



Introduction to Machine Learning and Tensorflow

Manu Mathew Thomas
Creative Coding Lab



What is Machine Learning ?

A class of algorithms which learns to performs a specific task based on sample data and without any explicit instruction

Example: Classifying whether an email is a spam or not



What is Machine Learning ?

A class of algorithms which learns to performs a specific task based on sample data and without any explicit instruction

Example: Classifying whether an email is a spam or not

Traditional Algorithms

```
if mail contains "..."  
{  
  
}  
else if sender == "  
{  
  
}  
else if ...  
{  
  
}
```

What is Machine Learning ?

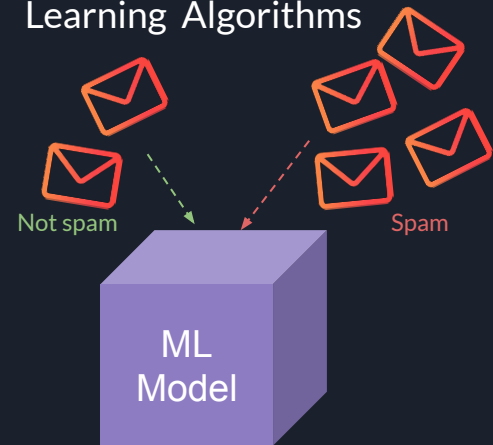
A class of algorithms which learns to performs a specific task based on sample data and without any explicit instruction

Example: Classifying whether an email is a spam or not

Traditional Algorithms

```
if mail contains "..."  
{  
  
}  
else if sender == "  
{  
  
}  
else if ...  
{  
  
}
```

Learning Algorithms



What is Machine Learning ?


A class of algorithms which learns to performs a specific task based on sample data and without any explicit instruction

Example: Classifying whether an email is a spam or not

Traditional Algorithms

```
if mail contains "..."  
{  
  
}  
else if sender == "  
{  
  
}  
else if ...  
{  
  
}
```


Learning Algorithms



Types of Machine Learning


Supervised learning - find patterns and insights from a labelled dataset

Regression



- Linear regression
- Support Vector
- Decision tree
- Random Forest
- Neural Networks

Classification

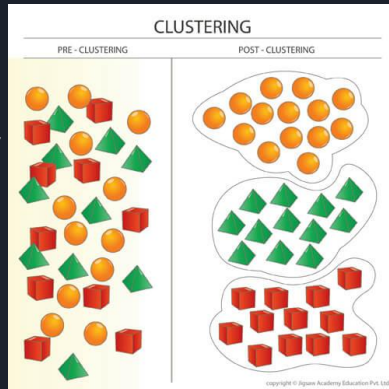


- K Nearest Neighbor
- Naive Bayes
- Support Vector Machine
- Decision tree
- Random Forest
- Neural Networks

Types of Machine Learning


Unsupervised learning - find patterns and insights from an unlabelled dataset

Clustering



- K Mean
- K Nearest Neighbor
- SVD
- Neural Networks

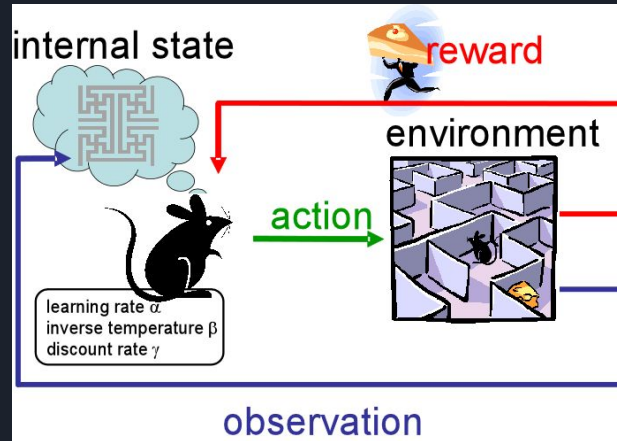
Feature Extraction



- Principal Component Analysis
- Gaussian Mixture Model
- Hidden Markov Model
- Neural Networks

