IFERP

ISSN (Online) 2394-2320

International Journal of Engineering Research in Computer Science and Engineering (IJERCSE)

Vol 3, Issue 7, July 2016

Deep Learning AI: is Changing over Technology World; What, How, Why?

[1]Rashmi Mothkur, [2] Poornima K M

[1] G.M.Institute of Technology, Davangere, [2] J.N.N College of Engineering,Shimoga

[1]rashmimothkur@gmail.com,[2]poornimarudresh1712@gmail.com

Abstract: **Artifiicial Intelligence technology is increasingly prevalent in our everyday lives. It is spread beyond the academic world with major players like Google, Microsoft and Face book. This resurgence has been powered in no small part by a new trend in Artificial Intelligence, specifically in machine learning known as Deep Learning. Deep learning is an approach and an attitude to learning, where the learner uses higher-order cognitive skills such as the ability to analyse, synthesize, solve problems and thinks meta cognitively in order to construct long-term understanding. It involves the critical analysis of new ideas, linking them to already known concepts and principles so that this understanding can be used for problem solving in new, unfamiliar contexts. Deep learning entails a sustained, substantial and positive influence on the way students act, think or feel. Deep learning promotes understanding and application for life. Deep learners reflect on the personal significance of what they are learning. They are autonomous –they virtually teach themselves. But they are also collaborative learners, with high meta-cognitive and learning skills. Deep Learning is going to teach us all the lesson of our lives: Jobs Are for Machines. This paper gives an overview of need for Deep Learning, Architecture, frameworks, applications and future scope of deep learning.**

I. INTRODUCTION

Artificial intelligence(ai) seems to have become ubiquitous in the technology industry. Ais, are replying to our emails on Gmail, learning how to drive our cars, and sorting our holiday photos. Mark zuckerberg is even building one to help out around the house. The problem is that the concept of "artificial intelligence" is way too potent for its own good, conjuring images of supercomputers that operate spaceships, rather than particularly clever spam filters. Ai is going to doom humanity. There are three layers. Neural networks are at the bottom- they are a type of computer architecture onto which artificial intelligence is built. Machine learning is next-it is a program run on a neural network, training computers to look for certain answers in pots of data; and deep learning is on top-it is a particular type of machine learning that's only become popular over the past decade, largely thanks to two new resources: cheap processing power and abundant data back in the 1950s, when computers were new, the first generation of ai researchers eagerly predicted that fully fledged ai was right around the corner. But that optimism faded as researchers began to grasp the vast complexity of real-world knowledge - particularly when it came to perceptual problems such as what makes a face a human face, rather than a mask or a monkey face.

The first linear perceptron model was proposed by frank rosenblatt, a neuro-biologist at cornell university [3] in 1958. Despite early popularity, perceptron received criticism due to its drawbacks, the major one being unable to classify non-linearly separable classes[18]. Neural nets then regained popularity with the discovery of the back-propagation algorithm [1], which was first discovered by paul werbos in 1974 and it was rediscovered in 1986 by rumelhart, hinton and williams[17]. Amongst them, Geoffrey Hinton was a believer in distributed representation and neural networks. He also invented the Boltzmann machine[2] in 1985 along with terry sejnowski that are widely used these days for training deep networks. It dates back to the work by geoffrey hinton in mid-2000 and his paper "learning multiple layers of representation [16]". Deep learning became a buzzword from the cat-finder project building high-level features using large scale unsupervised learning by andrew ng and a team at google research.

Deep learning is often thought of as a set of algorithms that 'mimics the brain'. A more accurate description would be an algorithm that 'learns in layers'. Deep learning involves learning through layers which allows a computer to build a hierarchy of complex concepts out of simpler concepts.

