Latest Trend - Deep Learning

P.V.K PRASANNA CHAITANYA

KESHAV MEMORIAL INSTITUTE OF TECHNOLOGY

ALEKYA SATTU

G.NARAYANAMMA INSTITUTE OF TECHNOLOGY AND SCIENCES

SAI KIRAN CHIKKALA

MAHATMA GANDHI INSTITUTE OF TECHNOLOGY

ABSTRACT

In this paper,we present about DEEP LEARNING .The rise of artificial intelligence is grounded in the success of deep learning. Neural networks are a broad family of algorithms that have formed the basis for deep learning. Early work on neural networks actually began in the 1950s and 60s. And just recently, has experienced a resurgence of interest, as deep learning has achieved impressive state-of-the-art results..Deep learning is becoming a mainstream technology in different applications .We have gathered the information from the surveys and from various organizations on how they have used the concept of Deep Learning in different aspects. Artificial Intelligence can be considered the all-encompassing umbrella. It refers to computer programs being able to “think,” behave, and do things as a human being might do them.Machine learning goes beyond that. It involves providing machines with the data they need to “learn” how to do something without being explicitly programmed to do it.Deep Learning is an artificial intelligence function that imitates the workings of the human brain in processing data and creating patterns for use in decision making. Deep learning is a subset of machine learning in Artificial Intelligence (AI) that has networks capable of learning unsupervised from data that is unstructured or unlabeled.Also known as Deep Neural Learning or Deep Neural Network.Deep learning allows computational models that are composed of multiple processing layers to learn representations of data with multiple levels of abstraction. These methods have dramatically improved the state-of-the-art in speech recognition, visual object recognition, object detection and many other domains such as drug discovery and genomics. We shall also highlight upon the fundamentals of deep learning ,types of neural networks and it’s practical applications in the present scenario.

KEYWORDS- *Deep Learning; Neural Networks; Algorithms; Stacking; Evolved; Computation; Prediction; Recognition; Unsupervised; Neurons; Mapping; Topology; Conventional; Connected Networks; Convolutional; Self-Driving; Dependencies; Transparency; Accuracy; Inferring; Training; Classifying; Potentially.*

INTRODUCTION

Deep learning is a specific subset of Machine Learning, which is a specific subset of Artificial Intelligence. In the 1980s, most neural networks were a single layer due to the cost of computation and availability of data. Nowadays we can afford to have more hidden layers in our Neural Nets, hence the name “Deep” Learning.Machine Learning has been used for classification on images and text for decades, but it struggled to cross the threshold – there’s a baseline accuracy that algorithms need to have to work in business settings. Deep Learning is finally enabling us to cross that line in places we weren’t able to before.Computer vision is a great example of a task that Deep Learning has transformed into something realistic for business applications. Using Deep Learning to classify and label images isn’t only better than any other traditional algorithms, it’s starting to be better than actual humans.Deep Learning is important because it finally makes these tasks accessible – it brings previously irrelevant workloads into the purview of Machine Learning.

WHAT IS DEEP LEARNING?

Deep Learning[1] is a subfield of machine learning concerned with algorithms inspired by the structure and function of the brain called artificial neural networks.It is part of a broader family of machine learning methods based on learning data representations, as opposed to task-specific algorithms. Learning can be supervised, semi-supervised or unsupervised.An algorithm is considered to be deep if the input data is passed through a series of nonlinearities or nonlinear transformations before it becomes output. In contrast, most modern machine learning algorithms are considered "shallow" because the input can only go only a few levels of subroutine calling.It removes the manual identification of features in data and, instead, relies on whatever training process it has in order to discover the useful patterns in the input examples. This makes training the neural network easier and faster, and it can yield a better result that advances the field of artificial intelligence.

