DEEP DEPTHS OF DEEP LEARNING

Shaina Mohite[1], Abdul Kalam[2], Mahima Yadav[3]

[1] Scholar, M.Des (UX), School of Architecture, Planning and Design, DIT University, Dehradun
[2,3] Assistant Professor School of Architecture, Planning and Design, DIT University, Dehradun

ABSTRACT

Deep learning, a subset of artificial intelligence (AI), utilizes artificial neural networks to extract knowledge from massive datasets. This technology fuels many current applications, from virtual assistants to fraud detection, and holds immense promise for the future of AI. This thesis explores the origins of deep learning, tracing its connection to advancements in artificial neural networks. It delves into the operational mechanisms of deep learning, explaining how these algorithms learn and process information. Finally, the thesis examines the future implications of deep learning for AI, considering its potential to achieve human-level capabilities. By comprehensively examining these aspects, this thesis aims to provide a thorough understanding of deep learning and its transformative role in shaping the future of AI.

Keyword: - *Deep Learning, Artificial Intelligence, Machine Learning, Neural Networks and perceptron.*

1. Introduction

Deep learning technology lies behind everyday products and services (such as digital assistants, voice-enabled TV remotes, and credit card fraud detection) as well as emerging technologies (such as self-driving cars).
Deep learning is all the rage today, as companies across industries seek to use advanced computational techniques to find useful information hidden across huge swaths of data. While the field of artificial intelligence is decades old, breakthroughs in the field of artificial neural networks are driving the explosion of deep learning.
Attempts at creating artificial intelligence go back decades. In the wake of World War II, the English mathematician and codebreaker Alan Turing penned his definition for true artificial intelligence. Dubbed the Turing Test, a conversational machine would have to convince a human that he was talking to another human. It took 60 years, but a computer finally passed the Turing Test back in 2014, when a chat bot developed by the University of Reading dubbed “Eugene” convinced 33% of the judges convened by the Royal Society in London that he was real. Since then, the field of deep learning and AI has exploded as computers get closer to delivering human-level capabilities.

1.1 The Concept and Working of Deep Learning

Deep Learning, more or less, mimics a human brain that uses neural networks to learn patterns in data. Neural networks are composed of layers of interconnected neurons for accuracy. These neural networks attempt to simulate the behavior of the human brain—albeit far from matching its ability—allowing it to “learn” from large amounts of data. It is a part of Machine Learning which in turn is the part of Artificial Intelligence.
The working of Deep Learning in simplest terms, through a combination of data inputs, weights, and bias. These elements work together to accurately recognize, classify, and describe objects within the data.
Deep neural networks consist of multiple layers of interconnected nodes, each building upon the previous layer to refine and optimize the prediction or categorization. This progression of computations through the network is called forward propagation. The input and output layers of a deep neural network are called visible layers. The input layer is where the deep learning model ingests the data for processing, and the output layer is where the final prediction or classification is made.
Another process called backpropagation uses algorithms, like gradient descent, to calculate errors in predictions and then adjusts the weights and biases of the function by moving backwards through the layers in an effort to train the

model. Together, forward propagation and backpropagation allow a neural network to make predictions and correct for any errors accordingly. Over time, the algorithm becomes gradually more accurate.
The diagram below describes the simplest type of deep neural network in the simplest terms. However, deep learning algorithms are incredibly complex, and there are different types of neural networks to address specific problems or datasets.