Types of Machine Learning


Reinforcement learning - agent does an action to increase the reward



Neural Networks

NNs are a collection of small processing units (neurons) arranged in a hierarchical fashion



Each processing units is a combination of a linear function followed by a nonlinear function



Neural Networks

NNs are a collection of small processing units (neurons) arranged in a hierarchical fashion


Each processing units is a combination of a linear function followed by a nonlinear function



Neural Networks


NNs are a collection of small processing units (neurons) arranged in a hierarchical fashion

Each processing units is a combination of a **linear function** followed by a nonlinear function



Linear Function: $Wx + b$

$Wx + b$ is the equation for a straight line ($y = Mx + c$) where M is the slope and c is the y-intercept




$Wx + b$ is preferred because:

- Straight lines are useful to model decision boundaries
- It's easier to work with

Neural Networks

NNs are a collection of small processing units (neurons) arranged in a hierarchical fashion



Each processing units is a combination of a linear function followed by a **nonlinear function**



Non-linear function or Activation function: $\sigma(Wx + b)$

σ activates neuron based on the output of the linear function


Creates non-linearity in the network



Neural Networks

NNs are a collection of small processing units (neurons) arranged in a hierarchical fashion


Each processing units is a combination of a linear function followed by a nonlinear function



Activation Functions


Sigmoid

$$\sigma(x) = \frac{1}{1+e^{-x}}$$




Leaky ReLU

$$\max(0.1x, x)$$



tanh

$$\tanh(x)$$



Maxout

$$\max(w_1^T x + b_1, w_2^T x + b_2)$$


ReLU

$$\max(0, x)$$



ELU

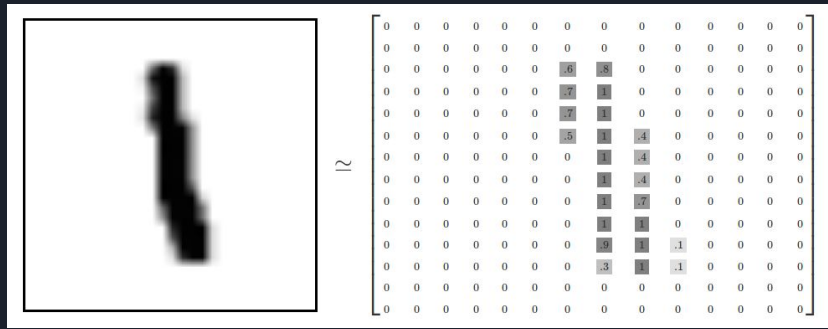
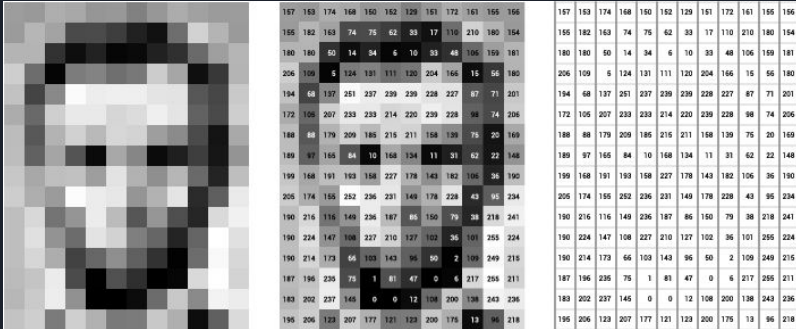
$$\begin{cases} x & x \geq 0 \\ \alpha(e^x - 1) & x < 0 \end{cases}$$