The obscure world of deep learning algorithms came into public limelight when google researchers fed 10 million random, unlabeled images from youtube into their experimental deep learning system. They then instructed the system to recognize the basic elements of a picture and how these elements fit together. The system comprising 16,000 cpus was able to identify images that shared similar characteristics such as images of cats. This canonical experiment showed the potential of deep learning algorithms. Deep learning algorithms apply to many areas including computer vision, image recognition, pattern recognition, speech recognition, behavior recognition etc.

II. HOW DOES A COMPUTER LEARN?

The gold standard of artificial intelligence is a computer that can learn the same way we do as humans. To understand the significance of deep learning algorithms, it's important to understand how computers think and learn. Since the early days, researchers have attempted to create computers that think. Until recently, this effort has been rules based adopting a 'top down' approach. The top-down approach involved writing enough rules for all possible circumstances. But this approach is obviously limited by the number of rules and by its finite rules base. To overcome these limitations, a bottom-up approach was proposed. The idea here is to learn from experience. The experience was provided by 'labeled data'. Labeled data is fed to a system and the system is trained based on the responses. This approach works for applications like spam filtering. However, most data such as pictures, video feeds, sounds is not labeled. A computer can "learn" in two general ways:

2.1 traditional coding: a specific algorithm is programmed into a machine and it uses "if/then" commands. A machine is shown every type of apple, and then when it sees a picture of an apple; it can tell it's an apple. This can create machines that "seem" to think results in weak ai.

2.2 machine learning: a general algorithm is programmed that gives the machine a foundation for recognizing programs and programming itself. An algorithm is programmed that allows a machine to understand what an apple is by looking at examples. It doesn't have to see every type of apple to know an apple is an apple. This can create machines that can actually think results in strong ai.

III HOW DEEP LEARNING ALGORITHMS LEARN?

Deep learning algorithms are modeled on the workings of the brain. The brain may be thought of as a massively parallel analog computer which contains about 10^{10} simple processors (neurons) - each of which require a few milliseconds to respond to input [4]. To model the workings of the brain, in theory, each neuron could be designed as a small electronic device which has a transfer function similar to a biological neuron. Each neuron is connected to many others neurons to imitate the workings of the brain.

It uses neural networks -simple computer simulations of how biological neurons behave—to extract rules and patterns from sets of data. Showing a neural network enough pictures of cats, for instance, or have it listen to enough German speech, and machine will be able to tell if a picture or sound recording, it has never seen before is a cat, or in German. The general approach is not new; the perception, mentioned above, was one of the first neural networks. But the ever-increasing power of computers has allowed deep learning machines to simulate billions of neurons. At the same time, the huge quantity of information available on the internet has provided the algorithms with an unprecedented quantity of data to chew on. The results can be impressive. Facebook's deep face algorithm [5], for instance, is about as good as a human being when it comes to recognizing specific faces, even if they are poorly lit or seen from a strange angle. E-mail spam is much less of a problem than it used to be, because the vast quantities of it circulating online have allowed computers to realize what a spam e-mail looks like, and divert it before it ever reaches your inbox.
the deep learning systems promised to learn their own rules from scratch, and offered the pleasing symmetry of using brain-inspired mechanics to achieve brain-like function. The strategy called for simulated neurons to be organized into several layers as shown in figure 1. Give such a system, a picture; the first layer of learning will simply notice all the dark and light pixels. The next layer might realize that some of these pixels form edges; the next might distinguish between horizontal and vertical lines. Eventually, a layer might recognize eyes, and might realize that two eyes are usually present in a human face.