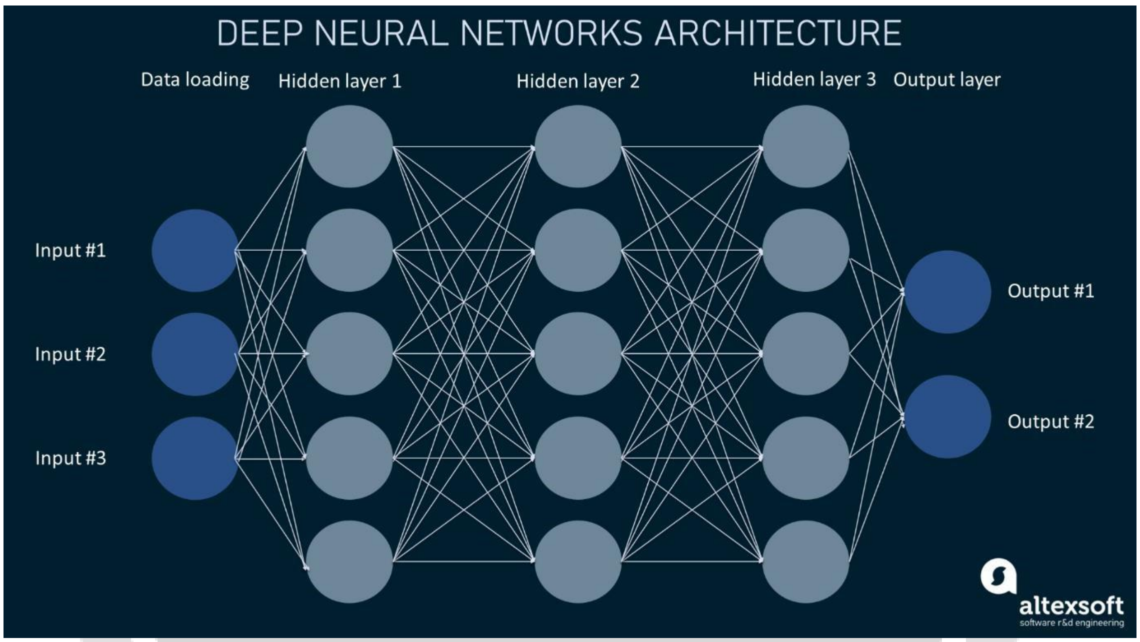

Fig - 1: Deep Neural Networks Architecture

Today, Deep learning networks can:

- Learn complex patterns in data, and this makes them very effective at tasks such as image recognition, natural language processing, and predictive modelling. Deep learning List Item - 2
- Drives many artificial intelligence (AI) applications and services that improve automation, performing analytical and physical tasks without human intervention.

Deep Learning was not feasible earlier because to achieve an acceptable level of accuracy, deep learning programs require access to immense amounts of training data and processing power, neither of which were easily available to programmers until the era of big data and cloud computing.
Also, people were of the opinion that Neural Networks were not working because the non-convex optimization problems they wanted to solve could not be solved perfectly

2. Tracing back to very first Neural Network

Neural networks are a set of algorithms that have been developed imitate the human brain in the way we identify patterns. In neurology, researchers study the way we process information. We have outstanding abilities to process information quickly and extract patterns.
The very initial stage of Artificial Neural Network was built on people's one point of view of understanding human mental phenomena: connectionism. Relevant work had started in 19th century. To make it easier to understand, it starts out as a philosophy question, rather than a computer science problem or a biology problem.

Alexander Bain first mentioned the connectionism ideas in his book BAIN, A. Mind and Body the Theories of Their Relation by Alexander Bain. Henry S. King & Company, 1873.
Even earlier, he offered an interpretation of mental phenomena within an association framework. Then, in his book, he argues a brand new point of view stating memory is a model capable of “putting required stuff together” rather than “storing everything beforehand”.
He named this memory structure as “neural groupings”, as a predecessor of “neural network” in our world, and further explained that under his view, connected nerve fibers are channeled to different parts of the network under different simulations.

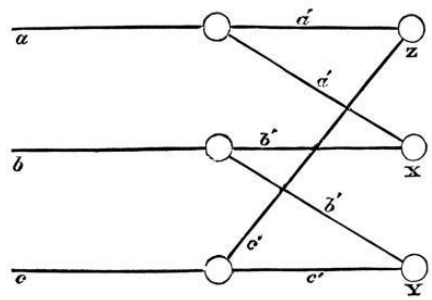

Fig -2: An example of Bain’s neural groupings structure.

