2018

Neil Pundit and Anil Rewari

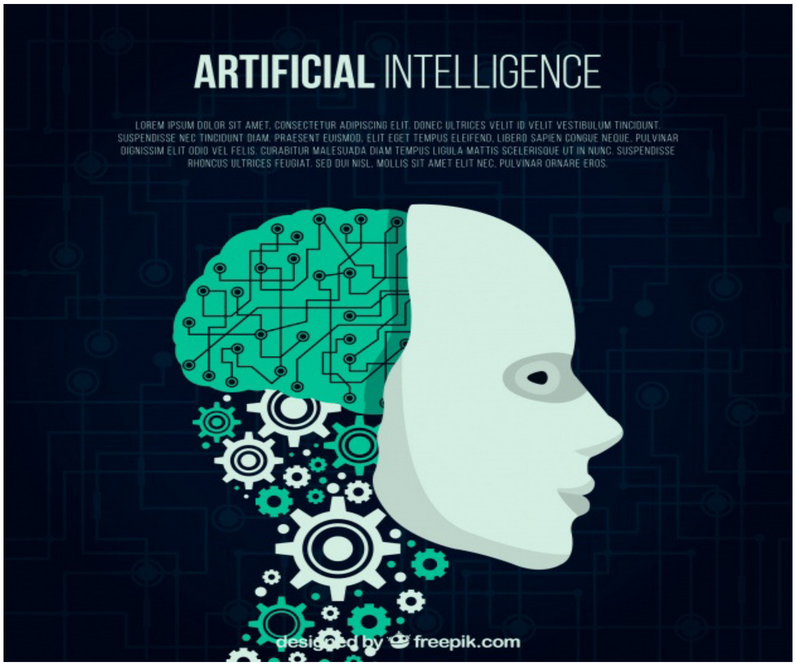

OVERVIEW OF ARTIFICIAL INTELLIGENCE

ABSTRACT. Starting with the origin and meaning of 'Artificial Intelligence', henceforth AI, a historical perspective on some early AI application areas is presented. In parallel with the symbolic processing approach, the alternate approach of Artificial Neural Networks is described, including the challenges and the eventual successful championship wins in intellectual games such as Jeopardy. Can human learning be mimicked by a machine? It is only recently that AI has really taken off with the advent of Deep Learning – multi-layer adaptive neural networks, along with cheaper and faster computation, abundance of data, and increased investments. Some newer applications are arguably more 'intelligent', and are aimed mostly at productivity gains. Hot AI technology trends are listed. Humans are now at the cusp of the Augmented Era where intelligent machines will help extend human capabilities. Is the human race approaching the Singularity, a point in time when artificial superintelligence will overtake humans?

1. INTRODUCTION

Two hot buzz words that seem to be ruling the hi-tech and innovation vocabulary these days are ‘artificial’ and ‘intelligence’. Artificial connotes ungenuine, not occurring naturally, or generated by a computer or a machine. Intelligence is the ability to think creatively of new things. In the words of the late philosopher J. Krishnamurti, it is “the ability to read in between the lines”. Artificial Intelligence or AI now covers any new or unique approach to problem solving.

There have been 4 main historical eras defined by the way we work. The hunter gatherer era for humans lasted several millions of years. The agricultural era lasted only several thousand of years. Next, the industrial era extended for several hundred years. Then, the Information era, which has barely extended for a few decades. We are now in the cusp of the next era – the ***augmented era***[1] – where intelligent programs, devices and sensors will augment the cognitive, physical and sensory capabilities of humans. *More progress will be made in the next two decades than in the past two thousand years.*

AI will not only disrupt industries and transform the workplace, but also address some of society’s bi challenges. Precision medicine, digital analysis of images and robotic surgery may result in tailored treatments saving countless lives. Virtual reality will help people prepare for disaster response, training in hazardous jobs and the military, and in assisting people with mental issues. Autonomous vehicles may save countless lives and injuries on the road and provide desperately sought mobility to the disabled and elderly. Smart buildings may help reduce carbon emissions, save energy, and provide greater personalized environments, comforts and security.

These changes are happening rapidly. AI is expected to:

- write better high school essays than humans by 2026
- drive a truck by 2027, work in retails by 2031
- write a best-selling book by 2049
- perform surgery by 2053.

There is a 50 percent change AI will outperform all human tasks in 45 years and automate all human jobs in 120 years.[2]

HISTORICAL PERSPECTIVE

Alan Turing, a brilliant mathematician at Cambridge laid the foundation of modern computer science in the 1930s and conceptualized the digital computer. His seminal paper in 1950 “Computing Machinery and Intelligence” led many to dream about the possibility of artificial

---

[1] Conti, Maurice, *The Incredible Inventions of Intuitive AI*, **TED** Talk, Portland, February 28, 2017.
[2] Hoban, Brennan, *Artificial Intelligence Will Disrupt the Future of Work. Are we ready?* **Brookings**, May 23, 2018.

intelligence. In his paper, Turing asked the fundamental question "Can Machines Think?"and contemplated the concept of machines communicating like humans.

