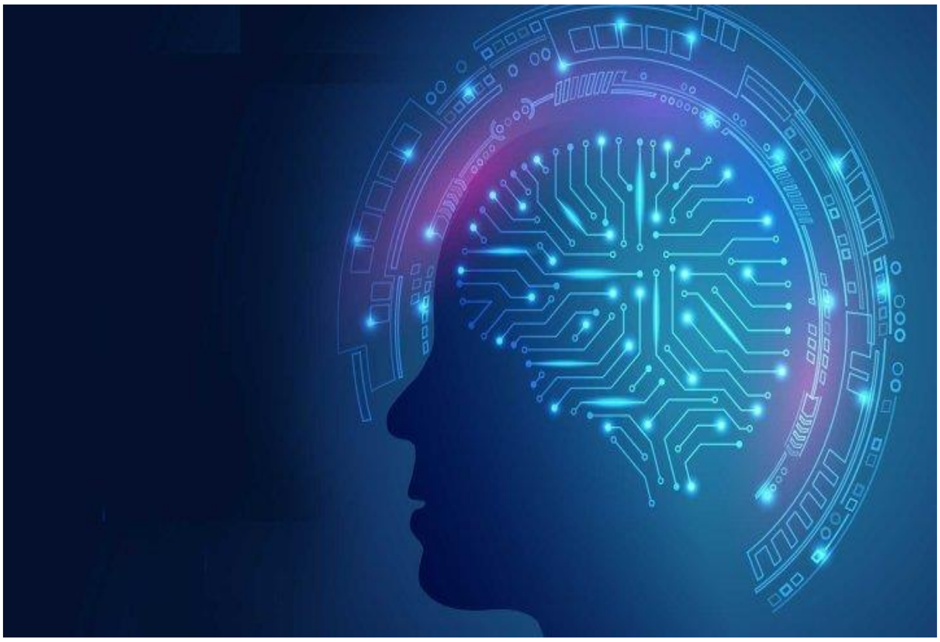

Artificial Intelligence: Predictions with the use of Neural Networks

Mentored and Supervised by:

Fotis Sarantopoulos, Diploma Naval Architecture & Marine Engineering NTUA, INSEAD MiM '23

Research and Writing by CLC's students:

Natalia Hatzigeorgiou / Emmanuel Kontos /

Niki Mazioti / Ilias Sohos / Perseus Georgiadis

Powered by:

TABLE OF CONTENTS

Introduction

Artificial intelligence has been a prominent scientific field for many decades. Machine learning is a subset of artificial intelligence, while artificial neural networks are a type of machine learning. Neural networks are based on the way the human brain works and have become popular for problems of classification, clustering, pattern recognition and regression. In turn, deep learning can be considered as a subset of machine learning. In deep learning, more complex ways of connecting the layers of the neural network are observed, more neurons are encountered, more computational power is needed, and even the possibility of automatic feature extraction is present. Deep machine learning algorithms find great application in big data management with notable success in voice recognition, pattern recognition, computer vision, natural language processing, recommendation systems, etc. The figure below graphically illustrates the above correlations.

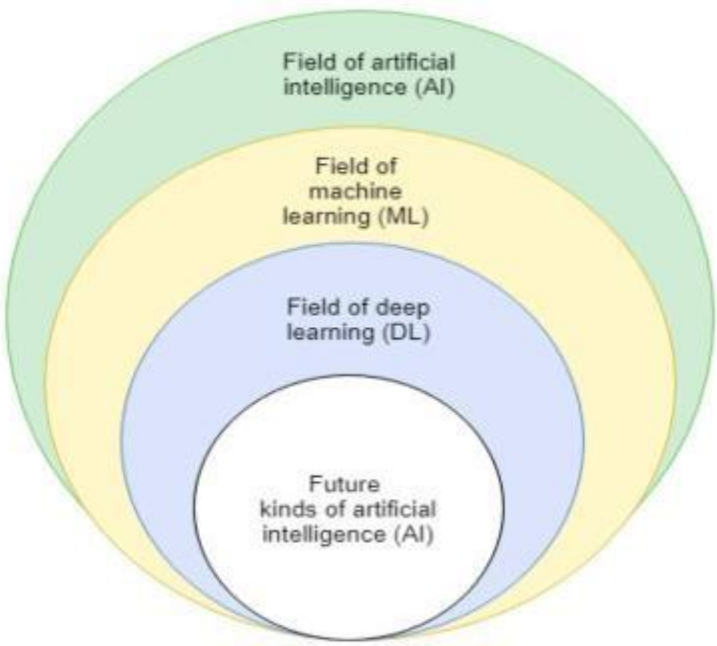

Artificial Intelligence

Artificial intelligence is the replication of human intelligence functions by machines, especially computer systems.