Introduction As a very simple example, showed in Figure 2. X, Y, Z is triggered by different combinations of simulations from a, b or c. Bain’s work is recognized as the earliest work of discovering neural networks of human brain, however, Bain was not convinced by himself and argued that this structure shows no practical value.

Donald O. Hebb was influenced by Bain. In 1949, he stated the famous rule: “Cells that fire together, wire together”, which emphasized on the activation behavior of co-fired cells. This theory attempts to explain neural groupings theory and provides a biological basis for learning methods in the rehabilitation of memory. In a more formal way, it says that when an axon of cell A is near enough to excite cell B and repeatedly or persistently takes part in firing it, some growth process or metabolic change takes place in one or both cells such that as efficiency as one of the cells firing B is increased. The weight of neurons/connection units get updated with this learning rule, indicating that the connection of two signals is magnified every time when they appear together, which is consistent with our daily experience as human.

3. Perceptron – the most basic neuron

As these pioneers successfully brought the connectionism and neural groupings into the world, other researchers started to work to complete the whole theory. One important completion is Frank Rosenblatt’s thoughts on how the information is stored and transferred in neural network. He had raised three questions regarding perceptrons: how perceptrons can help to:

- Detect information,
- Store information and
- Recognize something with the information.

The first question is quite clear. The answer to it belongs to some work within the sensing field. As for the following two questions, inspired by the connectionism, Rosenblatt believes that simple logic and Boolean algebra cannot explain the phenomenon very well and he preferred probability theory. He made several assumptions to the model accounting for randomness, plasticity and similarity and he proposed a model of perceptron with S-points (as input), A-units (as activation function in model neural networks) and R-cells (as output).

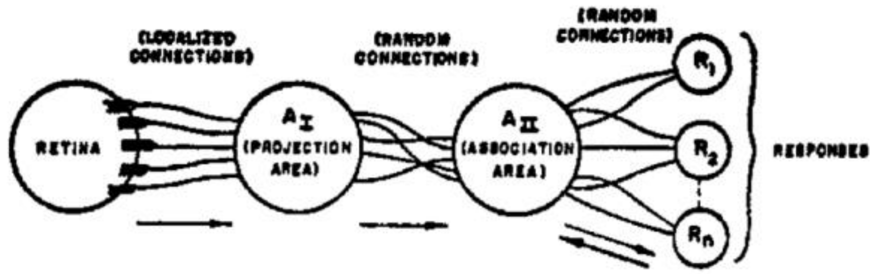

Fig -3: An illustration of Rosenblatt's model of perceptron

Accordingly, the perceptron, which is the basic building block of the neural network, was invented by Frank Rosenblatt in January 1957 at Cornell Aeronautical Laboratory, Inc at Buffalo, New York.

4. Minsky and Paper (1969)

The version of Perceptron we use nowadays was introduced by Minsky and Papert in 1969. They bring a major improvement to the classic model: they introduced an activation function. The activation function might take several forms and should "send" the weighted sum into a smaller set of possible values that allows us to classify the output. It's a smoother version than the thresholding applied before.

5. Adaline and Madaline (1959)

In 1959, Bernard Widrow and Marcian Hoff of Stanford developed models called "ADALINE" and "MADALINE."

5.1 Adaline (Adaptive Linear Neural):

A network with a single linear unit is called Adaline (Adaptive Linear Neural). A unit with a linear activation function is called a linear unit. In Adaline, there is only one output unit and output values are bipolar (+1,-1). Weights between the input unit and output unit are adjustable.