Let's build an image classifier

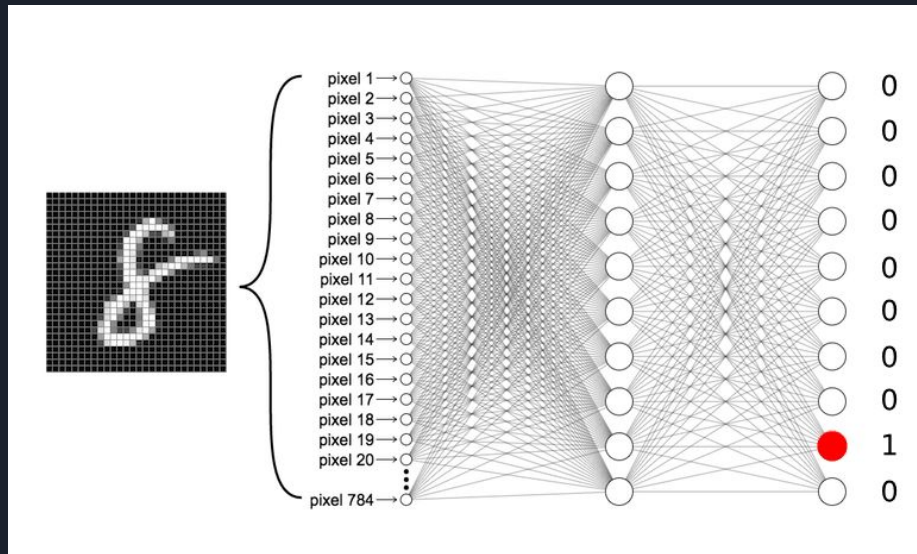
Images are just numbers (usually between 0 and 255)

We scale the images in the range $[0,1]$ as part of feature scaling



Let's build an image classifier

Fully Connected Network(FCN) - Each neuron in one layer is connected to every neuron in the next layer in the next layer



No spatial information

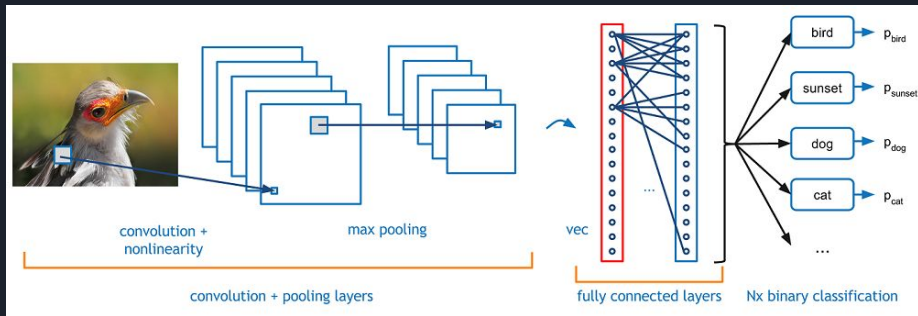
Need a large number of neuron (parameters)

Increases computations and memory footprint

Not commonly used anymore

Let's build an image classifier

Convolutional Neural Network - Each neuron is connected to a local region neurons of previous layer



Looks at spatial information



Reuse neurons (less # parameters)

Lower computations and memory footprint

Let's build an image classifier

Convolution in image processing

Kernel slides over the image and computes new pixel value as a sum/avg of element-wise multiplication