ISSN (Online) 2394-2320

International Journal of Engineering Research in Computer Science and Engineering (IJERCSE)
Vol 3, Issue 7, July 2016

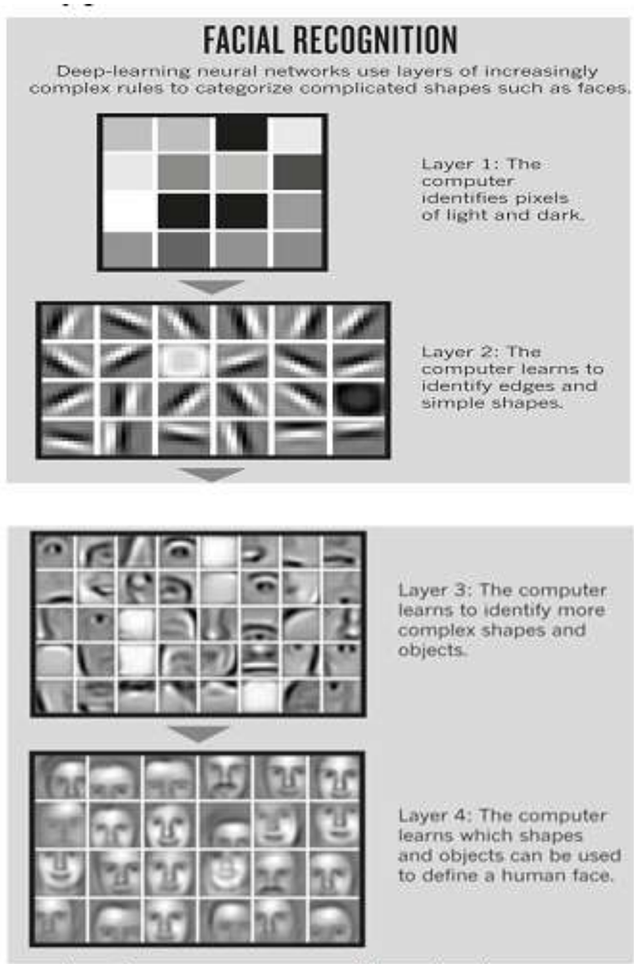

Fig 1: facial recognition layers

Deep learning systems are possible to implement now because of three reasons: high cpu power, better algorithms and the availability of more data. Over the next few years, the above factors will lead to more applications of deep learning systems.

IV ARCHITECTURE

Deep learning is a name for a certain set of stacked neural networks composed of several layers as shown in figure 2. The layers are made of nodes. A node is a place where computation happens, loosely patterned on the human neuron, and firing when it encounters sufficient stimuli. It combines input from the data with a set of coefficients or weights, that either amplify or dampen that input, thereby assigning significance to it in the task the algorithm is trying to learn. These input-weight products are summed and the sum is passed through a node's so-called activation function, to determine whether and to what extent that signal progresses further in the net to affect the ultimate outcome, say, an act of classification

A node layer is a row of neuron like switches that turn on or off as the input is fed through the net. Each layer's output is simultaneously the subsequent layer's input, starting from an initial input layer receiving your data.

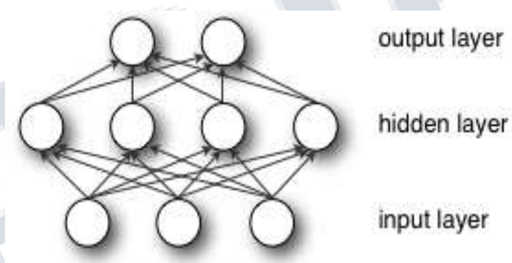

Fig 2: Architecture Of Deep Learning

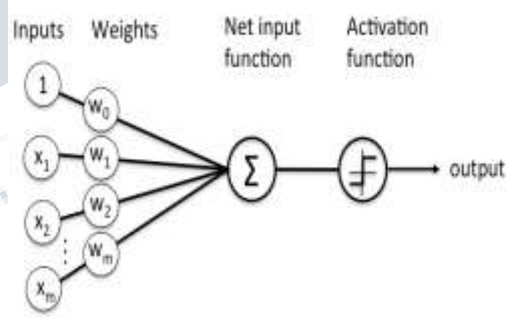

Fig 3: Function of Single Node

Each neuron in one layer is connected to every neuron in the next layer. As shown in figure 3, signals flow in only one direction. And finally, simplify the neuron design to 'fire' based on simple, weight driven inputs from other neurons. Such a simplified network called feed-forward neural network model is more practical to build and use.

a) Each neuron receives a signal from the neurons in the previous layer
B) Each of those signals is multiplied by a weight value.