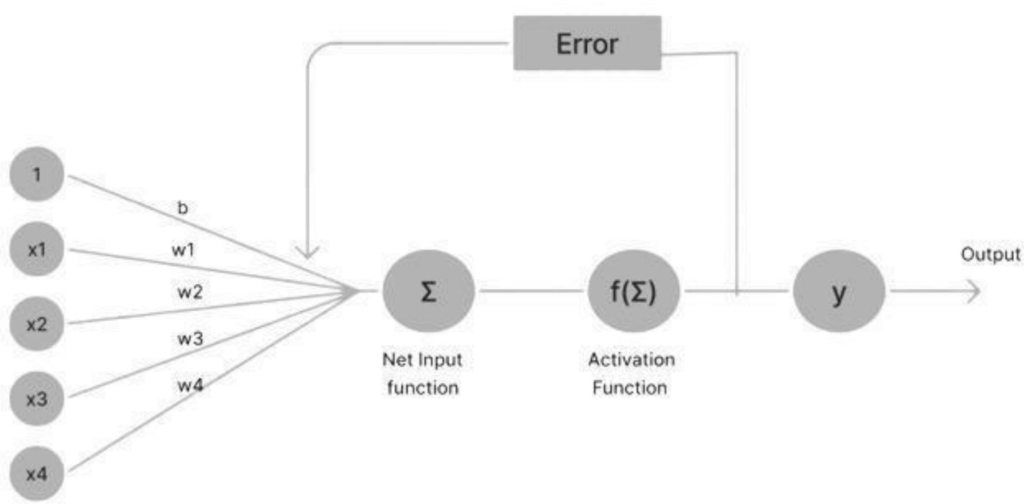

Fig 4- : Workflow of Adaline network

5.2 Madaline (Multiple Adaptive Linear Neuron):

The Madaline (supervised Learning) model consists of many Adaline in parallel with a single output unit. The Adaline layer is present between the input layer and the Madaline layer hence Adaline layer is a hidden layer. The weights between the input layer and the hidden layer are adjusted, and the weight between the hidden layer and the output layer is fixed.

Madaline was the first neural network applied to a real-world problem, using an adaptive filter that eliminates echoes on phone lines. While the system is as ancient as air traffic control systems, like air traffic control systems, it is still in commercial use.