AI applications

AI plays a crucial role in a variety of fields of our everyday lives. Some of them are the following:

1) ***Astronomy***. AI can be very helpful for understanding the universe and solving its complex problems.
2) ***Healthcare.*** AI is being used in the healthcare industry to make better and faster diagnoses than humans. AI can assist doctors with diagnoses and can alert them when a patient's condition is deteriorating, so that medical assistance can be provided before the patient needs to be hospitalized.
3) ***Finance.*** The finance and AI industries are the most compatible. Automation, chatbots, adaptive intelligence, algorithm trading, and machine learning are all being used in financial processes.
4) ***Data Security.*** Data security is critical for every business. AI can help in keeping data safe and secure. Some examples are AEG bot and AI2 Platform, which are used to better determine software bugs and cyber-attacks
5) ***Robotics***. In robotics, artificial intelligence plays a significant role. Typically, general robots are programmed to perform a repetitive task. However, with AI we can create intelligent robots that can perform tasks based on their own experiences rather than being pre-programmed. (Javatpoint, n.d.)

Advantages / Disadvantages of AI

There is a variety of advantages and disadvantages when it comes to AI, that should be both examined in order for a correct conclusion to be made.

Advantages:

- AI handles the information better than humans and it is more accurate
- The machine takes faster decisions than a human
- AI can be used for daily applications, making human’s life easier.

Disadvantages:

- The implementation cost of AI is very high.
- Unemployment increases, as AI replaces human beings.
- Lacking Out of Box Thinking

Machine Learning

Machine learning (ML) is a set of algorithms that aids in the solving of problems, the discovery of patterns in data, and the prediction of output values. Specifically, ML techniques are used to locate patterns in complicated data that would have otherwise been difficult to find; additionally, the hidden patterns and knowledge about a problem can be used to predict future events and perform all kinds of complex decision making.

Types of Machine Learning

The field of Machine Learning develops three modes of learning, similar to the ways by which humans learn: supervised learning, unsupervised learning and reinforcement learning.

● In **Supervised Learning**, for each input vector, there is a desired output vector. The goal is for the system to construct a function f that maps given inputs x to known desired outputs y, with the ultimate goal being the generalization of this function to inputs with unknown outputs. This is achieved by continuously varying the system parameters until the minimum desired error is reached. Error is the difference between the desired value and the output predicted by the system. For better generalisation, the practice of training and testing the results on different subsets of the data set is usually chosen. This avoids overfitting, i.e. the case in which the model trains very well for one particular dataset but fails to generalise to another. Supervised learning problems are divided into classification problems and regression problems. In the first case, the data is categorized into discrete classes (e.g., given an input of a set of symptoms of a patient, the output is 0 for a healthy patient and 1 for a sick patient). In the second case, a continuous value prediction is made (e.g., given an input of some meteorological metrics, the temperature of the next day is predicted).

● In **Unsupervised Learning**, a model is constructed for some sets of inputs in the form of observations without knowing the desired outputs. It is used in Association Analysis and Clustering problems. In Clustering, the goal is to organize samples into groups and the desired outcome is that the objects within a group are more similar to each other than to objects in other groups. The best-known unsupervised learning algorithm and more specifically Clustering algorithm is K-Means, which classifies samples into K classes. It does this in a series of steps. First, it initializes the K class/group centres, computes the distance between the sample and each class centre, classifies the sample into the class with the shortest distance from it, and updates the class centres by computing the mean of all vectors in each class. K-Means terminates after a certain number of iterations of the above process or when the class centres do not move enough in each iteration.

● In **Reinforcement Learning**, the algorithm tries to develop an autonomous agent that improves its performance through direct interaction with its environment. The evaluation is done through a reward function. The agent attempts to find a set of actions in its environment that maximises the reward function by trials and errors. It is mainly used in planning problems, such as robot motion control and optimisation of tasks in factories. It is also worth mentioning the very good effect of reinforcement learning in complex games such as chess.