Let's build an image classifier

Convolution in image processing

Examples: <http://setosa.io/ev/image-kernels/>

Some other kernel examples

1	1	1
1	1	1
1	1	1

Unweighted 3x3
smoothing kernel

0	1	0
1	4	1
0	1	0

Weighted 3x3 smoothing
kernel with Gaussian blur

0	-1	0
-1	5	-1
0	-1	0

Kernel to make
image sharper

-1	-1	-1
-1	9	-1
-1	-1	-1

Intensified sharper
image



Gaussian Blur



Sharpened image


Let's build an image classifier

Padding

3_0	3_1	2_2	1	0
0_2	0_2	1_0	3	1
3_0	1_1	2_2	2	3
2	0	0	2	2
2	0	0	0	1

without padding, $5 \times 5 \rightarrow 3 \times 3$


12.0	12.0	17.0
10.0	17.0	19.0
9.0	6.0	14.0




with padding, $5 \times 5 \rightarrow 5 \times 5$

Let's build an image classifier

Stride





Stride - 1



Stride - 2


Let's build an image classifier

Deconvolution



Let's build an image classifier

Classifying a number using hand-made image kernels



*


1	1	1
0	0	0
0	0	0



True

Let's build an image classifier

Classifying a number using hand-made image kernels



*


0	0	1
0	1	0
1	0	0



True

Let's build an image classifier

Classifying a number using hand-made image kernels



*

0	1	0
1	0	0
0	1	0




False

Let's build an image classifier

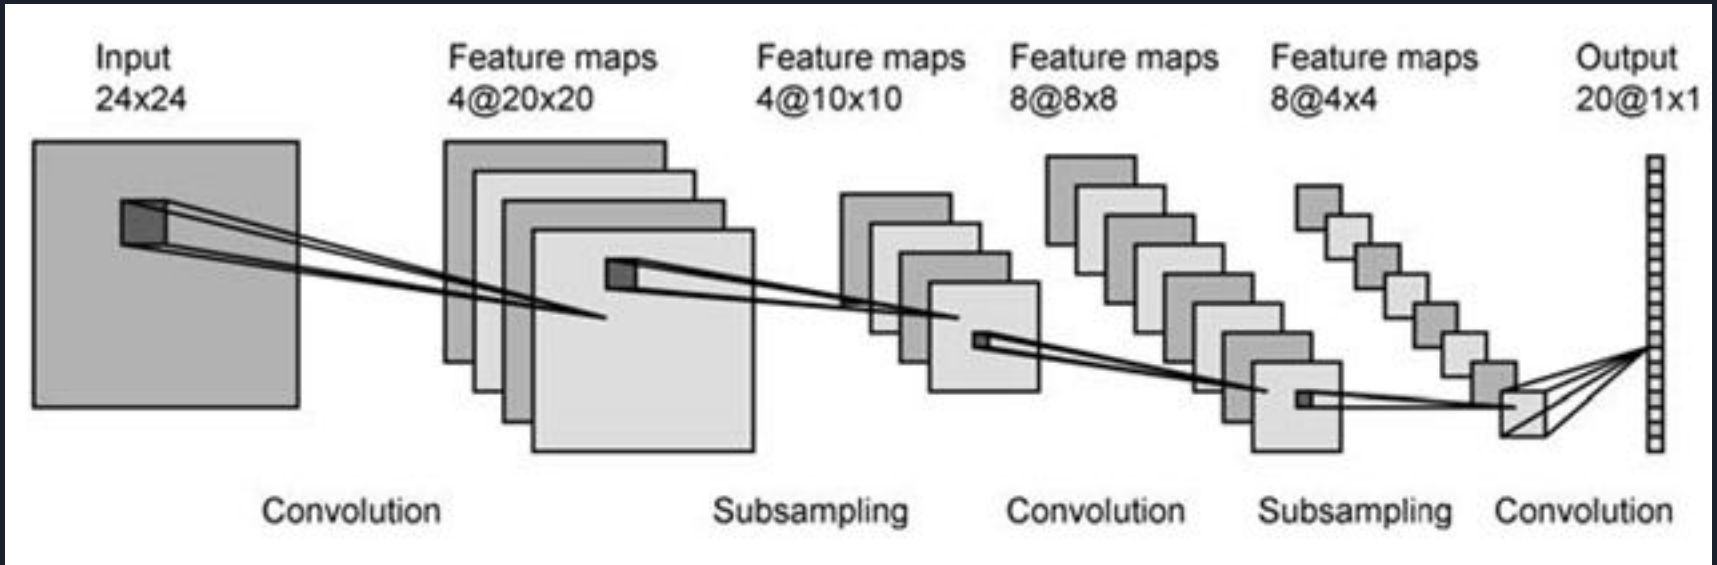
Classifying a face using hand-made image kernels