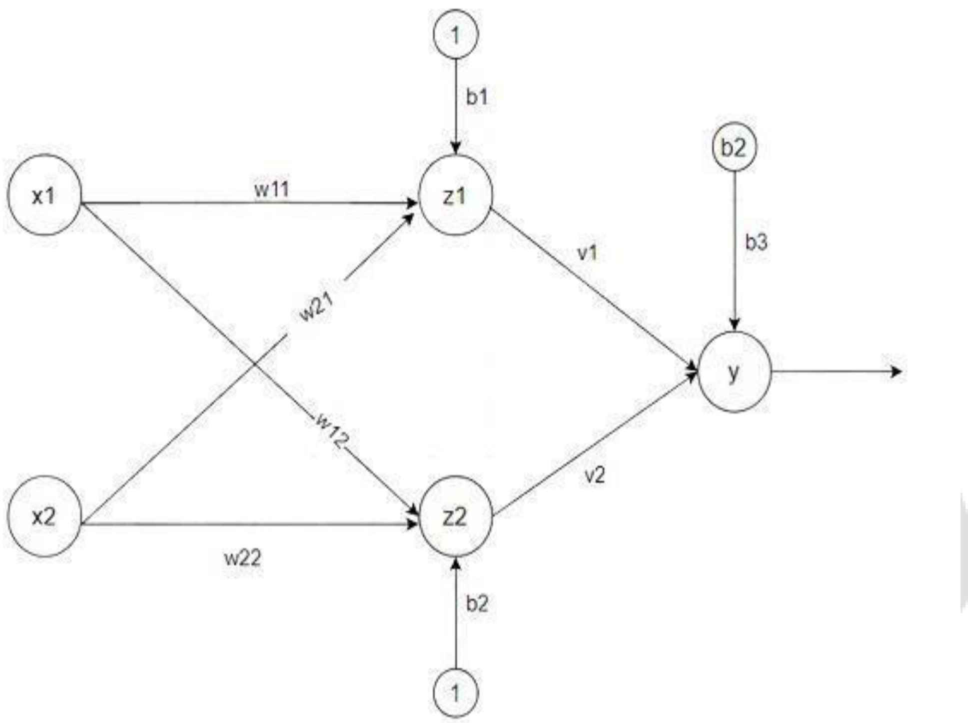

Fig 5- : Workflow of Madaline network

6. Minsky and Paper (1969)

The version of Perceptron we use nowadays was introduced by Minsky and Papert in 1969. They bring a major improvement to the classic model: they introduced an activation function. The activation function might take several forms and should “send” the weighted sum into a smaller set of possible values that allows us to classify the output. It’s a smoother version than the thresholding applied before. In 1969, ten years after the discovery of the perceptron—which showed that a machine could be taught to perform certain tasks using examples—Marvin Minsky and Seymour Paper published Perceptrons, their analysis of the computational capabilities of perceptrons for specific tasks.

7. Back Propagation (1974)

Backpropagation, developed in 1974 by is often viewed as a method for adapting artificial neural networks to classify patterns. Backpropagation is a process involved in training a neural network. It involves taking the error rate of a forward propagation and feeding this loss backward through the neural network layers to fine-tune the weights. Backpropagation is the essence of neural net training. Based on parts of the book by Rumelhart and colleagues, many authors equate backpropagation with the generalized delta rule applied to fully-connected feedforward networks.

8. Self Organizing Map (1980)

Self Organizing Maps or Kohenin's map is a type of artificial neural networks introduced by Teuvo Kohonen in the 1980s. SOM is trained using unsupervised learning, it is a little bit different from other artificial neural networks, SOM doesn't learn by backpropagation with SGD, it use competitive learning to adjust weights in neurons. And we use this type of artificial neural networks in dimension reduction to reduce our data by creating a spatially organized representation, also it helps us to discover the correlation between data.

9. Deep Belief Network (2006)

The Deep Belief Network (DBN) is a kind of Deep neural Network which is composed of stacked layers of Restricted Boltzman Machines (RBMs). It is a generative model and was proposed by Geoffrey Hinton in 2006. DBN can be used to solve unsupervised learning tasks to reduce the dimensionality of features, and can also be used to solve supervised learning tasks to build classification models or regression models.

Deep Belief Networks (DBNs) were invented as a solution for the problems encountered when using traditional neural networks training in deep layered networks, such as slow learning, becoming stuck in local minima due to poor parameter selection, and requiring a lot of training datasets.

10. GAN (2012)

By 2014, a generative adversarial network (GAN) was proposed by Goodfellow et al. as an intelligent deep-learning approach that could take the advantage of discriminative learners to build a well-behaved generative learner. This chapter dives into the details of the standard GAN model as the baseline member of the family of generative deep networks. By covering the principles of GANs, it looks at such early GANs and shows how to obtain satisfactory training. The chapter focuses on two well-known generative models, namely deep convolutional GAN and conditional GAN (CGAN). CGAN for simplicity, is a type of GAN that involves the conditional generation of data instances by a generator model. A conditional setup is used in CGANs, which means that both the generator and discriminator are contingent on auxiliary input from other modalities.

11. Capsule Networks (2017)

Symbolic Artificial Intelligence with its hard coding rules is incapable of solving these complex problems resulting in the introduction of Deep Learning (DL) models such as Recurrent Neural Networks and Convolutional Neural Networks (CNN). However, CNNs require lots of training data and are incapable of recognizing pose and deformation of objects leading to the introduction of Capsule Networks. Capsule Networks are the new sensation in Deep Learning. They have lived to this expectation as their performance in relation to the above problems has been better than Convolutional Neural Networks. Even with this promise in performance, lack of architectural knowledge and inner workings of Capsules serves as a hindrance for researchers to take full advantage of this breakthrough.