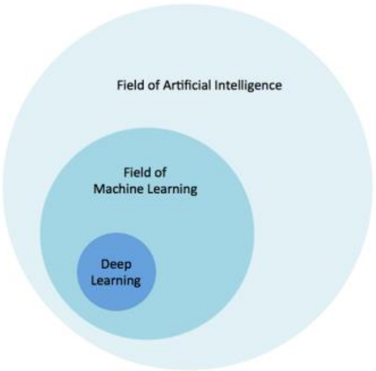

HISTORY OF DEEP LEARNING:

The history[2] of Deep Learning can be traced back to 1943, when Walter Pitts and Warren McCulloch created a computer model based on the neural networks of the human brain. They used a combination of algorithms and mathematics they called "threshold logic" to mimic the thought process. Since that time, Deep Learning has evolved steadily. The earliest efforts in developing Deep Learning algorithms came from Alexey Grigoryevich Ivakhnenko (developed the *Group Method of Data Handling*) and Valentin Grigor′evich Lapa (author of *Cybernetics and Forecasting Techniques)* in 1965. They used models with polynomial (complicated equations) activation functions, that were then analyzed statistically.

ARCHITECTURE OF DEEP LEARNING:

Deep learning isn't a single approach but rather a class of algorithms and topologies that you can apply to a broad spectrum of problems. While deep learning is certainly not new, it is experiencing explosive growth because of the intersection of deeply layered neural networks.

DEEP LEARNING AND THE RISE OF GPU:

Deep Learning consist of deep networks of varying topologies. Neural networks have been around for quite a while, but the development of numerous layers of networks (each providing some function, such as feature extraction) made them more practical to use.Adding layers means more interconnections and weights between and within the layers. This is where GPUs benefit deep learning, making it possible to train and execute these deep networks (where raw processors are not as efficient).

GPUs (Graphics Processing Unit) differ from traditional multicore processors in a few key ways. First, a traditional processor might contain 4 - 24 general-purpose CPUs, but a GPU might contain 1,000 - 4,000 specialized data processing cores.The high density of cores makes the GPU highly parallel (that is, it can perform many computations at once) compared with traditional CPUs.This makes GPUs ideal for large neural networks in which many neurons can be computed at once (where a traditional CPU could parallelize a considerably smaller number in parallel).

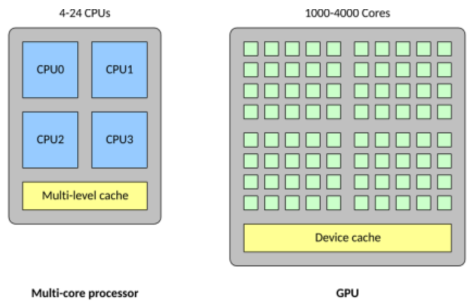

DIFFERENT ARCHITECTURES:

The number of architectures[3] and algorithms that are used in deep learning is wide and varied. This section explores five of the deep learning architectures spanning the past 20 years. Notably, LSTM and CNN are two of the oldest approaches in this list but also two of the most used in various applications.

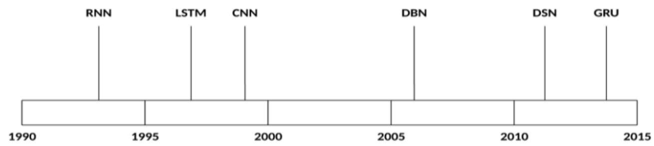

Fig 1: Evolution of different architectures

Architecture	Application
RNN	Speech Recognition, Handwriting Recognition
LSTM/GRU networks	Natural language text compression, handwriting recognition, speech recognition, gesture recognition, image captioning
CNN	Image recognition, video analysis, natural language processing
DBN	Image recognition, information retrieval, natural language understanding, failure prediction
DSN	Information retrieval, continuous speech recognition

Table 1: Different Architectures of Deep Learning

TYPES OF NEURAL NETWORKS AND THEIR ARCHITECTURE:

Artificial neural networks[4] are computational models which work similar to the functioning of a human nervous system. There are several kinds of artificial neural networks. These type of networks are implemented based on the mathematical operations and a set of parameters required to determine the output. Let's look at some of the neural networks:

RECURRENT NEURAL NETWORKS: The RNN[6] is one of the foundational network architectures from which other deep learning architectures are built. The primary difference between a typical multilayer network and a recurrent network is that rather than completely feed-forward connections, a recurrent network might have connections that feed back into prior layers (or into the same layer). This feedback allows RNNs to maintain memory of past inputs and model problems in time.Recurrent Neural Networks (RNN) are a powerful and robust type of neural networks and belong to the most promising algorithms out there at the moment because they are the only ones with an internal memory.RNN's are relatively old, like many other deep learning algorithms.In a RNN, the information cycles through a loop. When it makes a decision, it takes into consideration the current input and also what it has learned from the inputs it received previously.