Future of ML

Currently, ML is used in many industries such as banking and finance, healthcare, transportation, customer service, retail, and agriculture. In the future, ML will continue to be used in those industries and specifically in healthcare, it might facilitate clinical decision making by considering many different variables such as genes, lifestyle factors and the environment, using them for “disease prevention, treatment, and prognostication”. This is a way that ML can be lifesaving and it might something that will greatly benefit humanity in the future.

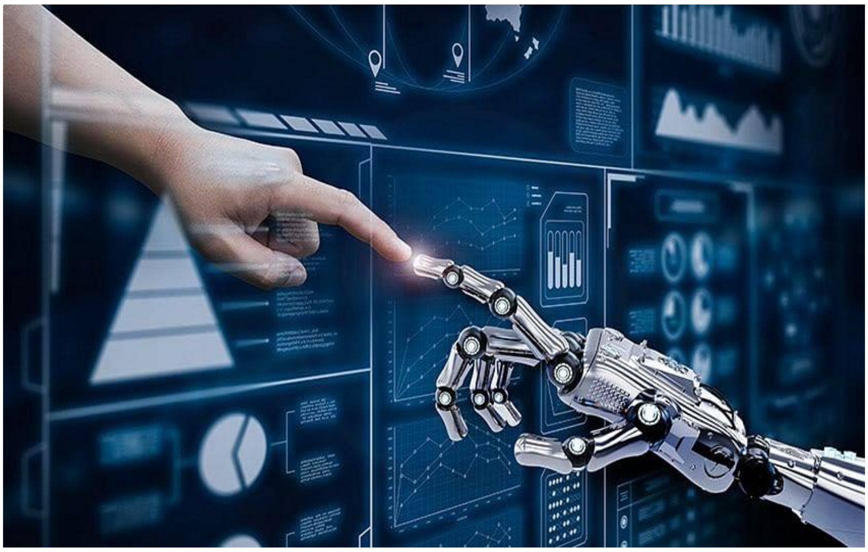

Deep Learning

The term "Deep Learning" refers to a sub-category of machine learning. Deep Learning is a machine learning algorithm utilizing deep neural networks, heavily influenced by, and based on, the human brain, or more precisely, the human brain's neurons. While this concept is not new, it is mostly prevalent in recent years due to the considerable increase of available data and computational power, which make the idea of a Deep Neural Network (DNN) more viable and practical. Compared to ML algorithms, DL algorithms are more specific in their outputs (while ML algorithms only output numerical values, DL algorithms can also output free-form elements), require less human intervention, and get much more accurate as data sizes increase. However, one main problem of DL algorithms is incomprehensability. As thousands of layers of neural networks build upon each other to produce an accurate prediction based on input value, it is extremely difficult for a human to understand and interfere with the algorithm. Thus, other ML algorithms are better suited when there is a smaller amount of data available and human intervention is required, while DL algorithms are better suited for more complex tasks with more available data and computational power.

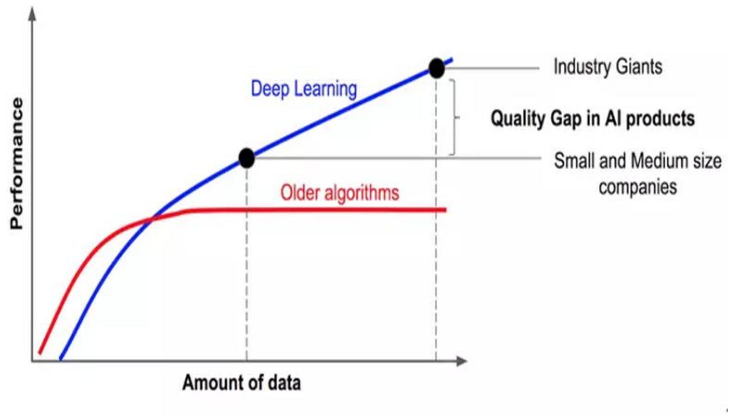

Applications

What really separates DL from other ML algorithms is its various complex aplications. Due to its ability to fit and predict outputs of non-linear functions, DL can be used to solve very complex human problems. DL is used in sectors like medicine, as a "powerful diagnostic tool", in the development and creation of new drugs, and in modern parts of medicine such as telemedicine. Another field where DL is implemented is robotics, where it has unlocked potential for human-like walking and running, automatic combat recovery, automatic aircraft inspection and maintenance and much more. Finally, the field of earthquake prevention saw a tremendous increase in accuracy after utilizing DL to predict earthquakes. Of course, implementations of DL do not end in those three areas.