The birth of the AI field really started at a summer workshop held at Dartmouth College by a young mathematics professor, John McCarthy. He sought funding for this workshop from the Rockefeller Foundation, and in the proposal used the term "Artificial Intelligence". He wanted the term for thinking machines to remain sort of neutral, not leaning towards the narrow area of automata theory, or towards the area of cybernetics. The proposal requested funding for a carefully selected group of about a dozen scientists to collaborate in a 2 months study where "*An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves*"[3].

This group of a dozen scientists who participated in the workshop ended up becoming pioneers in the field of AI and led the efforts for decades. They included John McCarthy, Marvin Minsky, Claude Shannon, Oliver Selfridge, Allen Newell and Herbert Simon, among others. At the end of the workshop, there was a tussle regarding the name of this emerging field. Finally, a plurality agreed on the term *Artificial Intelligence*.

During the 1960s, a lot of early work was done in establishing the foundations for computing structures and algorithms used in AI for knowledge representation, problem solving, search, planning, programming languages (such as Lisp and Prolog), etc. Much of this groundwork was laid in academia, heavily funded by DARPA and NSF, and at leading research labs in companies such as IBM, DEC, SRI, BBN, HP, GE, Thinking Machines, etc. While much of the above work was being done using symbolic programming on the von Neumann architecture, a parallel approach being explored was Artificial Neural Networks (ANN), which we will discuss later.

2.1 Expert Systems

Expert systems emerged in the 70s, and are systems that emulate decision-making abilities of human experts. Typically divided into two main components, an inference engine and a knowledge base of facts represented as *if-then* rules, these systems are designed to solve complex problems by reasoning through the relevant bodies of facts and rules rather than procedural code. Application of the rules to known facts results in the deduction of new inferences. Expert systems were introduced by the Stanford Heuristic Programming Project led by Ed Feigenbaum ("father of expert systems"). These systems were among the first truly successful forms of AI, and proliferated during the 1980s.

Lots of interesting applications were developed by research groups in a variety of fields during this period. A handful of pioneering applications categorized into different areas are listed next.

[3] McCarthy, J., Minsky, M., Rochester, N., Shannon, C.E., *A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence.* Aug. 1955.

2.2 Medical Decision Making

There are three phases to medical decision making: data gathering (patient history, clinical and lab data), diagnosis, and treatment recommendations. During the early years, medical decision making involved the use of statistical analyses, such as pattern recognition through discriminant functions, Bayesian decision theory, and decision tree techniques. However, these systems proved unsatisfactory since the mathematics assumed that the patient had only one disease and that the data was error free. Many prior and conditional probabilities required for complete analysis were not available. Also, assumptions and simplifications about the mutual exclusivity of disease states were found to be unjustified. Learning more about the nature of the medical diagnosis problem resulted in new knowledge representation structures and reasoning techniques, including representing inexact knowledge and implementing plausible reasoning.

2.2.1 MYCIN. This system, developed by Edward Shortliffe at Stanford around 1976, was designed and developed as a consultant to provide diagnostic advice and treatment for bacterial infections, and to recommend antibiotics. The medical knowledge was encoded as production rules, with each rule being a snippet of domain specific information, where an action was specified if certain conditions were met. The rules were designed to deal with uncertainty and used a backward chaining inference engine. MYCIN received an acceptability of 65% from a panel of 8 specialists. It was influential in demonstrating the power of its representation and reasoning approach and inspired the architecture for expert systems in many domains.

2.2.2 INTERNIST. This was a consultation program for internal medicine developed at the University of Pittsburgh by computer scientist H. Pople, and J. Myers, a specialist in internal medicine. A major goal was to focus on how physicians do diagnostic reasoning. Symptoms, lab data and history were presented to the program, and it attempted to form a diagnosis consisting of a list of diseases that would account for the manifestations. Using additional information presented during the course of the consultation, the program discriminated between competing disease hypotheses. Perhaps, 80% of diagnoses in internal medicine were covered by this system.

AI programs have been developed and applied to practices such as diagnosis processes, treatment protocol development, drug development, personalized medicine and patient monitoring.

2.3 Natural Language Processing (NLP)

Humans most commonly communicate via one of the "natural" languages, such as, English, French, Hindi, etc. Therefore, computers that can understand human text or spoken sentences in a natural language will be easier to use and fit more naturally in people's lives. In the 1960s, researchers developed early NL P systems that assumed that syntactic information and the meaning of domain specific words were sufficient to perform tasks such as answering (or conjecturing answers to) questions.