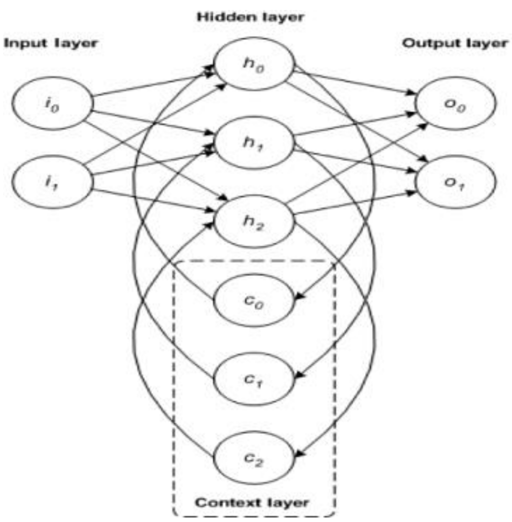

Fig 2: Recurrent Neural Networks

The two images below illustrate the difference in the information flow between a RNN and a Feed-Forward Neural Network.

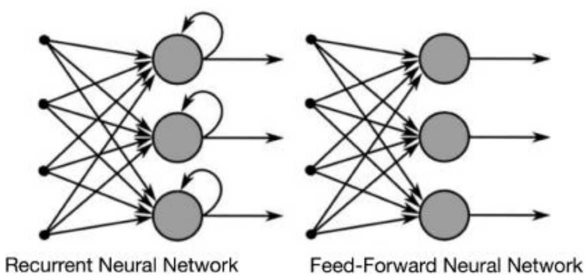

Fig 3:Difference between Recurrent and Feed-Forward Neural Networks

LSTM/GRU NETWORKS:

The LSTM was created in 1997 developed by Hochreiter and Schimdhuber. IBM applied LSTMs in IBM Watson® for milestone-setting conversational speech recognition.The LSTM departed from typical neuron-based neural network architectures and instead introduced the concept of a memory cell. The memory cell can retain its value for a short or long time as a function of its inputs, which allows the cell to remember what's important and not just its last computed value.

The LSTM memory cell contains three gates that control how information flows into or out of the cell. The input gate controls when new information can flow into the memory. The forget gate controls when an existing piece of information is forgotten, allowing the cell to remember new data. Finally, the output gate controls when the information that is contained in the cell is used in the output from the cell.

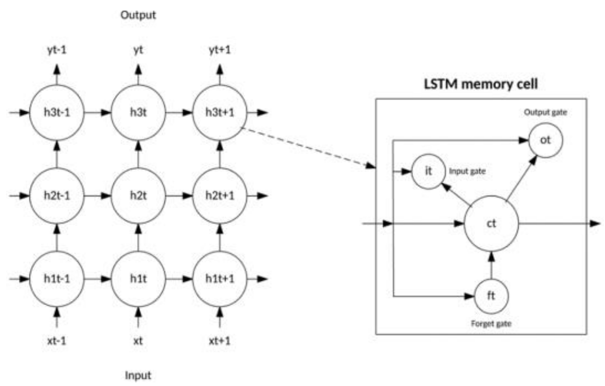

Fig 4: LSTM

In 2014, a simplification of the LSTM was introduced called the gated recurrent unit.The GRU includes two gates: an update gate and a reset gate. The update gate indicates how much of the previous cell contents to maintain. The reset gate defines how to incorporate the new input with the previous cell contents. A GRU can model a standard RNN simply by setting the reset gate to 1 and the update gate to 0.