Hard to build complex kernels




Let's build an image classifier

Kernels/Filters are learnable parameters in a CNN





Let's build an image classifier

Kernels/Filters are learnable parameters in a CNN





Let's build an image classifier

CNNs for RGB images



Let's build an image classifier

CNNs for RGB images



Let's build an image classifier

Learned kernels/filters



Let's build an image classifier


Feature maps - output of each convolution (linear function)



Input

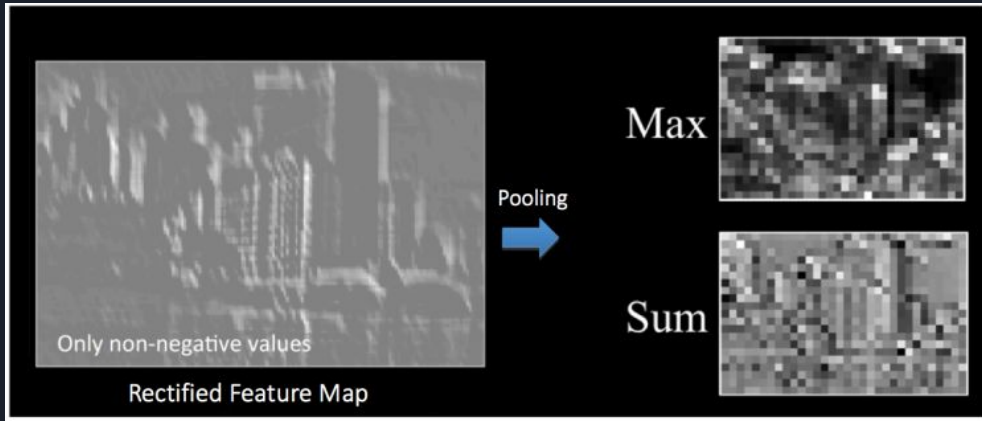
Let's build an image classifier

Activation function(nonlinear function)



Let's build an image classifier

Pooling/Downsampling




Pooling is good for extracting the most important features

Reduces the no. of parameters(weights) in the next layer

Another alternative is stride = 2 or 3

Let's build an image classifier


Our simple network



5 Layers are not deep enough (but CPU friendly)


Let's build an image classifier

Our simple network - Training




Let's build an image classifier


Our simple network - Training



64 x 64 x 3



16 kernels



16



16



32



2



Iteration 0

Output	Ground truth	
0.75	0	Cat
0.25	1	Dog

$\min(\text{distance}(\text{output}, \text{gt}))$


Let's build an image classifier

Our simple network - Training




64 x 64 x 3

Conv1



16 kernels

Conv2



16

Conv3



16

FC1



32

FC2



2



Iteration 500

Output	Ground truth
0.40	0 Cat
0.60	1 Dog

$\min(\text{distance}(\text{output}, \text{gt}))$

Cross entropy is a distance function(kind of) for probability distributions


Let's build an image classifier

Our simple network - Training




64 x 64 x 3

Conv1



16 kernels

Conv2



16

Conv3



16

FC1



32

FC2



2



Iteration 1000

Output	Ground truth
0.05	0 Cat
0.95	1 Dog

$\min(\text{distance}(\text{output}, \text{gt}))$

Cross entropy is a distance function(kind of) for probability distributions


Let's build an image classifier

Our simple network - inference



64 x 64 x 3

Conv1



16 kernels

Conv2



16

Conv3



16

FC1



32

FC2



2






Dog

Generative CNN - Autoencoder



Conv1 Conv2 Conv3 Conv4 Deconv1 Deconv2 Deconv3



Encoder

Decoder

Convolution + Bottleneck extracts the most significant features from the input to reconstruct the output