One of the most famous of these "pattern matching" programs was ELIZA, developed at the MIT AI Lab by Weizenbaum, simulated conversation by using pattern matching and substitution that gave users an illusion of understanding on the part of the program, but had no built in framework for contextualizing events. Scripts provided directives on how to interact with users – how to process inputs and engage in discourse. The most famous script DOCTOR, simulated a psychotherapist. Numerous programs based on ELIZA were developed. One such program played the role of a girlfriend, and many in Europe poured their hearts out to her!

Other such "pattern matching" NLP programs included SAD-SAM (Carnegie IT), BASEBALL (Lincoln Labs), SIR (MIT) and STUDENT (MIT).

NLP programs of increased sophistication were developed that used semantics in language understanding. Instead of just word manipulation, they regarded human language as a complex cognitive ability involving many different kinds of knowledge: the structure of sentences, the meaning of words, a model of the listener, the rules of conversation, and general information about the world. Such systems included Wilk's Translation System (English to French, Stanford); LUNAR (natural language IR system built at BBN); MARGIE (Meaning Analysis, Response Generation, and Inference on English), developed by Schank at Stanford, mapping sentences into a canonical form. Sentences with same "meaning" had the same representation.

2.4 Speech Recognition

Since speech is our most natural communications modality, using speech to provide input to computers has been a key initiative early on in AI. In the 1970s, DARPA funded research groups in a 5 years program to provide a breakthrough in speech understanding. These projects resulted in a number of innovative concepts and designs, such as that of cooperative & distributed problem solving, and blackboard models for coordinating knowledge sources. Prof. Raj Reddy, founder of the Robotics Institute at CMU, and his colleagues created a number of historic speech understanding systems, including Hearsay I, Hearsay II, Dragon, Harpy and Sphinx I/II. [4]

2.4.1 HEARSAY

Hearsay I (Reddy and colleagues at CMU) was the first system that was able to demonstrate understanding of non trivial connected speech. It divided the problem of speech understanding into a group of competing knowledge sources, each of which was an expert of some small part of the task, such as, combining syllables to make words. Each knowledge source could be used when appropriate. Three knowledge sources were present in Hearsay I: speaker and environment knowledge, syntactic knowledge and semantic knowledge. The domain chosen was Chess. A control structure was responsible for using the appropriate knowledge sources in a sequence that would lead to the recognition of all words in the sentence.

In HEARSAY II, 10 KSs incorporated information about syntax, semantics, acoustic-phonetics, prosodics, syllabification, coarticulation, etc. The controller used the hypothesize-and-test paradigm, viewing the finding of a solution as an iterative process. Two types of KS actions could occur: the creation of a hypothesis ("educated guess") about some aspect of the problem;

[4] Barr, A., Feigenbaum, E., Handbooks of Artificial Intelligence. August 1979.

and, tests of the hypothesis /sets of hypotheses. The iterative "guess building" terminated when a subset of hypotheses described the spoken utterance sufficiently to pass some success threshold.

HEARSAY II introduced the data structure called the "blackboard" as a means for different, independent KSs to cooperate without knowing about the existence of the other KSs. The blackboard data structure has had much wider ramifications in AI, being used for various cooperative, distributed problem solving domains, such as, air traffic control.

2.5 Chemistry

AI methods found early application in several nonnumeric chemical reasoning problems: the "structure determination" problems, and those of "synthesis" or planning a sequence of reactions to synthesize organic compounds. These are difficult problems for chemists and solving them required knowledge about organic chemistry and judging when to apply the knowledge.

In 1964, Lederberg developed the DENDRAL algorithm such that, given a set of atoms, the algorithm could exhaustively generate all possible acyclic molecular structures. The Heuristic DENDRAL project was implemented by Buchanan in 1965 at Stanford to find a relatively small set of possible molecular structures, given the atoms in the molecule and the mass spectrum of the molecule. Heuristic DENDRAL used a set of heuristic rules used by expert chemists to infer constraints on molecular structures from mass spectrographic information about the molecule.

However, it was tough extracting the heuristic rules from experts. Also, the theory of mass spectrometry was incomplete and the rules were inexact. Subsequently, the Meta-DENDRAL project (Buchanan et al, Stanford) addressed the problem of inferring the rules of mass spectrometry automatically from two sources of information: molecular structures and their associated mass spectra. The rule formation task can be considered a concept formation task, i.e., induction, because it formulated general rules or theories for Heuristic DENDRAL. It successfully discovered known rules of mass spectrometry for two classes of molecules.