ISSN (Online) 2394-2320

International Journal of Engineering Research in Computer Science and Engineering (IJERCSE)
Vol 3, Issue 7, July 2016

C) The weighted inputs are summed, and passed through a limiting function which scales the output to a fixed range of values.
D) The output of the limiter is then broadcast to all of the neurons in the next layer.

4.1 classes of deep architecture:

Deep learning refers to a rather wide class of machine learning techniques and architectures, with the hallmark of using many layers of non-linear information processing stages that are hierarchical in nature. Depending on how the architectures and techniques are intended for use, e.g., synthesis/generation or recognition/classification, one can broadly categorize most of the work in this area into three classes:

4.1.1 Generative deep architectures, which are intended to characterize the high-order correlation properties of the observed or visible data for pattern analysis or synthesis purposes, and/or characterize the joint statistical distributions of the visible data and their associated classes [6]. In the latter case, the use of bayes rule can turn this type of architecture into a discriminative one.

4.1.2 Discriminative deep architectures, which are intended to directly provide discriminative power for pattern classification, often by characterizing the posterior distributions of classes conditioned on the visible data.

4.1.3 Hybrid deep architectures, where the goal is discrimination but it is assisted with the outcomes of generative architectures via better optimization or/and regularization, or discriminative criteria [6] are used to learn the parameters in any of the deep generative models in category generative deep architectures above.

V. ADVANTAGES OF DEEP LEARNING

Compared to many traditional machine learning techniques, deep learning excels in four closely related areas:

5.1 Detecting Complex Nonlinear Patterns

The accurate car resale value cannot be determined solely by a feature like "engine volume". One must include more data like how old the car is, what brand it is, its upkeep and more. In addition one may choose include more volatile features such as what state the economy is in, trending colors, marketing campaigns, and so on. What is important to understand is that all these variables are interrelated and the effect of changing one depends on the values of the others? The effect of having a 5l engine depends on whether it's about a sports car or an old pickup truck. The age of a car relates to how collectible and well maintained it is. When all these features are taken into consideration a pretty complex picture is drawn.

5.2 learn the importance of variables automatically

Because deep learning algorithms fold their problem space so well, variables that do the same thing end up canceling each other out. When the deep auto encoder[7] algorithm was developed by hinton et al., the paper included a great plot showing the difference between a low level object function and deep learning. The data included roughly 800,000 newswire articles categorized into eight different classes.

5.3 learns highly variable patterns from few examples

Deep architectures allow for a feature to be learned without having to study an exponentially large number of configurations. If only *n* out of d variables are required to correctly classify something and there are *k* classes, a deep learning algorithm needs at most n*k observations in order to learn the classes. In contrast, traditional methods can require exponential nd observations! This huge difference has been researched extensively by bengio et al. Deep learning has the potential to correctly classify things which it has never seen before, and is not vulnerable to out-of-distribution examples like simpler models can be.

5.1.4 Learns Abstract Representations

The philosophical principle of compositionality lies at the very core of deep learning. It is probably impossible to describe the concept of "dog" using a single layer of binary pixel intensity values. Rather, one should use a stack of increasingly abstract representations of a dog. Andrew ng et al., discovered that these abstractions are not only important to solving complicated problems, but once learned, they can be transferred across problems.