There are many everyday examples where DL is implemented. One of those examples is autonomous cars, like Tesla's. A Tesla's DL algorithm must perfect accuracy in 3 key areas: computer vision, prediction, and driving policy. Computer vision refers to object detection. The car must swiftly and accurately identify nearby objects like cars, buses, pedestrians, traffic lights/signs etc. Prediction means that the algorithm must be able to anticipate how those nearby objects are going to change. For example, the movement of nearby vehicles, if a pedestrian is crossing the road, or a change in traffic lights. This part is what most rival companies' algorithms fail to perfect, mainly due to lack of training data. Driving policy refers to the actions a human driver would take while on the road, based on the input they take. This includes passsing a slow car, speeding up/slowing down depending on the speed limit etc.

Neural Networks

Artificial neural networks or simply neural networks are an abstract algorithmic construct, which belongs to the family of machine learning, capable of "learning" to perform specific functions by reading examples of them, and without being specially programmed. The structure of the system is inspired by the neurons of the brain, hence its specific name. With artificial neural networks, the scientific community is trying to imitate and model this non-linear behaviour of the brain, which is directly related to decision-making. The methods used in these networks provide a liberated model, which can be adapted to a wide range of applications, being fault-tolerant and able to use parallel and distributed systems for faster extraction of the solution to the problem. They are essentially "black boxes" which "learn" the internal relationship of an unknown system without attempting to "guess" the functions governing the interconnections of that system. They are widely used in areas such as mathematical function approximation, problems of data classification, feature extraction from models, clustering, quantization of vectors, optimization, etc. They are one of the key methodologies in information processing and are applied in almost all areas of science and engineering.

Historical Background

The history of artificial neural networks is highlighted by four main dates. At first, in 1943 Warren S. McCulloch – a neurologist - and Walter Pitts - a mathematician- conducted thorough research on the function of the human brain and the transmission of its complex patterns through its neurons. They did all of this by modeling a neural network with the help of electrical circuits. In this way, a potential connection between the brain's neurons and a binary threshold based on the Boolean logic (binary numbers and true/false statements) was discovered and developed in the future. After that, in 1958, when computers were starting to advance and become more equipped, Frank Rosenblatt – a neurobiologist at the University of Cornell- worked on the perceptron, a computer model which simulates the ability of the brain to recognize items. He was credited with the accomplishment of making a computer to differentiate cards marked on the left and cards marked on the right. One year later, Windrow and Hoff developed the first neural network, which was applied as a solution to a real-life problem, that of minimizing the echoes on phone lines. Moreover, in 1985 the first international Conference on Neural networks was established, with a wide and devoted audience. Since that year, additional proposals on the potential use of neural networks were made.

Today, research on neural networks is advancing and is used by established companies, such as Google, Apple, Instagram, and Facebook. However, further hardware development is necessary to be worked on so that the key to the future of neural networks is provided.

Issues of Neural Networks

Thanks to their huge number of parameters, artificial neural networks can be easily modified to fit almost any complex dataset. This ability has allowed them to find applications in many areas, in which it has been difficult to make progress, such as image recognition, natural language processing, etc. However, sometimes, the complexity of these networks may become a potential weakness, since it can lead to overfitting or underfitting.

1) Overfitting

This situation occurs when the neural network is so closely fitted to the training set that it is difficult to generalize and make predictions for new, unseen data. In practice, detecting that a trained model is overfitting is a difficult task, so it is necessary that some steps should be followed during training. Specifically, it is advisable that a dataset should be divided into three parts – training set, cross-validation set and test set. The model learns by only considering data from the first part, while cross-validation set is used to track progress of training and optimize the model. Furthermore, at the end of the training process, the test set is used, in order to evaluate the performance of the trained network.