For a detailed timeline of Deep Learning throughout the history, we can refer the following timeline chart:

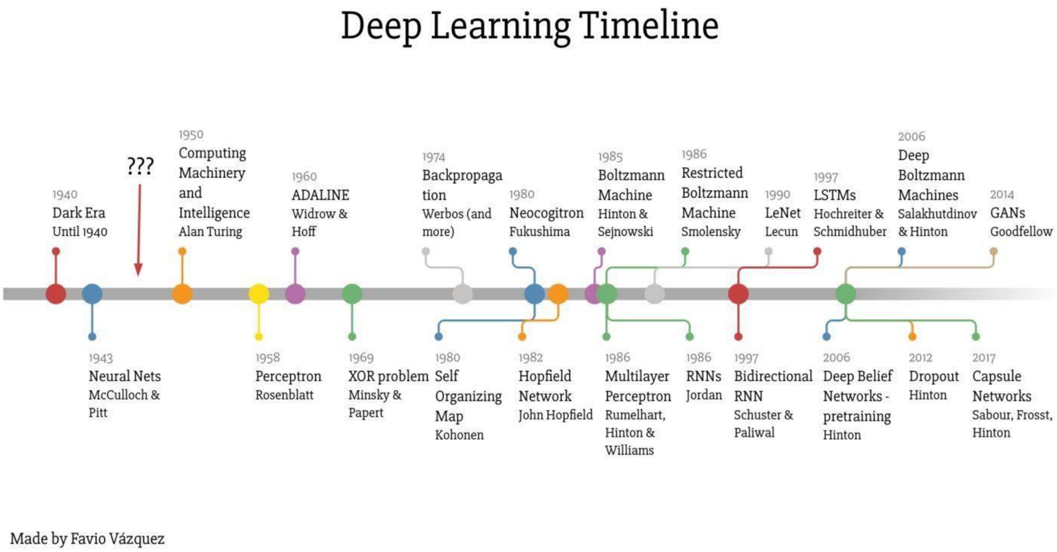

12. Future Implications of Deep Learning:

12.1 Forging global learning loops with and for AI [Article published by McKinsey and Company

For employees to keep raising the performance bar as required to maintain growth, there must be a mechanism for capturing the experiences, experiments, and learning occurring across the organization. At many businesses, learning typically gets stuck in the mind of one individual, team, business unit, or silo, rather than contributing to the organization as a whole. In contrast, AI-enabled companies develop the skills, processes, and technical systems to build global learning loops that turn individual knowledge and local insights into an ever-increasing flow of collective wisdom that everyone in the organization shares and contributes to. These learning systems codify valuable knowledge gained from the frontline business systems (operated with insights derived from AI at scale and speed) and the AI teams' approaches to processing data and developing AI models for solving business problems.

One of the best ways to create this global learning loop on the business side is through the development of an AI-driven nerve center for managing operations.

The global pharmaceutical company, for instance, developed what it calls its "clinical control tower" that continually updates and shares findings derived from the diverse data gathered from hundreds of clinical trials across thousands of sites around the world. This system enables decision makers to understand in detail what drives variations among clinical trials (in speed, quality, and cost) and delivers predictions that enable interventions to reallocate resources and avoid delays and waste.

12.2 Foundation models for Computer Vision [Interview with Andrew Ng- a British American computer scientist

We have now reached a point where we have so-called foundation models for NLP: humongous models like GPT3, trained on tons of data, that people can use to fine-tune for specific applications or domains. However, those NLP foundation models don't really utilize domain knowledge.

What about foundation models for computer vision? Are they possible, and if yes, how and when can we get there, and what would that enable? Foundation models are a matter of both scale and convention, according to Ng. He thinks they will happen, as there are multiple research groups working on building foundation models for computer vision.

"It's not that one day it's not a foundation model, but the next day it is," he explained. "In the case of NLP, we saw development of models, starting from the BERT model at Google, the transformer model, GPT2 and GPT3. It was a sequence of increasingly large models trained on more and more data that then led people to call some of these emerging models, foundation models."

Ng said he believes we will see something similar in computer vision. “Many people have been pre-training on ImageNet for many years now,” he said. “I think the gradual trend will be to pre-train on larger and larger data sets, increasingly on unlabeled datasets rather than just labelled datasets, and increasingly a little bit more on video rather than just images.”