VI. FRAMEWORKS.

6.1 Caffe: Convolution Architecture for Fast Feature Embedding: Caffe provides multimedia scientists and practitioners with a clean and modifiable framework for state-of-the-art deep learning algorithms and a collection of reference models. The framework is a BSD-licensed C++ library with Python and MATLAB bindings for training and deploying generalpurpose convolutional neural networks and other deep models efficiently on commodity architectures. Caffe fits industry and internet-scale media needs by CUDA GPU computation, processing over 40 million images a day on a single K40 or Titan GPU[8] (≈ 2.5 ms per image). By separating model representation from actual implementation, Caffe allows experimentation and seamless switching among platforms for ease of development and deployment from prototyping machines to cloud environments. Caffe is maintained and developed by the Berkeley Vision and Learning Center (BVLC) with the help of an active community of contributors on GitHub. It powers ongoing research projects, large-scale industrial applications, and startup prototypes in vision, speech, and multimedia.

6.2 OverFeat:OverFeat is a Convolutional Network-based image features extractor and classifier. OverFeat was trained on the ImageNet dataset and participated in the ImageNet 2013 competition. The package allows researchers to use OverFeat to recognize images and extract features.

6.3 Torch: Torch is a scientific computing framework with wide support for machine learning algorithms that puts GPUs first. It is easy to use and efficient, thanks to an easy and fast scripting language, LuaJIT- Just in Time compiler for Lua programming language[14]and an underlying C/CUDA implementation .Core features are a powerful N-dimensional array, mazing interface to C, via LuaJIT, linear algebra routines, neural network and energy-based models, numeric optimization routines, fast and efficient GPU support.

6.4 TensorFlow: TensorFlow is an open source software library for numerical computation using data flow graphs. Nodes in the graph represent mathematical operations, while the graph edges represent the multidimensional data arrays (tensors) communicated between them. The flexible architecture allows you to deploy computation to one or more CPUs or GPUs in a desktop, server, or mobile device with a single API. TensorFlow was originally developed by researchers and engineers working on the Google Brain Team within Google's Machine Intelligence research organization for the purposes of conducting machine learning and deep neural networks research, but the system is general enough to be applicable in a wide variety of other domains as well.

6.5 Theano:A math expression compiler that efficiently defines, optimizes, and evaluates mathematical expressions involving multi-dimensional arrayas.Theano is a numericalcomputation library for Python.InTheano, computations are expressed using a NumPy-like syntax and compiled to run efficiently on either CPU or GPUarchitectures.

6.6 MXNet: is a deep learning framework designed for both efficiency *and f*lexibility. It allows to mix the flavours of symbolic programming and imperative programming to maximize efficiency and productivity. In its core, a dynamic dependency scheduler that automatically parallelizes both symbolic and imperative operations on the fly. A graph optimization layer on top of that makes symbolic execution fast and memory efficient. The library is portable, lightweight and it scales to multiple GPUs and multiple machines.

6.7 CUDA: CUDA Deep Neural Network library (cuDNN) is a GPU-accelerated library of primitives for deep neural networks. cuDNN provides highly tuned implementations for standard routines such as forward and backward convolution, pooling, normalization and activation layers. cuDNN is part of the NVIDIA Deep Learning SDK. Deep learning researchers and framework developers worldwide rely on cuDNN for high-performance GPU acceleration. It allows them to focus on training neural networks and developing software applications rather than spending time on low-level GPU performance tuning.

6.8 Keras: is a minimalist, highly modular neural networks library, written in Python and capable of running on top of either TensorFlow or Theano. It was developed with a focus on enabling fast experimentation.

VII APPLICATIONS OF DEEP LEARNING:

Below are the outstanding applications that use deep learning techniques.

7.1 google deepmind: is a british artificial intelligence company founded in september 2010 as deepmind technologies. It was renamed when it was acquired by google in 2014. The company has created a neural network that learns how to play in a fashion similar to that of humans[9], as well as a neural turing machine, or a neural network that may be able to access an external memory like a conventional turing machine, resulting in a computer that mimics the short-term memory of the human brain[10].