Although until recently the most frequently recommended division of dataset would be: 80% training set and 20% test set, when a dataset comprises of millions of entries, these proportions are no longer appropriate. In short, everything depends on the total size of the dataset, and in cases of millions of data, it could even be better to divide the set in 98/2 ratio.

2) Underfitting

This situation occurs when the trained model can neither fit the training data, nor generalize to new unseen data. Detecting an underfit model is obvious, since it will have poor performance on training data, and it is, of course, an unsuitable.

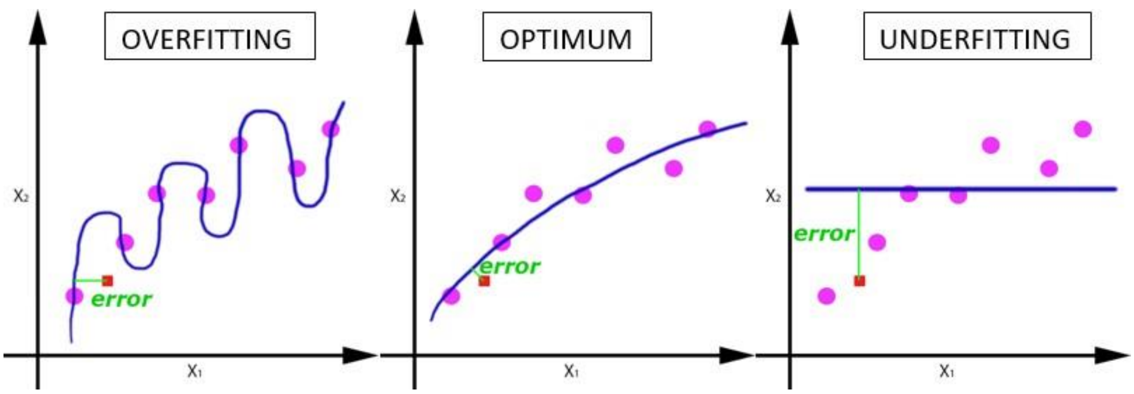

Types of Neural Networks

1. Feedforward Neural Networks

Feed-forward networks, as they are referred to in Greek literature, are the simplest neural networks. In these networks, information flows only in one direction (feed-forward), from the input layer to the hidden layers and from there to the output layer. There are no cycles or iterations in the network. More specifically, in feed-forward networks, the input layer is of equal dimensions to the input vector being fed into the network. Each input node is connected to all nodes in the next layer, while the output layer has equal dimensions to the output y. There can be multiple hidden layers connecting the input layer and the output layer, but in its simplest form, the perceptron connects the input directly to the output. The role of the hidden layers is to extract useful results, helping the network extract higher order information than that of the input data. When all nodes of each layer are connected to all nodes of the next layer, a fully connected neural network exists. Sometimes, the output signals of the intermediate neurons seem incomprehensible to humans, but they lead computers to better results.