As a computer vision insider, Ng is very much aware of the steady progress being made in AI. He believes that at some point, the press and public will declare a computer vision model to be a foundation model. Predicting exactly when that will happen, however, is a different story. How will we get there? Well, it’s complicated.

12.3 Possible Risks

Advances in deep learning algorithms overshadow their security risk in software implementations. This paper discloses a set of vulnerabilities in popular deep learning frameworks including Caffe, TensorFlow, and Torch. Contrary to the small code size of deep learning models, these deep learning frameworks are complex, and they heavily depend on numerous open-source packages. Recent studies show vulnerabilities by studying their impact on common deep learning applications such as voice recognition and image classification. By exploiting these framework implementations, attackers can launch denial-of-service attacks that crash or hang a deep learning application, or control-flow hijacking attacks that lead to either system compromise or recognition evasions.

13. Conclusion

In conclusion, deep learning has emerged as a powerful engine driving innovation in artificial intelligence. Its ability to learn from vast amounts of data has revolutionized numerous industries and holds the potential to tackle even more complex challenges in the future. As deep learning algorithms continue to evolve and computational power increases, the boundaries between human and machine intelligence will likely continue to blur.

14. References

1. Sandra Mahoney and Jon Schamber. (2011). “Integrative and Deep Learning Through a Learning Community: A Process View of Self.” The Journal of General Education, 60(4), 234–2471(https://doi.org/10.5325/jgeneeduc.60.4.0234)

2. LeCun, Y., Bengio, Y., & Hinton, G. E. (2015). “Deep learning.” Nature, 521(7553), 436–4442(https://doi.org/10.1038/nature14539)

3. Bostan, A., Ekin, C. C., Şengül, G., Karakaya, M., & Tirkeş, G. (2018, September 1). “Deep Learning and Current Trends in Machine Learning.” 2018 3rd International Conference on Computer Science and Engineering (UBMK)3(https://doi.org/10.1109/ubmk.2018.8566365)

4. Wang, H., & Raj, B. (2015). “A survey: Time travel in deep learning space: An introduction to deep learning models and how deep learning models evolved from the initial ideas.” 4(https://doi.org/10.48550/arXiv.1510.04781)

5. Mohamed Abdel-Basset, Nour Moustafa, Hossam Hawash. “Generative Adversarial Networks (GANs)” in Deep Learning Approaches for Security Threats in IoT Environments, IEEE, 2023, pp. 271-285, doi: 10.1002/9781119884170.ch12.

6. Q. Xiao, K. Li, D. Zhang, and W. Xu. “Security Risks in Deep Learning Implementations.” 2018 IEEE Security and Privacy Workshops (SPW), San Francisco, CA, USA, 2018, pp. 123-128, doi: 5(https://doi.org/10.1109/SPW.2018.00027)

7. Neural Networks - history. (n.d.).https://cs.stanford.edu/people/eroberts/courses/soco/projects/neural networks/History/history1.html

8. Waldrop, M. M. (2019). What are the limits of deep learning? Proceedings of the National Academy of Sciences of the United States of America, 116(4), 1074–1077. https://doi.org/10.1073/pnas.1821594116

9. Nielsen, M. A. (2015). Neural networks and deep learning. http://neuralnetworksanddeeplearning.com/

10. GfG. (2023, April 20). Adaline and Madaline Network. GeeksforGeeks. https://www.geeksforgeeks.org/adaline-and-madaline-network/

11. Minsky, M., & Papert, S. (2017). Perceptrons. In The MIT Press eBooks. https://doi.org/10.7551/mitpress/11301.001.0001

12. Meakin, T., Palmer, J., Sartori, V., & Vickers, J. (2021, April 30). Winning with AI is a state of mind. McKinsey & Company. https://www.mckinsey.com/capabilities/quantumblack/our-insights/winning-with-ai-is-a-state-of-mind

13. Anadiotis, G. (2022, March 22). Andrew Ng predicts the next 10 years in AI. VentureBeat. https://venturebeat.com/ai/andrew-ng-predicts-the-next-10-years-in-ai/