7.2 Cortana: is an intelligent personal assistant created by Microsoft for Windows 10, Windows 10 Mobile,, Xbox One, iOS and Android. It has been launched as a key ingredient of Microsoft's planned "makeover" of the future operating systems for Windows PhoneWindows Phone 8 and Windows Cortana can set reminders, recognize natural voice without the requirement for keyboard input, and answer questions using information from the Bing search engine such as current weather and traffic, sports scores.If Firefox is the default browser, Cortana uses the Firefox default search engine instead of Bing. Searches will only be made with Microsoft Bing search engine and all links will open with Microsoft Edge.

7.3 Siri: is a computer program that works as an intelligent personal assistant and knowledge navigator, part of Apple Inc.'s iOS [11], watchOS, and tvOS operating systems. The feature uses a natural language user interface to answer questions, make recommendations and perform actions by delegating requests to a set of Web services.

7.4 DeepFace: is a deep learning facial recognition system created by a research group at Facebook. It identifies human faces in digital images. It employs a nine-layer net with over 120 million connection weights, and was trained on four million images uploaded by Facebook users[15] The system is said to be 97% accurate, compared to 85% for the Face Book's Next Generation Identification system[16] .

The potential applications where deep learning can be applied are discussed below:

7.5 Medicine:

Medical fields should be developed by deep learning. Tumors or cancers are detected on scanned images. This means nothing else but being able to utilize one of the strongest features of deep learning-the technique of image recognition. It is possible to dramatically increase precision using deep learning to help with the early detection of an illness and identifying the kind of illness. Since cnn (convolution neural network) can be applied to 3d images, 3d scanned images should be able to be analyzed relatively easily [12]. By adopting deep learning more in the current medical field, deep learning should greatly contribute.

7.6 Robotics:

Autonomous robotics recently gains attention in many areas, from the success of autonomous driving vehicles such as google car, tesla to flying droids, to industrial assistant robots such as baxter robots. The ultimate goal is that the robots can learn through the interactions with the environment and the human subjects automatically, and reduce the stress of human subjects in various industrial and home style tasks.

7.7 Automobiles:

Deep learning techniques can be applied in automobiles, and it is possible to reduce the risk of accidents by improving driving assistance functions. It can be said that the ultimate type of driving assistance is self-driving cars, which is being tackled mainly by google and tesla.

7.8 Agriculture:

7.8.1 Real-Time Plant Phenotyping: Deep Learning Techniques For Image Analysis And Object Classification, Farmers Are Now Able To Accurately Measure And Characterize Crops.

7.8.2 Advanced Plant Breeding: Deep Learning Is Powering The Next Generation Of Plant Breeding Techniques, That Can Predict Future Performance Of New Hybrids.

7.8.3 Predictive weather analytic: deep learning is already being used to make sense of a wide variety of sensor data including temperature, soil and humidity. these analysis combined with aerial images from satellites and drones build detailed weather models. with these solutions, farmers are able to make decisions that maximize yield based on future weather conditions.

7.9 Sports:

Deep learning can certainly contribute to sports as well. In the study field known as sports science, it has become increasingly important to analyze and examine data from sports.sport can be seen as a chunk of image sequences and number data. Since deep learning is good at identifying features that humans can't find, it will become easier to find out why certain players get good scores while others do not.

7.10 advert (ad) technologies:

Ad technologies could expand their coverage with deep learning. When we say ad technologies, this currently means recommendation or ad networks that optimize ad banners or products to be shown. on the other hand, when we say advertising, this doesn't only mean banners or ad networks. there are various kinds of ads in the world depending on the type of media such as tv ads, radio ads, newspaper ads, posters, flyers, and so on. we have also digital ad campaigns with youtube, vine, facebook, twitter, snapchat, and so on. advertising itself has changed its definition and content, but all ads have one thing in common, they consist of images and/or language. this means the deep learning is good at above discussed fields . until now, there was only use of user-behavior-based indicators, such as page view (pv),click through rate (ctr) and conversion rate (cvr)[13], to estimate the effect of an ad, but if we apply deep learning technologies, we might be able to analyze the actual content of an ad and auto generate ads going forward. especially since movies and videos can only be analyzed as a result of image recognition and natural language processing, video recognition, not image recognition, will gather momentum besides ad technologies.