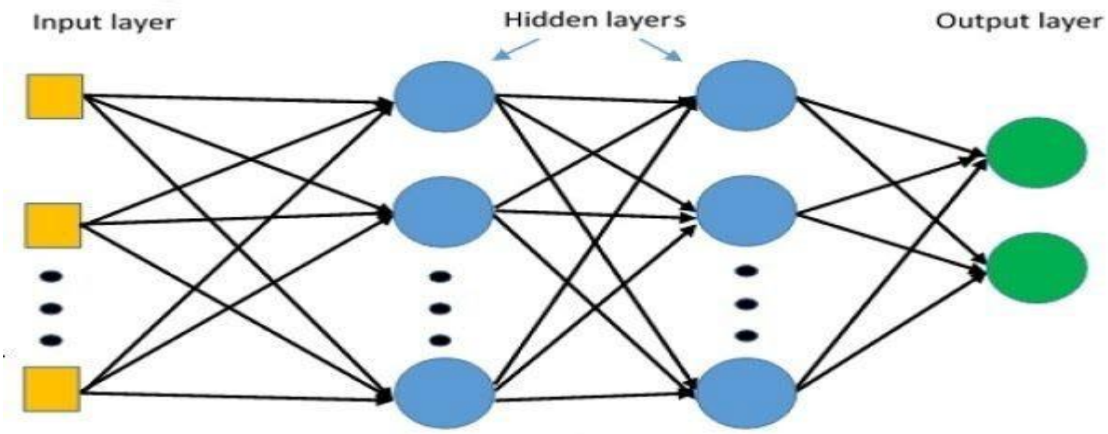

2. Recurrent Neural Networks

Recurrent neural networks (RNNs) are essential for sequential data prediction and therefore they find effective application in time series problems. The basic idea behind RNNs is that for an output to be produced, not only input data is used but also the previous outputs. RNNs differ from traditional feedforward neural networks as they do not have neurons in a single direction. In other words, they can pass information both in the direction of a previous layer and in the same layer. In this way information does not flow only in one direction and as a result this enables us to have short term memory. The RNN can be thought of as a Multi-Layer Perceptron (MLP) network with the addition of loops in its construction. The term MLP essentially refers to simple feedforward neural networks with at least three layers. In RNNs there is an input layer, a hidden layer and an output layer. This structure is similar to the MLP architecture except that the hidden layers are interconnected. In a vanilla (basic) RNN/LSTM model the nodes are connected in one direction. This type of architecture ensures that the output t=n depends on the inputs t=n, t=n-1, ..., and t=1. In other words, the output will depend on a sequence of data rather than on a data as shown in the following figure.

(Input1) → Output1

(Input2, Input1) → Output2

(Input3, Input2, Input1) → Output3

(Input4, Input3, Input2, Input1) → Output4

The following diagram is indicative of the operation of a RNN, showing the correlation of inputs and outputs. The repetitive process allows information to pass from one step to the next.

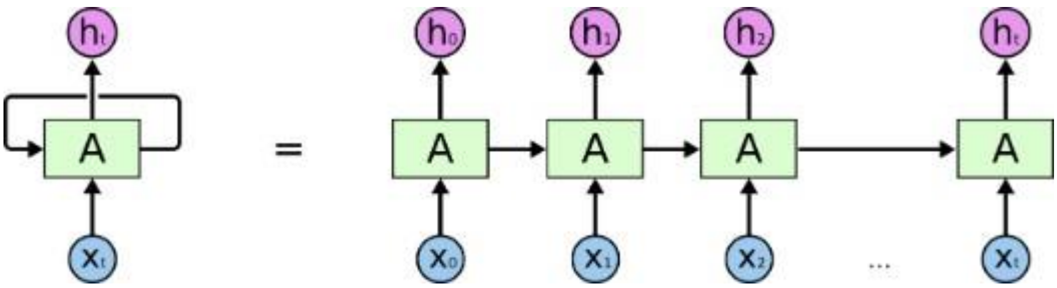

Conclusion

In this project, the field of Artificial Intelligence was thoroughly analysed and was divided into several subcategories: Machine Learning, Deep Learning and Neural Networks. General information about machine learning was stated, followed by the different types of ML and the future of ML was examined. In addition, deep learning was inspected and its future applications were mentioned. Finally, neural networks were analysed, their historical background was presented, and they were divided into two types: feedforward and recurrent neural networks. For future work, a deeper analysis of the different types of neural networks can be conducted, as well as the application of neural networks to solve real world problems.

