

Deep Learning & Neural Networks

Dr. Shahira Ahmed

DS710

Lecture 1





Faculty of Engineering
جامعة ميدوcean
المدينة التعليمية


Deep Learning & Neural Networks

DS710


Lecture 1: Introduction

- Artificial Intelligence
- Machine Learning
- ML Learning life Cycle
- Perceptron-Neuron

Artificial intelligence (AI) makes it possible for machines to mimic human intelligence.



AI





المملكة العربية السعودية
الجمهورية العربية السورية
الجمهورية الجزائرية الديمقراطية الشعبية

Learning


$$\exists u. \forall v. \{q(v, u) \& \sim[[r(v) \vee p(v)] \& s(u, v)]\}$$
$$q(w, a) \& \{[\sim r(v) \& \sim p(v)] \vee \sim s(a, v)\}$$

What is Machine Learning?


- Samuel's Checkers-Player “Machine Learning: Field of study that gives computers the ability to learn without being explicitly programmed.” - Arthur Samuel (1959)



- **Traditional Programming**




- **Machine Learning**



When Do We Use Machine Learning?

- ML is used when:
 - Human expertise does not exist (navigating on Mars)
 - Humans can't explain their expertise (speech recognition)
 - Models must be customized (personalized medicine)
 - Models are based on huge amounts of data (genomics)






وزارة التعليم
المملكة العربية
السعودية

Applications of ML

E-commerce & Recommender Systems : Amazon, netflix, ...

Inspired by your shopping trends



Related to items you've viewed [See more](#)



- Nevada made it legal for autonomous cars to drive on roads in June 2011
 - As of 2013, four states (Nevada, Florida, California, and Michigan) have legalized autonomous cars



Penn's Autonomous Car
(Ben Franklin Racing Team)



Waymo undergoing testing in the
San Francisco Bay Area



Roborace autonomous racing car
on display at the 2017 New York City
ePrix



Alphabet's DeepMind uses A.I. to detect signs of eye disease



Jonas Hafner | EyeEm | Getty Images






Alphabet's DeepMind uses A.I. to detect signs of eye disease



Jonas Hafner | EyeEm | Getty Images


- for image classification, the key breakthrough came in 2012 when a team led by Geoff Hinton at the University of Toronto won the ImageNet Large Scale Visual Recognition Challenge (ILSVRC) with a deep convolutional neural network.
- The competition involves classifying images into one of a thousand categories and issued as a benchmark to compare the latest state-of-the-art techniques. The deep learning model had an error rate of 16%—a massive improvement on the next best model, which only achieved a 26.2% error rate.
- The 2015 winner achieved superhuman performance for the first time, with an error rate of 4%, and the current state-of-the-art model achieves an error rate of just 2%.








Deep Fake?




[Barack Obama: Intro to Deep Learning | MIT 6.S191 - YouTube](#)

Living portraits



<https://www.youtube.com/watch?v=P2uZF-5F1wl>

Application of Machine Learning



<https://www.youtube.com/watch?v=mlZYRwJ2oJg>

Application of Machine Learning



"man in black shirt is playing guitar."




"construction worker in orange safety vest is working on road."





"two young girls are playing with lego toy."

“an armchair in the
shape of an
avocado”




“a living room with two white armchairs and a painting of the colosseum. the painting is mounted above a modern fireplace”






What is Deep Learning?

- Deep learning is a class of machine learning algorithm that uses multiple stacked layers of processing units to learn high-level representations from unstructured data.
- Deep Learning involves taking large volumes of structured or unstructured data and using complex algorithms to train neural networks.





[AI, ML, or DL – learn what it means \(datasciencecentral.com\)](https://datasciencecentral.com)

Structured vs Unstructured Data


STRUCTURED DATA				
id	age	gender	height (cm)	location
0001	54	M	186	London
0002	35	F	166	New York
0003	62	F	170	Amsterdam
0004	23	M	164	London
0005	25	M	180	Cairo
0006	29	F	181	Beijing
0007	46	M	172	Chicago

UNSTRUCTURED DATA		
		This service is terrible!
		Your website is great!
images	audio	text


- Structured, tabular data as input, arranged into columns of features that describe each observation.
 - For example, a person's age, income, and number of website visits in the last month are all features that could help to predict if the person will subscribe to a particular online service in the coming month.

- Unstructured data refers to any data that is not naturally arranged into columns of features, such as images, audio, and text. There is of course spatial structure to an image, temporal structure to a recording, and both spatial and temporal structure to video data, but since the data does not arrive in columns of features, it is considered unstructured.

Classic Machine Learning




Deep Learning



Aspect	Machine Learning	Deep Learning
Data Volume	Hundreds or thousands of data	Millions of data
Computational Cost	ML model doesn't need a lot of computational power (Using CPU is often sufficient)	DL model needs a lot of computational power (The use of GPU is often necessary)
Training Time	ML model takes less time to train	DL model takes more time to train
Feature Engineering	Needs to be done explicitly by human	DL model can automatically learn important feature during training
Interpretability	ML models' behavior are easier to interpret	DL models' behavior are more difficult to interpret
Application Examples	Customer segmentation, Recommendation system, Fraud detection	Natural Language Processing, Computer Vision

<https://www.stratascratch.com/blog/data-science-vs-machine-learning-vs-deep-learning-the-difference/#:~:text=Machine%20learning%20models%20work%20with,comparison%20with%20deep%20learning%20models.>

Why now?