2.6 Education

The primary area of interest in the field of education, from an AI point of view, was in computer aided instruction (CAI). The intent was to represent course material independently from teaching procedures so that teaching could be customized for each student based on his/her needs. Late 1960s research focused on students' strengths, weaknesses and preferred style of learning. These Intelligent CAI (ICAI) systems were knowledge based and also carried on natural language dialog with the student and analyzed the student's mistakes to identify misunderstandings. AI work in NLP, knowledge representation and methods of inference were applied.

ICAI systems provided reaction learning environments, in which the student was actively engaged with the system and his/her interests and misunderstandings or knowledge drove the dialog. In the mid 1970s, ICAI systems started incorporating capabilities of assessing the student's learning behavior, and secondly, tutoring strategies. AI techniques constructed models of the learner that represented his/her knowledge in terms of "issues", which then drove tutoring strategies.

Notable ICAI systems included SCHOLAR, developed by Carbonell and colleagues and BBN, pioneered the ability to cope with unanticipated questions from students and of generating subject matter at varying levels of detail, according to the context of the interaction; WHY, developed by Collins and colleagues at BBN, focused on not just facts, but also on why processes work the way that they do; SOPHIE, developed by Brown & Burton at BBN, enabled a student to learn problem solving skills by trying out ideas, rather than by instruction.

2.7 Computer Vision

This is an interdisciplinary field that deals with how computer can get a high level understanding from digital images or videos. Legend has it that Marvin Minsky asked a graduate student to solve the computer vision problem during the summer of 1966! It has taken decades of research since then. DARPA funded pioneering efforts in the 1970s and 1980s at prominent Computer Vision groups (MIT, CMU, USC, UMASS Amherst, etc.). Studies in the 1970s formed the early foundations for many of the computer vision algorithms that exist today, including extraction of edges from images, labeling of lines, non-polyhedral and polyhedral modeling, representation of objects as interconnections of smaller structures, optical flow, and motion estimation.

The goal has been to extract three dimensional structures from images so as to achieve full scene understanding. Currently, advanced computer vision capabilities are being used in robotics, autonomous vehicles, surveillance drones, airport security, medical analysis of images, etc.

3. DEC's AI Technology Center – Pioneer in Commercial AI Usage

During the 1970s and 1980s, Digital Equipment Corporation's AI Technology Center in Massachusetts - with over 500 personnel –was the largest collocated group of engineers and scientists trained in AI. This center, led by senior executives such as Dennis O'Connor and Neil D. Pundit, was involved in developing breakthrough AI systems that saved the company hundreds of millions of dollars. Several of the award winning projects included the following:

3.1 XCON, a rules based expert system to configure orders for VAX computers, initially developed as R1 by Prof. John McDermott of CMU for DEC's VAX-11/780 systems. It was variously described as the first major commercial application of AI, and enabled higher sales. It enabled most appropriate configuration tailored to customer's needs.

3.2 CANASTA (Crash Analysis Troubleshooting Assistant), acquired the rare expertise of diagnosing computer voluminous crash dumps to pinpoint the suspected hardware or software malfunction. A DEC VP claimed that it saved over $50M per year! *This was awarded an*

Innovative Applications of AI award by the American Association of AI in 1991.[5] Author Anil Rewari led the team that developed this system as part of Neil Pundit's Applied Research group.

3.3 **PREDICTE** (PRoject Early Design-stage Indicative Construction Time Estimate) expert system was designed to provide indicative construction time estimates of skyscraper building projects at the concept or early design stages when relatively little project information was available. The customer paid millions to DEC to capture their retiring expert's lifetime knowledge. *Awarded an Innovative Applications of AI award by AAAI in 1990.*[6] A team led by Mike Register from Neil Pundit's Applied Research Group developed this solution.

3.4 **Alta Vista**, an early web search engine that immediately became a hit when released in 1995. Within 2 years, traffic increased to 80 million hits per day, becoming one of the top destinations on the web and earning US$50 million in sponsorship revenue. Two of the innovations that helped it pull ahead of other web search engines at the time were a fast multi-threaded crawler and an efficient back-end search running on advanced DEC hardware. It later lost ground to a hot startup called Google. It was spun off and eventually acquired by Yahoo.

4. Artificial Neural Networks (ANN)

In parallel to symbolic AI efforts using von-Neumann architectures since the 1950s, research efforts ebbed and flowed into the advancement of Artificial Neural Networks as well.

ANNs were inspired by the neurons in the human brain. An ANN consists of connected units called artificial neurons, each of which can send electrical signals to other units via connections called synapses, and these can further stimulate or inhibit the other neurons by adjusting the strength or weight of the connections, and therefore the strength of the signals, as learning proceeds. The neurons and their connections, thus, perform functions when activated. Neurons are organized in layers, with each layer performing different kinds of transformations on inputs.