Bibliography

1) "What Is Artificial Intelligence (AI)?". *Investopedia*, 2022, https://www.investopedia.com/terms/a/artificial-intelligence-ai.asp.
2) "Advantages And Disadvantages Of Artificial Intelligence". *Medium*, 2022, https://towardsdatascience.com/advantages-and-disadvantages-of-artificial-intelligence-182a5ef6588c.
3) Burns, Ed. "What Is Machine Learning And Why Is It Important?". SearchEnterpriseAI, https://www.techtarget.com/searchenterpriseai/definition/machine-learning-ML#:~:text=Machine%20learning%20(ML)%20is%20a,to%20predict%20new%20output%20values.
4) Choy, Garry et al. "Current Applications And Future Impact Of Machine Learning In Radiology". *Radiology*, 2018, https://doi.org/10.1148/radiol.2018171820.
5) DeBrusk, Chris. *The Risk Of Machine-Learning Bias (And How To Prevent It)*. Massachusetts Institute Of Technology, 2018, https://mitsmr.com/2Gefdre
6) Deng, Jia et al. "Imagenet: A Large-Scale Hierarchical Image Database." *Ieeexplore.ieee.org*. N.p., 2022. Web. 29 May 2022.
7) Nepal, Prabin. "Alexnet Architecture Explained." *Medium*. N.p., 2020. Web. 29 May 2022.
8) Ross, Philip E. "Alphago Wins Game One Against World Go Champion." *IEEE Spectrum*. N.p., 2016. Web. 29 May 2022.
9) Massachusetts Institute of Technology. *MIT Introduction To Deep Learning | 6.S191*. 2022. Web. 29 May 2022.
10) Wehle, Hans-Dieter. "Machine Learning, Deep Learning, And AI: What'S The Difference?." *Data Scientist Innovation Day*. Hans-Dieter Wehle, 2017. Web. 29 May 2022.
11) Team, Editorial. "Understanding The Universal Approximation Theorem." *Towards AI*. N.p., 2020. Web. 29 May 2022.
12) Mannam, Sai. "Applications Of Deep Learning In Healthcare." *Young Investigators* (2020): 1-3. Web. 29 May 2022.
13) Pierson, Harry A., and Michael S. Gashler. "Deep Learning In Robotics: A Review Of Recent Research." *Advanced Robotics* 31.16 (2017): 821-835. Web. 29 May 2022.
14) Ali, Javed. "Earthquake Prediction Using Deep Learning - Live AI." *Live AI*. N.p., 2021. Web. 29 May 2022.
15) Bouchard, Yarrow. "Tesla'S Deep Learning At Scale: Using Billions Of Miles To Train Neural Networks." *towardsdatascience.com*. N.p., 2019. Web. 29 May 2022.

16) Abiodun, Oludare Isaac, et al. “State-of-the-Art in Artificial Neural Network Applications: A Survey.” *Heliyon*, Elsevier, 23 Nov. 2018, www.ncbi.nlm.nih.gov/pmc/articles/PMC6260436/. Accessed 22 Apr. 2022.
17) By: IBM Cloud Education. “What Are Neural Networks?” *IBM*, www.ibm.com/cloud/learn/neural-networks#toc-history-of-rIfu5uF2. Accessed 22 Apr. 2022.
18) Chen, James. “Neural Network Definition.” *Investopedia*, Investopedia, 8 Feb. 2022, www.investopedia.com/terms/n/neuralnetwork.asp. Accessed 22 Apr. 2022.
19) DACS Staff, Lon R. Dean. “Artificial Neural Networks Technology.” *History of Neural Networks*, www2.psych.utoronto.ca/users/reingold/courses/ai/cache/neural4.html. Accessed 22 Apr. 2022.
20) “History: The 1940's to the 1970s.” *Neural Networks - History*, www.cs.stanford.edu/people/eroberts/courses/soco/projects/neural-networks/History/history1.html. Accessed 22 Apr. 2022.
21) Jaspreet. “A Concise History of Neural Networks.” *Medium*, Towards Data Science, 8 Apr. 2019, www.towardsdatascience.com/a-concise-history-of-neural-networks-2070655d3fec. Accessed 22 Apr. 2022.
22) Larry Hardesty | MIT News Office. “Explained: Neural Networks.” *MIT News | Massachusetts Institute of Technology*, www.news.mit.edu/2017/explained-neural-networks-deep-learning-0414. Accessed 22 Apr. 2022.
23) Opperman, Artem. “Underfitting and Overfitting in Deep Learning.” *Medium*, MLearning.ai, 19 July 2021, www.medium.com/mlearning-ai/underfitting-and-overfitting-in-deep-learning-687b1b7eb738. Accessed 22 Apr. 2022.
24) “Overfitting vs. Underfitting: What's the Difference?” *V7*, www.v7labs.com/blog/overfitting-vs-underfitting. Accessed 22 Apr. 2022.
25) Patel, Khush. “Overfitting vs Underfitting.” *Medium*, Towards Data Science, 14 Sept. 2019, www.towardsdatascience.com/overfitting-vs-underfitting-ddc80c2fc00d. Accessed 22 Apr. 2022.
26) Schmidhuber, Jürgen. “Deep Learning in Neural Networks: An Overview.” *Neural Networks*, Pergamon, 13 Oct. 2014, www.sciencedirect.com/science/article/pii/S0893608014002135. Accessed 22 Apr. 2022.