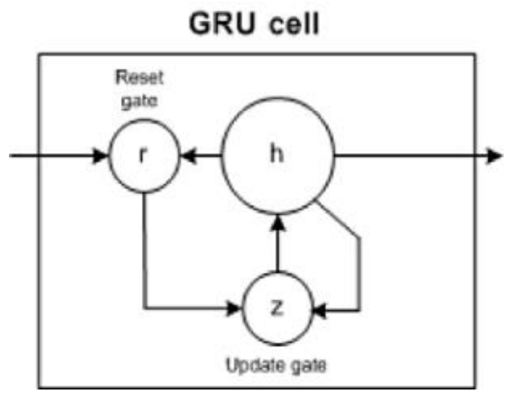

Fig 5: GRU Networks

CONVOLUTIONAL NEURAL NETWORKS:

A CNN[7] is a multilayer neural network that was biologically inspired by the animal visual cortex. The architecture is particularly useful in image-processing applications. The first CNN was created by Yann LeCun; at the time, the architecture focused on handwritten character recognition, such as postal code interpretation. As a deep network, early layers recognize features (such as edges), and later layers recombine these features into higher-level attributes of the input.The architecture of a CNN is designed to take advantage of the 2D structure of an input image (or other 2D input such as a speech signal). This is achieved with local connections and tied weights followed by some form of pooling which results in translation invariant features. Another benefit of CNNs is that they are easier to train and have many fewer parameters than fully connected networks with the same number of hidden units.

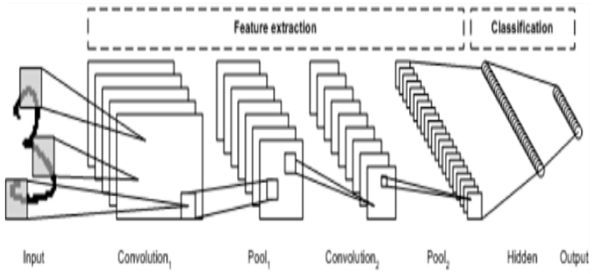

Fig 6: Convolutional Neural Networks

DEEP BELIEF NEURAL NETWORKS:

The DBN is a typical network architecture but includes a novel training algorithm. The DBN is a multilayer network (typically deep, including many hidden layers) in which each pair of connected layers is a restricted Boltzmann machine (RBM). In this way, a DBN is represented as a stack of RBMs.

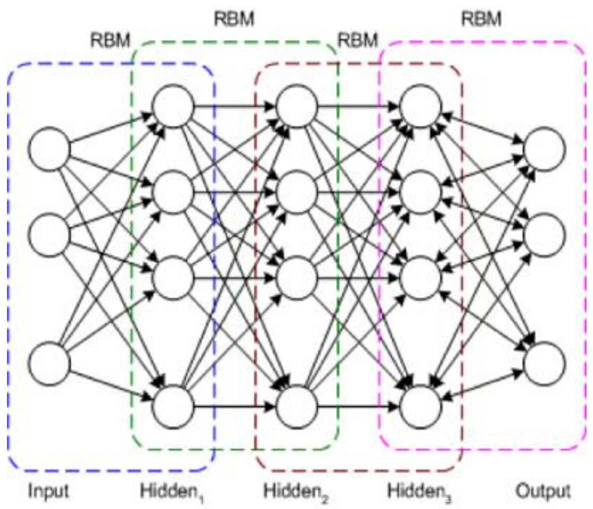

Fig 7: Deep Belief Neural Networks

DEEP STACKING NETWORKS:

The final architecture is the DSN, also called a deep convex network. A DSN is different from traditional deep learning frameworks in that although it consists of a deep network, it's actually a deep set of individual networks, each with its own hidden layers. This architecture is a response to one of the problems with deep learning: the complexity of training. Each layer in a deep learning architecture exponentially increases the complexity of training, so the DSN views training not as a single problem but as a set of individual training problems.

The DSN consists of a set of modules, each of which is a subnetwork in the overall hierarchy of the DSN. In one instance of this architecture, three modules are created for the DSN. Each module consists of an input layer, a single hidden layer, and an output layer. Modules are stacked one on top of another, where the inputs of a module consist of the prior layer outputs and the original input vector. This layering allows the overall network to learn more complex classification than would be possible given a single module.