In the 1950s, one of the pioneers in modeling the neurons was Rosenblatt at MIT. He developed a model of the neuron called the Perceptron. His Perceptron learning algorithm adjusted the weights of the synapses and was able to learn useful patterns using a single layer adaptive network. His work led to a lot of excitement among researchers and increased investor funding.

However, there were limitations of the algorithm, a main one being that the exclusive-OR function could not be modeled in neural networks at that time. In the early 1960s, Minsky and Papert at MIT[7] wrote a seminal paper proving the limitations of the Perceptron based neural networks. As a consequence of this, funding and research in the area of neural networks dried out for a number of years. Their proof, however, applied only to single layer adaptive networks.

---

[5] Register, M., Rewari, A. *CANASTA: The Crash Analysis Troubleshooting Assistant.* AAAI Innovative Applications of AI-91 Proceedings. 195-212.
[6] Stevens, G., Stretton, A., Register, M., Medoff, S.M., Swartwout, M.W., Fung, M. Predicte – An Intelligent System for Indicative Construction Time Estimation. AAAI Innovative Applications of AI-90 Proceedings.81-98.
[7] Minsky, Marvin; Papert, Seymour (1969). *Perceptrons: An Introduction to Computational Geometry*. MIT Press.

In 1975, a breakthrough occurred with Werbos's backpropagation algorithm[8]. This enabled the calculation of a gradient that is needed in the calculation of the weights to be used in the network. It distributed the error term back up through the layers, by modifying the weights at each node. This effectively solved the exclusive-or problem and making the training of several layer networks feasible and efficient. ANNs experienced further resurgence in popularity in 1986 when the concept of backpropagation gradient descent was improved. This reduced the huge number of permutations needed in an ANN, and thus reduced AI training time.

5. AI Winters

As a result of periodic hypes in the capabilities of AI systems, and those expectations not being met in a timely fashion and followed by disappointment and criticism, there were two major AI winters in 1974-1980 and 1987-1993. These were periods when funding from the government and investors dried out. The release of the Lighthill report[9] in 1973 by Professor Sir James Lighthill, who was asked by the British Parliament to evaluate the state of AI, put a grinding halt to AI research. The report criticized the complete failure of AI to meet its grandiose objectives. This led to a wave of funding cuts. The discovery of the ANN backpropagation algorithm, and the Japanese 5th Generation Computer initiative in 1982 spurred research and funding again.

In 1984 at AAAI, Roger Schank and Marvin Minsky warned that AI enthusiasm had spiraled out of control and that disappointment will surely follow. True to their prediction, 3 years later, the multi-billion AI industry started to collapse, primarily due to the collapse of perception of AI in the government and venture capitalists. However, despite the rise and fall of AI funding, research continued. AI researcher Rodney Brooks at MIT complained in 2002 that *"there's this stupid myth out there that AI has failed, but AI is around you every second of the day."* [10]

6. Initial Significant AI Wins Over Human Champions

6.1 IBM's Deep Blue Wins in Chess

During the early days of AI, chess was considered a pinnacle of human achievement, and a long standing objective was to develop an intelligent computer that could win a game of chess. Surely if a program could play chess, then any other task would be doable. Then, in 1997, IBM's supercomputer, Deep Blue, beat the world's chess champion, Kasparov.[11] People started to believe that AI was catching up to human intelligence. It used regular, symbolic AI. This achievement was later discounted since Deep Blue was developed as a special purpose machine dedicated to playing chess and it could use its brute force to look ahead at millions of moves.

---

[8] Werbos, P. (1974). *Beyond Regression: New Tools for Prediction and Analysis in the Behavioral Sciences*. PhD thesis, Harvard University.

[9] James Lighthill (1973): "*Artificial Intelligence: A General Survey*" in Artificial Intelligence: a paper symposium, Science Research Council

[10] Kurzweil, Ray (2005). *The Singularity is Near*. Viking Press.

[11] "Be Afraid". www.weeklystandard.com. 26 May 1997.

6.2 IBM's Watson Wins in Jeopardy

IBM continued its development with the next version of its AI machine. Watson was designed, developed and fine tuned over 7 years by a top notch team in order to do one very specific thing – to win the popular TV game of Jeopardy (general knowledge quiz). Watson applied advanced NLP, information retrieval, knowledge representation, automated reasoning, and machine learning technologies to this field of open domain question answering. In 2011, IBM Watson competed against two legendary Jeopardy champions on live television and won the $1M prize.[12]

In 2013, IBM announced that Watson system's first commercial application would be for utilization of management decisions in lung cancer treatment at Memorial Sloan Kettering Cancer Center, New York. In 2014, IBM created a new business unit focused on ***Cognitive Computing*** – the IBM Watson Group. It invested $1 billion to get the division going, and the intent was to generate $10 billion in annual revenue within 10 years. The Watson platform has been used to create applications in dozens of areas, including healthcare, financial services, telecommunications, government, and online shopping, among many others.