7.11 Banking

A bank wants to analyze large volumes of video footage from the atm machines to detect abnormal activity. deep learning techniques can be used to extract images from video footage and identify normal and suspicious or fraudulent activities.

7.12 Insurance:

An insurance company wants to be able to compute claims costs directly from images. a deep learning technique can be trained to take all the accident images to classify the accident by severity of damage and the parts that were damaged.

7.13 earth observation and satellite images:

Earth observation has been a little overlooked this year but integrating deep learning technologies could power rapid developments across the entire field. When looking down from above, one longstanding challenge with managing satellite data concerns hyper spectral data classification.

VIII. CONCLUSION

Deep learning is at the root of some incredible new developments in image processing, natural language processing, biology, autonomous driving, artificial intelligence and more. Deep learning has revolutionized pattern recognition and machine learning. deep learning systems can automatically detect and generate more complex, high-level features from raw data sources. 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.

Future Scope

In terms of optimization, data quantities are also important. This means that deep neural networks require a significant amount of time for each training. but we are still quite far from emulating the learning abilities of animal of humans. a key element we are missing is predictive or unsupervised learning: the ability of a machine to model the environment, predict possible futures and understand how the world works by observing it and acting in it. there is a need to integrate reasoning with deep learning and need a good architecture for “episodic” (short-term) memory. there is a need to find good principles for unsupervised learning. better language understanding, dialog, and translation better auto-pilots for cars, medical image analysis, and robot perception are very active topics of research at the moment.

REFERENCES

[1] R.Rojas,“chapter7-neural networks: a systematic introduction*”*, *springer*, 1996.

[2]david h ackley, geofrey e hinton, “a learning algorithm for boltzmann machines”, *cognitive science*, vol:9, pp:147-169,1985.

[3] samir roy, udit chakraborthy, “introduction to soft

computing: neuro fuzzy and genetic algorithms",pearson, 2013.

[4] winfred philips, "chapter 2-computers and intelligence", *consortium on cognitive science instruction.*

[5]yaniv taigman, ming yang, marc'aurelio ranzato, "deepface: closing the gap to human-level performance in face verification", *ieee conference on computer vision and pattern recognition*, pp: 1701-1708, 2014.

[6] li deng, "three classes of deep learning architectures and their applications: a tutorial survey", *transactions on signal and information processing*, 2011.

[7]pascal vincent, hugo larochelle, "stacked denoising autoencoders: learning useful representations in a deep network with a local denoising criterion" *journal of machine learning research*, pp: 3371-3408, 2010.

[8]yang qing jia , evan shelhamer , jeff donahue, sergey karayev, "caffe: convolutional architecture for fast feature embedding", *22nd international conference on multimedia*,pp:675-678,acm,2014.

[9]"the last ai breakthrough deepmind made before google bought it",the physicsar arxivblog,retrieved 12 october 2014.

[10]best of 2014: google's secretive deepmind startup unveils a "neural turing machine", mit technology review.

[11]hay,timothy,"siri inc. launches 'do engine'application for iphone", retrieved 9 october 2011.

[12]yusuke sugumori, "practical applications of deep learning", packt publishers, may 2016.

[13]http://www.wordstream.com/blog/ws/2016/02/29/google- adwords-industry-benchmarks.

[14]http://luajit.org/luajit.html.

[15]tom simonite,"facebook creates software that matches faces almost as well as you do", massachusetts institute of technology review, march 17, 2014.

[16]geoffrey e hinton, "learning multiple layers of representation",trends in cognitive sciences, vol:11, no:10, pp:428-434, 2007.

[17]marvin minsky and seymore papert, "perceptrons: an introduction to computational geometry", mit press,1987.

[18]d.e.rumelhart, g.e.hinton, "learning internal representation by error Propagation",pp:318-362,1986.