No one knows how they work:
we cannot predict when outputs will be
catastrophically wrong



“panda”

+ 0.007 ×




=



“gibbon”

No one knows how they work:
we cannot predict when outputs will be
catastrophically wrong



Limitation of ML



Google apologises for Photos app's racist blunder

🕒 1 July 2015

Google 'fixed' its racist algorithm by removing gorillas from its image-labeling tech

[Lecture 1 Fall 2020.pdf](#) | Powered by Box



diri noir avec banan @jackyalcine · Jun 29
Google Photos, y'all [redacted] My friend's not a gorilla.

👍 813 ⭐ 394 +1 ...

TWITTER

Women less likely to be shown ads for high-paid jobs on Google, study shows

theguardian

Microsoft 'deeply sorry' for racist and sexist tweets by AI chatbot


theguardian

Facebook fires human editors, algorithm immediately posts fake news

ars TECHNICA

- Data
 - Data collection
 - Data Labeling
 - Data Cleaning
 - Data processing

- Data
 - Data collection
 - Data Labeling
 - Data processing



- (1) Companies and Proprietary Data Sources
 - Large companies like Facebook, Google, Amazon, American Express, and Blue Cross have amazing amounts of data .
 - Many responsible companies like The New York Times, Twitter, Facebook, and Google do release certain data, typically by rate-limited application program interfaces (APIs).
 - Some organizations do provide bulk downloads of interesting data for off line analysis, as with the Google Ngrams, IMDb, and the taxi fare data sets.

- (2) Government Data Sources
 - The website <http://Data.gov> is an initiative by the federal government to centrally collect its data sources, and at last count points to over 100,000 data sets!



Where Does Data Come From?

- [Kaggle](#)
- [UCI Machine Learning Repository](#)
- [Google Dataset Search](#)
- [AWS Open Data Registry](#)
- [Wikipedia ML Datasets](#)
- [**Microsoft Datasets**](#)

- DATA: It can be any unprocessed fact, value, text, sound, or picture that is not being interpreted and analyzed.
- Data is the most important part of all Data Analytics, Machine Learning.
- Without data, there is no machine learning.

- Data preprocessing is an important step in the machine learning pipeline. This step can include cleaning and normalizing the data, handling missing values, and feature selection or engineering.



Iris Versicolor



Iris Setosa




Iris Virginica

iris	sepal length	sepal width	petal length	petal width
Iris-virginica	6.500	3.200	5.100	2.000
Iris-virginica	6.500	3.000	5.500	1.800
Iris-virginica	6.500	3.000	5.200	2.000
Iris-versicolor	6.600	2.900	4.600	1.300
Iris-versicolor	6.600	3.000	4.400	1.400
Iris-versicolor	6.700	3.100	4.400	1.400
Iris-versicolor	6.700	3.000	5.000	1.700
Iris-versicolor	6.700	3.100	4.700	1.500
Iris-virginica	6.700	2.500	5.800	1.800
Iris-virginica	6.700	3.300	5.700	2.100
Iris-virginica	6.700	3.100	5.600	2.400
Iris-virginica	6.700	3.300	5.700	2.500
Iris-virginica	6.700	3.000	5.200	2.300
Iris-versicolor	6.800	2.800	4.800	1.400
Iris-virginica	6.800	3.000	5.500	2.100
Iris-virginica	6.800	3.200	5.900	2.300
Iris-versicolor	6.900	3.100	4.900	1.500
Iris-virginica	6.900	3.200	5.700	2.300

iris	sepal length	sepal width	petal length	petal width
Iris-virginica	6.500	3.200	5.100	2.000
Iris-virginica	6.500	3.000	5.500	1.800
Iris-virginica	6.500	3.000	5.200	2.000
Iris-versicolor	6.600	2.900	4.600	1.300
Iris-versicolor	6.600	3.000	4.400	1.400
Iris-versicolor	6.700	3.100	4.400	1.400
Iris-versicolor	6.700	3.000	5.000	1.700
Iris-versicolor	6.700	3.100	4.700	1.500
Iris-virginica	6.700	2.500	5.800	1.800
Iris-virginica	6.700	3.300	5.700	2.100

- Data can be divided into training and testing sets.
- The training set is used to train the model
- The testing set is used to evaluate the performance of the model.

- **Training** The idea is to give a set of inputs(features) and it's expected outputs(labels), so after training, we will have a model (hypothesis) that will then map new data to one of the categories trained on.
- **Prediction** Once our model is ready, it can be fed a set of inputs to which it will provide a predicted output(label).



- Supervised (inductive) learning
 - Given: training data+ desired outputs (labels)
- Unsupervised learning
 - Given: training data (without desired outputs)
- Semi-supervised learning
 - Given: training data + a few desired outputs
- Reinforcement learning
 - Rewards from sequence of actions




وزارة التعليم
المملكة العربية
السعودية




Neuron


Biological Neuron




Artificial Neuron




Artificial Neuron



Artificial Neuron




Recap : Nnets and the brain




- Neural nets are composed of networks of computational models of neurons called perceptrons

So what are neural networks??




- What are these boxes?
 - Functions that take an input and produce an output
 - What are these functions?

The human perspective



- In a human, those functions are computed by the brain...

Example



<https://towardsdatascience.com/what-the-hell-is-perceptron-626217814f53>

- Machine vs Deep Learning
- Applications
- Biological neuron
- Artificial neuron
- Artificial neural network



مجموعة جامعة
الجميع الذكية
للتعليم القابضة

MIDOCEAN
UNIVERSITY
BY EVERYONE'S SMART UNIVERSITY

*Thank
you*



www.midocean.ae