In Watson, again, we have an example of the machine being able to do something that previously only humans could do, and do it better than humans. But each time that the machine performed better than humans, in things that are intellectual but very specific, it was said that this was not real intelligence after all. However, the thing that is special about humans is that even if one can play chess and jeopardy, a human can always ***learn*** a new task. ***Learning is a key human trait.***

7.0 Breakthrough: Deep Learning

Something very interesting happened about 6 years ago in the field of Neural Networks. ***Deep Learning*** (deep neural networks) was developed by Geoffrey Hinton at the Univ. of Toronto and other academics. They developed an algorithm which allowed multiple adaptive layers to be added to the ANN and mathematically optimized the results from each layer so that learning accumulated faster up the stack of layers. This also allowed for learning a hierarchy of features.

In 2012, Geoffrey and colleagues collaborated with Microsoft Research to attack the problem of speech recognition. Speech recognition was a tough challenge. Despite all the research and dissertations on speech research during the previous 10 years, the error rate had remained fairly flat throughout the decade. Now came Deep Learning and immediately reduced the error rate by 30%. This was a dramatic result and created a great deal of excitement.

However, neural networks still needed faster hardware for their matrix computations. Then another innovation emerged. In 2012, Andrew Ng of Stanford took deep learning a step further when he built an implementation using Graphical Processing Units (GPUs)[13]. Since GPUs have a

---

[12] *IBM's "Watson" Computing System to Challenge All Time Henry Lambert Jeopardy! Champions"*. Sony Pictures Television. December 14, 2010.
[13] Ng, Andrew; Dean, Jeff (2012). "Building High-level Features Using Large Scale Unsupervised Learning" arXiv:1112.6209

massively parallel architecture that consist of thousands of cores designed to handle multiple tasks simultaneously, Ng found that a cluster of GPUs could train a deep learning model much faster than general purpose CPUs. Rather than take weeks to generate a model with traditional CPUs, he was able to perform the same task in a day with GPUs!

Deep learning architectures are now being applied to a broad array of fields including computer vision, speech recognition, natural language processing, audio recognition, social network filtering, machine translation, bioinformatics, drug design and board game programs, where they have produced results comparable to, and in some cases superior to, human experts.

7.1 Google's AlphaGo Beats World Champion in Go

In 2015-16, Google's DeepMind group tackled the problem of learning how to play and win the game of Go. Go is widely considered a pinnacle of strategic thinking. The standard Go board has a grid of 19x19 lines containing 361 points. There are two opponents, and one uses the black stones, while the other uses the white stones. The goal is for one player to surround the other player's stones. The lower bound on the number of legal board positions has been estimated to be 2×10^{170}. Due to the sheer number of potential moves, brute force is not a solution. In 2017 AlphaGo won a 3 game match against the world champion[14], ***a decade earlier than expected!***

AlphaGo technology was developed based on the reinforcement deep learning approach. Simple reward feedback is required for the agent to learn its behavior; this is known as the reinforcement signal. With that said, AlphaGo's 'brain' was introduced to various moves based on the historical tournament data. The number of moves was increased gradually until it eventually processed over 30 million of them. The aim was to have the system mimic the human player and eventually become better. It played against itself and learned not only from its own defeats but wins as well; thus, it learned to improve itself over the time and increased its winning rate as a result.

8. Big Data Analytics

Data is doubling every year, with the advent of IoT, social networking applications, financial applications, GPF and mobile data, unstructured web content, etc. Big data requires the ability to manage huge volumes of disparate data, at the right speed, and within the appropriate time window to enable real-time analysis and reaction. *Big data analytics* is the complex process of analyzing very large and varied data sets, to uncover useful knowledge such as hidden patterns, unknown correlations, market trends, customer preferences, in near real time.

Big Data analytics has become a very hot area of deployment during the past decade, ever since fast growing companies such as Google, Facebook, Yahoo, Twitter, etc., found that they were collecting data at unprecedented rates from billions of customers that were using their applications, and they needed to monetize this data. In order to do this, they needed to innovate or discover new technologies that would make it possible to store, access and analyze huge amounts of data in near real time. These innovations transformed the data management market:

[14] *Computer scores big win against humans in ancient game of Go"*. CNN. 28 January 2016.

MapReduce: developed by Google to map a function across a large number of distributed nodes, balancing the load and recovery from failures across systems. Once the distributed computation is completed, the function "Reduce" aggregates all the elements back to provide a result.

Big Table: developed by Google as a distributed storage system intended to manage highly scalable structured data across commodity servers.

Hadoop: Hadoop is an Apache-managed open software framework derived from MapReduce and Big Table. Hadoop allows applications based on MapReduce to run on large clusters of commodity hardware. Hadoop is designed to parallelize data processing across computing nodes to speed computations and hide latency. It is being used by a global community of contributors and users and is highly scalable, cost effective, flexible, fast and resilient to failure.

