{[1]} \right) =$$

		...	
		...	
...	...	...	...
		...	

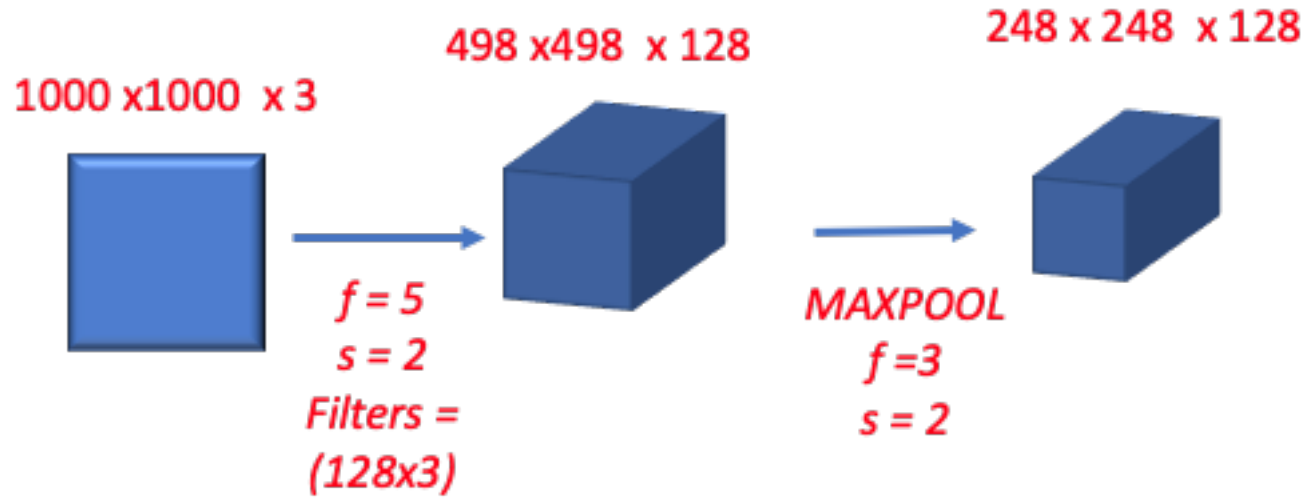
$$\text{ReLU}\left( \begin{array}{cccc} 4 & 3 & \dots & 2 \\ 255 & 42 & \dots & 211 \\ \dots & \dots & \dots & \dots \\ 54 & 192 & \dots & 6 \end{array} + b_4^{[1]} \right) =$$

		...	
		...	
...	...	...	...
		...	

$n - f + 1$   
98 x 98 x 4

These are the activations  
for the next layer ( $a^{[1]}$ )