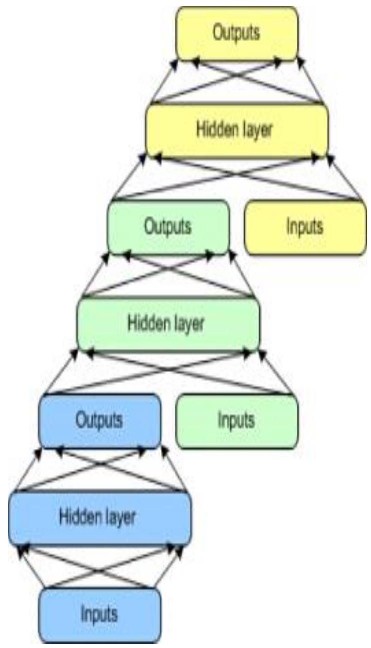

Fig 8: Deep Stacking Neural Networks

FEED FORWARD NEURAL NETWORK– Artificial Neuron:

This neural network is one of the simplest form of ANN, where the data or the input travels in one direction. The data passes through the input nodes and exit on the output nodes. This neural network may or may not have the hidden layers. In simple words, it has a front propagated wave and no back propagation by using a classifying activation function usually.

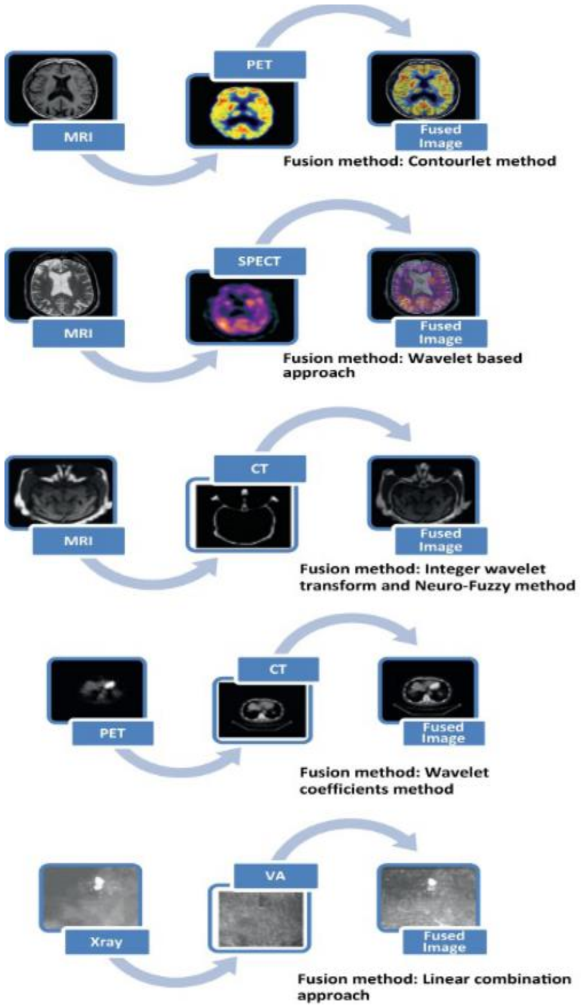

Fig 9: Feed-Forward Neural Networks

RADIAL BASIS FUNCTIONS NEURAL NETWORKS:

The radial basis function (RBF)[5] networks are inspired by biological neural systems, in which neurons are organized hierarchically in various pathways for signal processing, and they tuned to respond selectively to different features/characteristics of the stimuli within their respective fields. In general, neurons in higher layers have larger receptive fields and they selectively respond to more global and complex patterns. For example, neurons at different levels along the visual pathway respond selectively to different types of visual stimuli:

Neurons in the primary visual cortex (V1) receive visual input from the retina and selectively respond to different orientations of linear features;

Neurons in the middle temporal (MT) area receive visual input from the V1 area and selectively respond to different motion directions;

Neurons in the medial superior temporal area (MST) receive visual input from the MT area and selectively respond to different motion patterns (optic flow) such as rotation, expansion, contraction, and spiral motions.