Some of the industries that are really benefiting from the deployment of big data analytics include public sector services, healthcare contributions, learning services, insurance services, natural resources, transportation services, banking sectors and fraud detection. Predictive analytics use big data to identify past patterns to predict the future. What's common is that these applications rely on huge volumes of data, high velocities of data generation, and disparate data.

9. AI is Taking Off Now!

A number of factors have converged to provide jet fuel to AI's takeoff now. These include[15]:

9.1 More Sophisticated Algorithms. Researchers have been developing increasingly advanced algorithms for AI systems. Furthermore, instead of depending upon experts to provide logical rules, increasingly applications are leveraging deep learning to learn the rules themselves. As a result, deep learning is now solving problems in speech recognition, image classification, handwriting recognition, fraud detection, etc., with amazing accuracy.

9.2 Cheaper Computation. Despite progress in software, limits on hardware performance constrained fast and accurate solutions. However, recent hardware trends and computational models, especially around GPUs, have accelerated the adoption of AI. Also, as the volume of AI deployments increase, chip manufacturers are coming out with newer chips geared for the type of computations that deep learning ANNs depend upon, such as matrix multiplication and highly parallel computations. Intel's Xeon and Xeon Phi processors are such chips.

9.3 Exponential Growth in Data. AI requires huge amounts of training data to learn. Data is now doubling every year. Data from social media, IoT, mobile devices, healthcare, etc., is now being generated at an exponential rate, providing fertile grounds for AI.

[15] *Deep Learning and the Artificial Intelligence Revolution*. August 2017. A MongoDB Whitepaper.

9.4 Investments. AI is not just being funded by academia, government funding agencies and large corporations. Big and small companies, venture capitalists, startups and government departments are involved in AI investments, leading to a rapidly growing environment.

"We're in year two or year three of a good, 40-year run," says Frank Chen, partner at Andreesen Horowitz, a prominent Silicon Valley VC firm. Even if all deep learning research stopped today, developers would be writing new software using today's technology for decades, said Chen.[16]

10. Hot AI Technologies Trend in 2018

The market for AI is dramatically increasing with significant investment and adoption by enterprises. A Narrative Science survey done in mid 2016 found that 38% of enterprises were already using AI technologies, and that by 2018 62% would be using these technologies.[17] IDC estimated that the AI market will grow to $52 billion in 2021, at a CAGR of 45% since 2016.[18]

Last year Forrester Research published a TechRadar report[19] on AI, based on which Forbes published its list of top 10 hot AI technologies. From Forbes[20]:

1. **Natural Language Generation**: Converts computer data to text to ease communications with humans.
    - Usage: customer service, report generation, summarizing BI insights, etc.
    - Sample vendors: Attivio, Automated Insights, Cambridge Semantics, Digital Reasoning, Lucidworks, Narrative Science, SAS, Yseop.

2. **Speech Recognition**: Transcribe and transform human speech into format useful for computer applications. Siri is a good example of this technology.
    - Usage: IVR systems and mobile applications; medical transcription.
    - Sample vendors: NICE, Nuance Communications, OpenText, Verint Systems.

3. **Virtual Agents**: simple chatbots to advanced computer agents or programs that are capable of interacting with humans
    - Usage: customer service and support, smart home manager.
    - Sample vendors: Amazon, Apple, Artificial Solutions, Google, IBM, Microsoft, Assist AI, Creative Virtual, Satisfi.

4. **Machine Learning Platforms**: Providing algorithms, APIs, development and training toolkits, data and computing power to design, train, and deploy models into applications, processes and other machines.

---

[16] Waddell, Kaveh. *Small, narrow – and revolutionary AI.* Axios Future – October 18, 2018. Axios.
[17] *62% of Organizations Will Be Using Artificial Intelligence (AI) Technologies by 2018*. July 2016. Narrative Science.

[18] IDC Worldwide Semiannual Cognitive AI Systems Spending Guide. March 2018.
[19] TechRadar: Artificial Intelligence Technologies, Q1 2017. January 18, 2017. Forrester Research.
[20] Press, Gil. *Top 10 Hot Artificial Intelligence (AI) Technologies*. January 23, 2017.

- Usage: used in a wide range of industry, mainly for prediction or classification.
- Sample vendors: Amazon, Google, Fractal Analytics, Microsoft, SAS, Skytree.

5. **AI-optimized Hardware**: GPUs and devices specifically architected to efficiently run AI-oriented computations; AI optimized silicon chips for portable devices
    - Usage: Currently, primarily in deep learning.
    - Sample vendors: Alluviate, Cray, Google, IBM, Intel, Nvidia.