# MaxPooling



# Pooling

Goal: Reduce number of parameters in the model

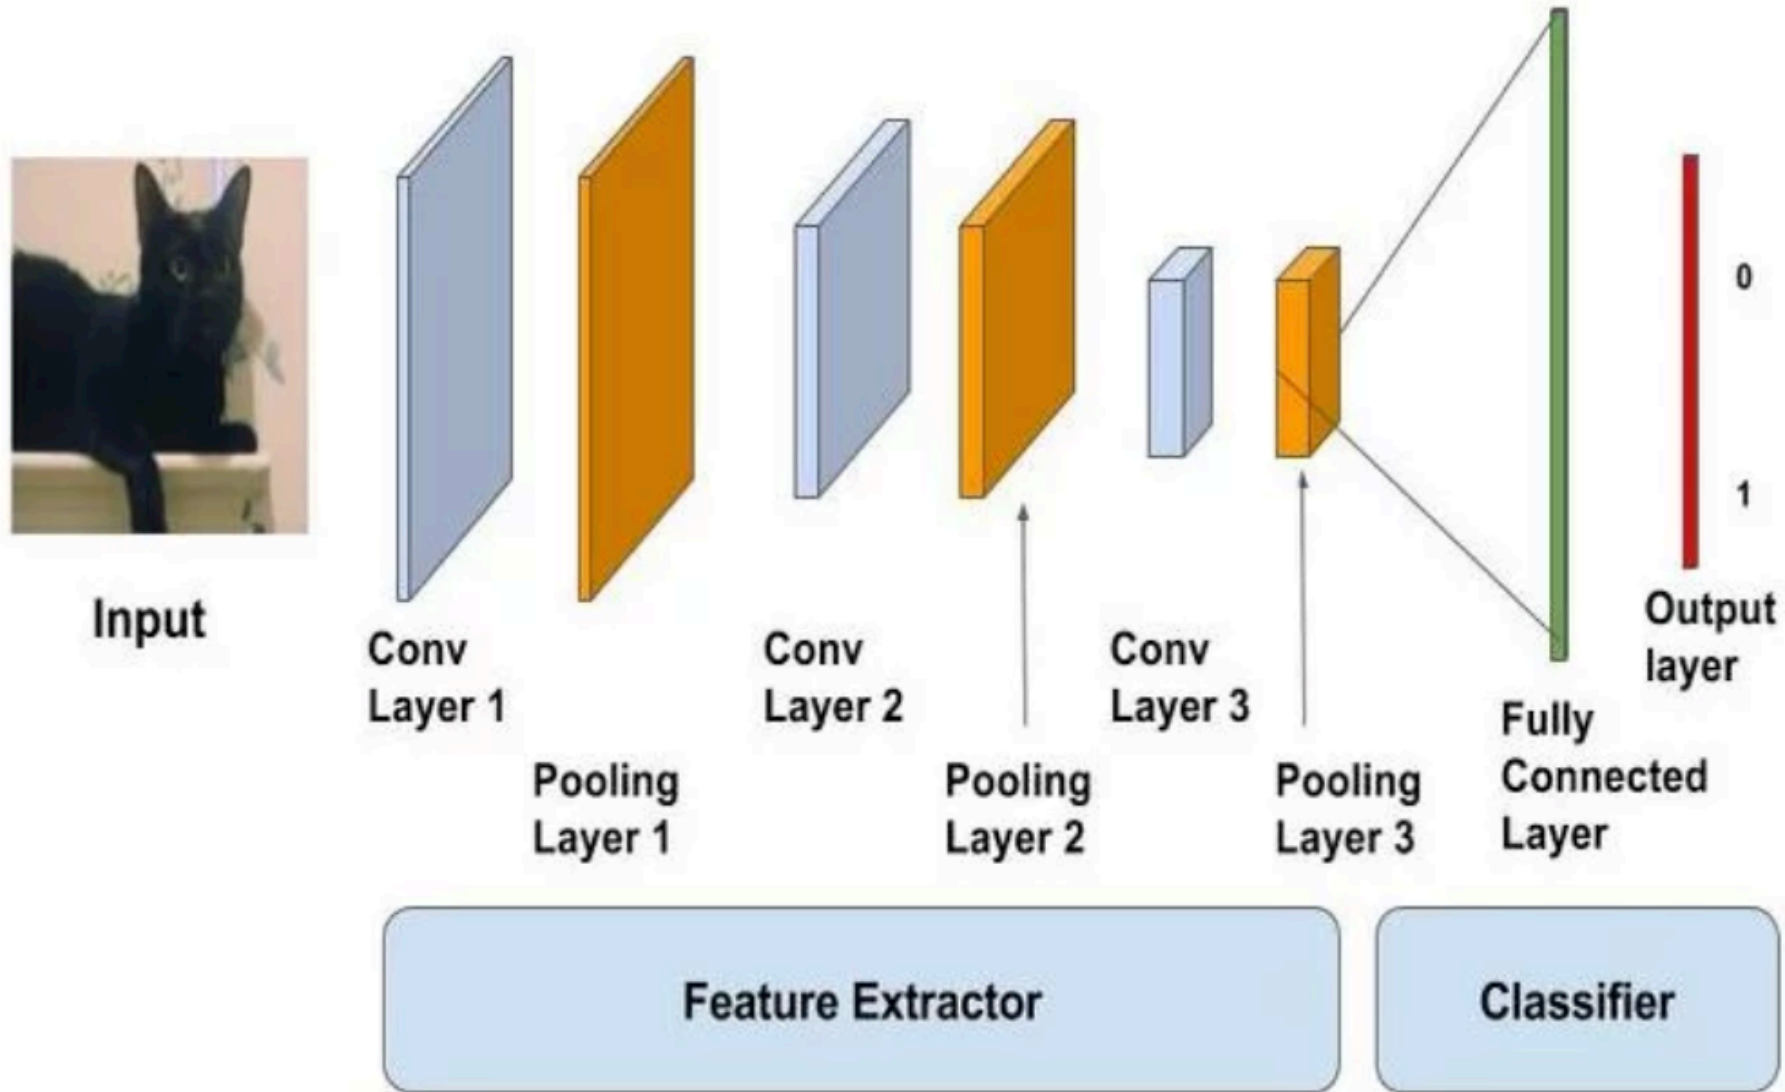


$$480 \times 640 = 307,200$$



$$84 \times 111 = 9,324$$

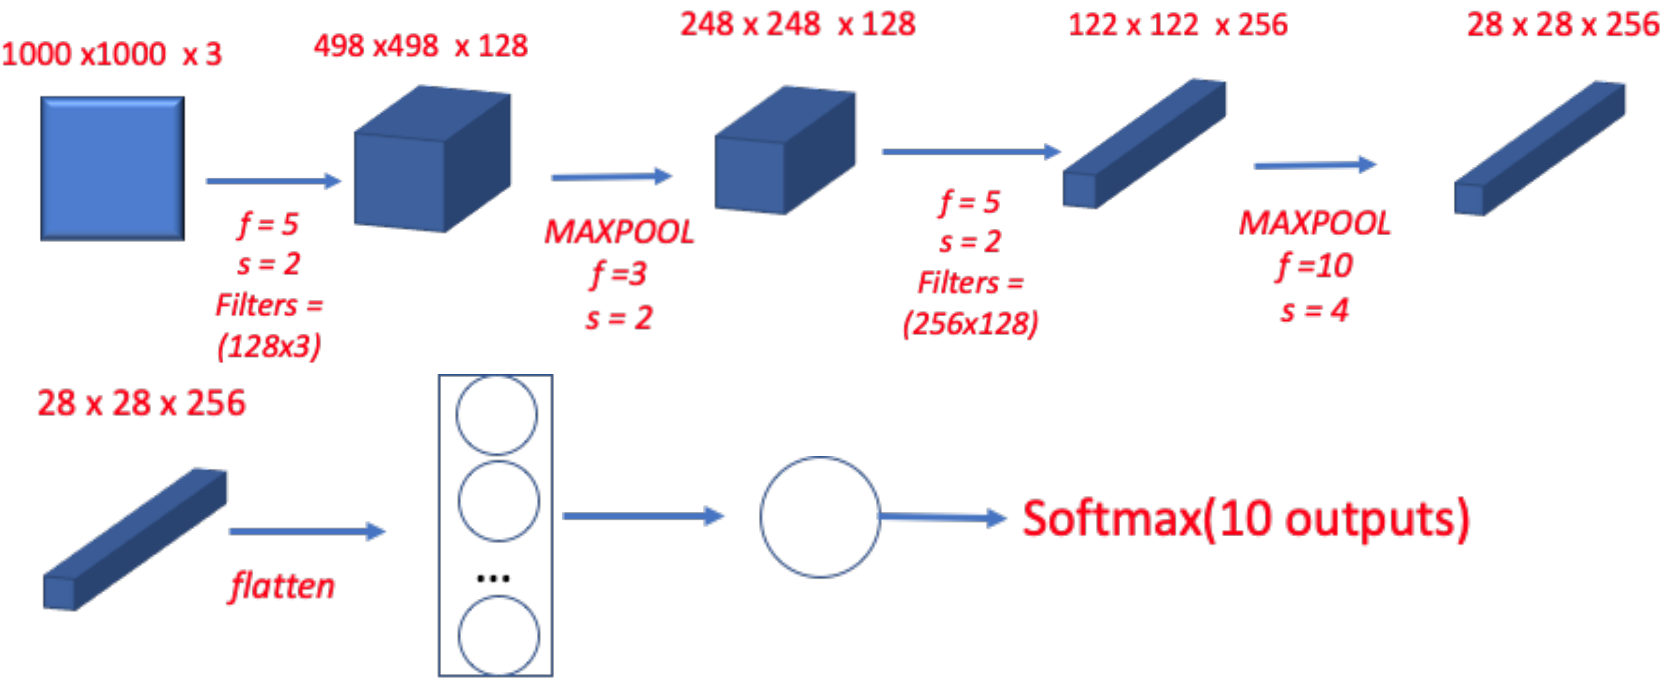
Think of it as down sampling an image.



**Trend:** Each layer shrinks the dimensions of the data/image but the number of filters increase



# Why Convolutions Work



**Parameter sharing:** Each filter is used by all the data from the previous layer. Each act as a feature detector

**Translation Invariance:** Maxpooling helps find the features wherever they reside in the input data.

# Normalization

Do we need to normalize data for NNs?

# Overfitting


Ideas on Preventing overfitting?

What about Ensemble learning?


Dropout layers

```
Dropout layers (Mode.add(Dropout(0.2)))
```

# Image Augmentation to Prevent Overfitting



# Image Augmentation to Prevent Overfitting



```
from keras.preprocessing.image import
ImageDataGenerator

train_gen = ImageDataGenerator(rotation_range=8,
                               width_shift_range=0.08, shear_range=0.3,
                               height_shift_range=0.08, zoom_range=0.08 )

training_set= train_gen.flow(X_train, y_train,
                             batch_size=64)

classifier.fit_generator(training_set,
                        steps_per_epoch=60000//64, epochs=5)
```

# Hyperparameters

What are the hyper-parameters for CNNs?

Number of epochs?

Dropouts

Maxpooling and network size and shape