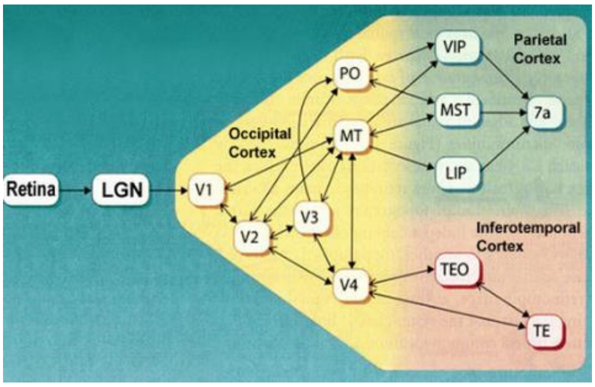

Fig 10: Radial Basis Function Neural Networks

SELF ORGANIZING MAP (SOM):

The Self-Organizing Map is one of the most popular neural network models. It belongs to the category of competitive learning networks. The Self-Organizing Map is based on unsupervised learning, which means that no human intervention is needed during the learning and that little needs to be known about the characteristics of the input data. We could, for example, use the SOM for clustering data without knowing the class memberships of the input data. The SOM can be used to detect features inherent to the problem and thus has also been called SOFM, the Self-Organizing Feature Map. The SOM algorithm is based on unsupervised, competitive learning. It provides a topology preserving mapping from the high dimensional space to map units. Map units, or neurons, usually form a two-dimensional lattice and thus the mapping is a mapping from high dimensional space onto a plane.

HOW DEEP LEARNING WORKS?

Deep Learning is a machine learning method. It allows us to train an AI to predict outputs, given a set of inputs. Both supervised and unsupervised learning can be used to train the AI.A deep machine learning process consists of two main phases: training and inferring. You should think about the training phase as a process of labeling large amounts of data and determining their matching characteristics. The system compares these characteristics and memorizes them to make correct conclusions when it faces similar data next time.A deep learning training process includes following stages:

- ANNs ask a set of binary false/true questions or.
- Extracting numerical values from data blocks.
- Classifying data according to the answers received.
- Labeling Data.

INFERRING PHASE

During the inferring phase, the deep learning AI makes conclusions and label new unexposed data using their previous knowledge.

Example:

Airplane ticket price estimation service.

We want our airplane ticket price estimator to predict the price using the following inputs (we are excluding return tickets for simplicity):

Origin Airport

Destination Airport

Departure Date

Airline

TRAINING THE NEURAL NETWORK

Training the AI is the hardest part of Deep Learning because,

We need a large data set.

We need a large amount of computational power.

For our airplane ticket price estimator, we need to find historical data of ticket prices. And due to the large amount of possible airports and departure date combinations, we need a very large list of ticket prices.To train the AI, we need to give it the inputs from our data set, and compare its outputs with the outputs from the data set. Since the AI is still untrained, its outputs will be wrong.Once we go through the whole data set, we can create a function that shows us how wrong the AI's outputs were from the real outputs. This function is called theCost Function Ideally, we want our cost function to be zero. That's when our AI's outputs are the same as the data set outputs.

HOW DEEP LEARNING IS DIFFERENT FROM MACHINE LEARNING?

Machine learning and deep learning on a rage! All of a sudden everyone is talking about them – irrespective of whether they understand the differences or not!

What is Machine Learning and Deep Learning?

Machine Learning: Machine learning is a subset of artificial intelligence in the field of computer science that often uses statistical techniques to give computers the ability to "learn" with data, without being explicitly programmed

Deep Learning: Deep learning is a machine learning technique that teaches computers to do what comes naturally to humans.

OVERVIEW OF DIFFERENCES:

DATA DEPENDENCIES: The most important difference[8] between deep learning and traditional machine learning is its performance as the scale of data increases. When the data is small, deep learning algorithms don't perform that well. This is because deep learning algorithms need a large amount of data to understand it perfectly. On the other hand, traditional machine learning algorithms with their handcrafted rules prevail in this scenario. Below image summarizes this fact.