Generative CNN - Autoencoder

Iteration 0

Conv1

Conv2

Conv3

Conv4




Deconv1

Deconv2

Deconv3

Output

Ground truth



Encoder

Decoder

$\min(\text{distance}(\text{output}, \text{gt}))$

L1 and L2 norms are used for computing pixel distance

Generative CNN - Autoencoder

Iteration 10000

Conv1

Conv2

Conv3

Conv4




Deconv1

Deconv2

Deconv3

Output

Ground truth



Encoder



Decoder

$\min(\text{distance}(\text{output}, \text{gt}))$

L1 and L2 norms are used for computing pixel distance

Generative CNN - Autoencoder

Iteration 0

Conv1

Conv2

Conv3

Conv4




Deconv1

Deconv2

Deconv3

Output

Ground truth



Encoder



Decoder

$\min(\text{distance}(\text{output}, \text{gt}))$

L1 and L2 norms are used for computing pixel distance

Generative CNN - Autoencoder

Iteration 10000

Conv1

Conv2

Conv3

Conv4


Deconv1

Deconv2

Deconv3

Output

Ground truth




Encoder

Decoder

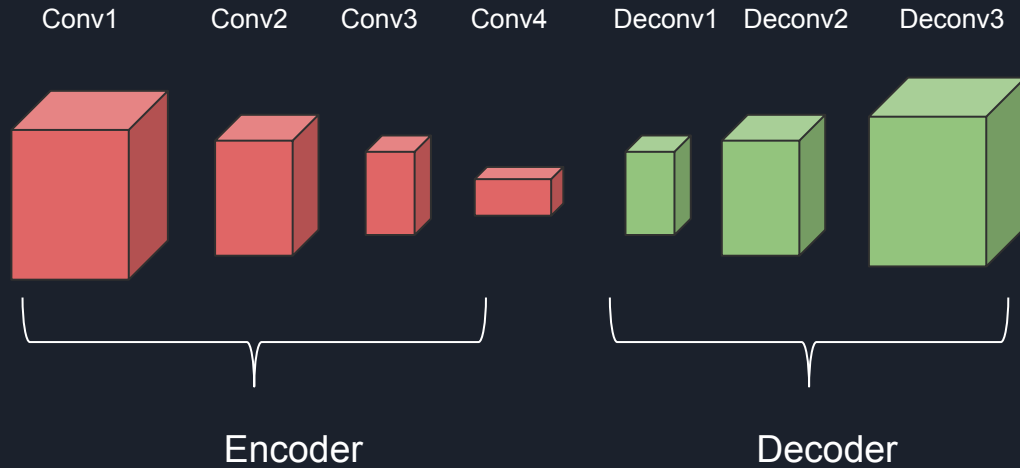
$\min(\text{distance}(\text{output}, \text{gt}))$

L1 and L2 norms are used for computing pixel distance

Generative CNN - Autoencoder





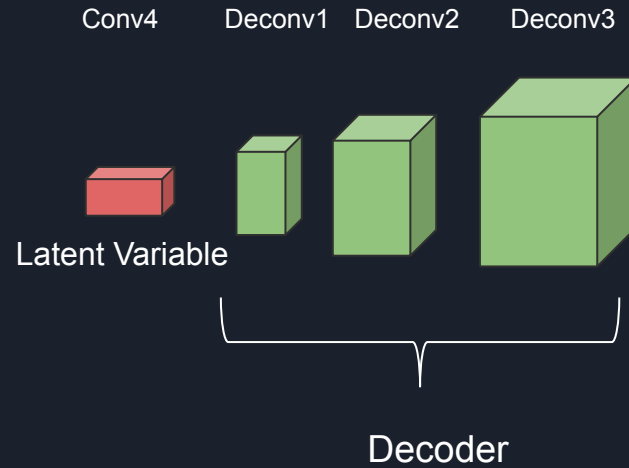
Ansel Adams, Yosemite Valley Bridge



Ansel Adams, Yosemite Valley Bridge


Dataset not used in training

Generative CNN - Autoencoder



Changing the latent variable randomly will generate new images

Generative CNN - Variational Autoencoder





Next time

Setup tensorflow environment

Building a simple image classifier in tensorflow

Maybe Gans?