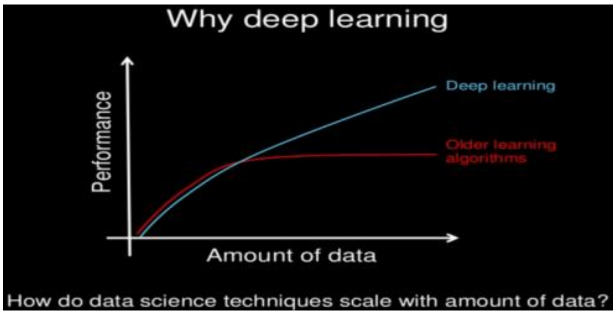

Fig 11:Scale of Data Increase

HARDWARE DEPENDENCIES:

Deep learning algorithms heavily depend on high-end machines, contrary to traditional machine learning algorithms, which can work on low-end machines. This is because the requirements of deep learning algorithm include GPUs which are an integral part of its working. Deep learning algorithms inherently do a large amount of matrix multiplication operations. These operations can be efficiently optimized using a GPU because GPU is built for this purpose.

FEATURE ENGINEERING

Feature engineering is a process of putting domain knowledge into the creation of feature extractors to reduce the complexity of the data and make patterns more visible to learning algorithms to work. This process is difficult and expensive in terms of time and expertise.In Machine learning, most of the applied features need to be identified by an expert and then hand-coded as per the domain and data type.Deep learning algorithms try to learn high-level features from data. This is a very distinctive part of Deep Learning and a major step ahead of traditional Machine Learning. Therefore, deep learning reduces the task of developing new feature extractor for every problem. Like, Convolutional NN will try to learn low-level features such as edges and lines in early layers then parts of faces of people and then high-level representation of a face.

PROBLEM SOLVING APPROACH:

When solving a problem using traditional machine learning algorithm, it is generally recommended to break the problem down into different parts, solve them individually and combine them to get the result. Deep learning in contrast advocates to solve the problem end-to-end.

EXECUTION TIME:

Usually, a deep learning algorithm takes a long time to train. This is because there are so many parameters in a deep learning algorithm that training them takes longer than usual. State of the art deep learning algorithm ResNet takes about two weeks to train completely from scratch. Whereas machine learning comparatively takes much less time to train, ranging from a few seconds to a few hours.

ADVANTAGES:

The deep learning term forms multiple questions among people who have never faced this technology in practice. What is it, what is its technical background, and what benefits can it bring to technological companies?As a part of artificial intelligence (AI), deep learning stands behind numerous innovations: self-driving cars, both voice and image recognition, etc. This technology has occupied multiple aspects of human lives.Some of its beneficial[9] areas are:

CREATING NEW FEATURES:

One of the main benefits of deep learning over various machine learning algorithms is its ability to generate new features from limited series of features located in the training dataset. Therefore, deep learning algorithms can create new tasks to solve current ones.Since deep learning can create features without a human intervention, data scientists can save much time on working with big data and relying on this technology. It allows them to use more complex sets of features in comparison with traditional machine learning software.

ADVANCED ANALYSIS:

Due to its improved data processing models, deep learning generates actionable results when solving data science tasks. While machine learning works only with labeled data, deep learning supports unsupervised learning techniques that allow the system become smarter on its own. The capacity to determine the most important features allows deep learning to efficiently provide data scientists with concise and reliable analysis results.

CHALLENGES:

Deep learning is an approach that models human abstract thinking (or at least represents an attempt to approach it) rather than using it. However, this technology has a set of significant disadvantages despite all its benefits.

CONTINUOUS DATA MANAGEMENT:

In deep learning, a training process is based on analyzing large amounts of data. Although, fast-moving and streaming input data provides little time for ensuring an efficient training process. That is why data scientists have to adapt their deep learning algorithms in the way neural networks can handle large amounts of continuous input data.

ENSURING CONCLUSION TRANSPARENCY:

Another important disadvantage of deep learning software is that it is incapable of providing arguments why it has reached a certain conclusion. Unlike in case of traditional machine learning, you cannot follow an algorithm to find out why your system has decided that it is a cat on a picture, not a dog. To correct errors in DL algorithms, you have to revise the whole algorithm.

RESOURCE DEMANDING TECHNOLOGY:

Deep learning is a quite resource-demanding technology. It requires more powerful GPUs, high-performance graphics processing units, large amounts of storage to train the models, etc. Furthermore, this technology needs more time to train in comparison with traditional machine learning.

Despite all its challenges, deep learning discovers new improved methods of unstructured big data analytics for those with the intention to use it. Indeed, businesses can gain significant benefits from using deep learning within their tasks of data processing.

APPLICATIONS:

It's predicted that many deep learning applications[10] will affect your life in the near future. Actually, they are already making an impact. Within the next five to 10 years, deep learning development tools, libraries, and languages will become standard components of every software development toolkit.

1. **Self-driving cars**: Companies building these types of driver-assistance services, as well as full-blown self-driving cars like Google's, need to teach a computer how to take over key parts (or all) of driving using digital sensor systems instead of a human's senses. To do that companies generally start out by training algorithms using a large amount of data.

2. **Deep Learning on Health Care:** Breast or Skin-Cancer diagnostics? Mobile and Monitoring Apps? Or prediction and personalized medicine on the basis of Biobank-data? AI is completely reshaping life sciences, medicine, and healthcare as an industry. Innovations in AI are advancing the future of precision medicine and population health management in unbelievable ways. Computer-aided detection, quantitative imaging, decision support tools and computer-aided diagnosis will play a big role in years to come.

3. **Voice Search & Voice Activated Assistants:** One of the most popular usage areas of deep learning is voice search & voice-activated intelligent assistants. With the big tech giants have already made significant investments in this area, voice-activated assistants can be found on nearly every smartphone. Apple's Siri is on the market since October 2011. Google Assistant, the voice-activated assistant for Android, was launched less than a year after Siri. The newest of the voice-activated intelligent assistants is Microsoft Cortana.

4. **Automatic Machine Translation:** This is a task where given words, phrase or sentence in one language, automatically translate it into another language.Automatic machine translation has been around for a long time, but deep learning is achieving top results in two specific areas: *Automatic Translation of Text & Automatic Translation of Images*

5. **Image Recognition:** Another popular area regarding deep learning is image recognition. It aims to recognize and identify people and objects in images as well as to understand the content and context. Image recognition is already being used in several sectors like gaming, social media, retail, tourism, etc.

CONCLUSION

Deep Learning is a vast subject. Many data scientists solely focus only on Neural network techniques.There are many algorithms other than backpropagation.Neural networks particularly work well on some particular class of problems like image recognition.The neural network algorithms are very calculation intensive. They require highly efficient computing machines. Large datasets take a significant amount of runtime on R. We need to try different types of options and packages.Currently, there is a lot of exciting research going on, around neural networks.In contrast to more conventional machine learning and feature engineering algorithms, Deep Learning has an advantage of potentially providing a solution to address the data analysis and learning problems found in massive volumes of input data. More specifically, it aids in automatically extracting complex data representations from large volumes of unsupervised data.Deep Learning models, improved formulation of data abstractions, distributed computing, semantic indexing, data tagging, information retrieval, criteria for extracting good data representations, and domain adaptation. Future works should focus on addressing one or more of these problems often seen in it's practical applications,Big Data, thus contributing to the Deep Learning Analytics research corpus.

REFERENCES

[1]https://machinelearningmastery.com/what-is-deep-learning/

[2]http://www.dataversity.net/brief-history-deep-learning/

[3]https://www.ibm.com/developerworks/library/cc-machine-learning-deep-learning-architectures/index.html

[4]https://analyticsindiamag.com/6-types-of-artificial-neural-networks-currently-being-used-in-todays-technology/

[5]http://fourier.eng.hmc.edu/e161/lectures/nn/node11.html

[6]https://towardsdatascience.com/recurrent-neural-networks-and-lstm-4b601dd822a5

[7]http://ufldl.stanford.edu/tutorial/supervised/ConvolutionalNeuralNetwork/

[8]https://www.analyticsvidhya.com/blog/2017/04/comparison-between-deep-learning-machine-learning/

[9]https://stfalcon.com/en/blog/post/deep-learning-benefits-and-challenges

[10]https://medium.com/@vratulmittal/top-15-deep-learning-applications-that-will-rule-the-world-in-2018-and-beyond-7c6130c43b01