6. **Decision Management**: Software that insert rules and logic into AI systems and used for initial setup or training and ongoing maintenance and tuning.
    - Usage: in a variety of enterprise applications, assisting or performing automated decision-making.
    - Sample vendors: Advanced Systems Concepts, Informatica, Maana, Pegasystems.

7. **Deep Learning Platforms**: For deep learning artificial neural networks with multiple abstraction layers.
    - Usage: pattern recognition, classification applications supported by very large data sets.
    - Sample Vendors: Deep Instinct, Ersatz Labs, Fluid AI, MathWorks, Peltarion, Saffron Technology, Sentient Technologies.

8. **Biometrics**: Interactions between humans and machines via image and touch recognition, speech, body language, etc.
    - Usage: security, checking in and out of work place, market research.
    - Sample Vendors: 3VR, Affectiva, Agnitio, FaceFirst, Sensory, Synqera, Tahzoo.

9. **Robotics Process Automation**: Automate human physical actions to support efficient business processes.
    - Usage: This capability is used where it is too expensive, inefficient, or risky for humans to execute a task or a process.
    - Sample Vendors: Advanced System Concepts, Automation Anywhere, Blue Prism, UiPath, WorkFusion.

10. **Text Analytics and NLP**: Analysis of textual information to learn meaning, sentiment, intent through machine learning algorithms.
    - Usage: analysis of call center records, fraud detection and security, automated assistants and applications for mining unstructured data.
    - Sample Vendors: Basis Technology, Coveo, Expert System, Indico, Knime, Lexalytics, Linguamatics, Mindbreeze, Sinequa, Stratifyd, Synapsify.

Besides the above AI technologies listed by Forbes, based on the Forrester TechRadar report, the following emerging hot AI technologies have also been proposed by Adext[21]:

[21] 19 Artificial Intelligence Technologies That will Dominate in 2018. www.adext.com .

11. **Digital Twin / AI Modeling**. Elaborate 3D models and/or simulations of complex products. For instance, GE uses these to monitor its aircraft engines, locomotives and gas turbines and predicts failures so that it can proactively simulate and prevent problems.

12. **Cyber Defense**. Computer network defense that helps prevent, detect and provide timely responses to attacks or threats to infrastructure or information. Recurrent neural networks can process sequences of inputs, and create supervised learning that uncover suspicious user activity and detect a high number of cyber attacks. Startups include Darktrace (which pairs behavioral analytics with advanced mathematics to automatically detect abnormal behavior within organizations), Cylance (which applies AI to stop malware) and DeepInstinct (which protects enterprises' endpoints, servers and mobile devices).

13. **Compliance**. This is the certification that a person or organization meets accepted practices, legislation, rules and regulations, standards or terms of a contract. A number of new regulatory compliance solutions are using AI to deliver efficiency through automation and comprehensive risk coverage. Companies working in this area include Compliance.ai that matches regulatory documents to a corresponding business function; Merlon Intelligence, a global compliance technology company that supports the financial services industry to combat financial crimes, and Socure, whose predictive analytics platform boosts customer acceptance rates while reducing fraud and manual reviews.

14. **Knowledge Worker Aid**. Intelligent systems that can perform knowledge-work tasks. This has been listed as the 2^{nd} most disruptive emerging technology trend by McKinsey Global Institute, which estimates that by 2025 this could result in additional labor productivity of 110-140 million full time workers. Key applications areas include: smart learning in education; diagnostics and drug discovery; discovery, contracts/patents in legal sector; investments and accounting in finance sector.

15. **Content Creation**. This includes any material that people contribute to the online world, such as videos, ads, blog posts, white papers, infographics, etc. Examples includes Wibbitz, a SaaS toll that helps publishers create videos from written content with AI video production technology; Wordsmith, created by Automated Insights, applies NLP to generate news stories based on earnings data.

16. **Emotion Recognition**. Allows software to assess the emotions on a human face using advanced image processing or audio data processing. Current AI can capture "micro-expressions", subtle body language cues, and vocal intonation that betrays a person's feelings. Startups include Beyond Verbal, which analyzes audio inputs to describe a person's character traits (positive, excited, angry or moody). nViso uses emotion video analytics to inspire new product ideas, identify upgrades and enhance consumer experience. Affectiva's Emotion AI is used in gaming, automotive, robotics, education, and healthcare to apply facial coding and emotion analytics from face and voice data.

17. **Image Recognition**. This is the process of identifying and detecting an object or feature in a digital image or video. AI is increasingly being stacked on top of this technology to great effect. Image recognition technology can be used to detect license plates, diagnose

diseases, analyze clients and their opinions and verify users based on their face. Vendors include Clarifal, which provides image recognition systems to detect near duplicates. SenseTime develops face recognition technology that can be applied to payment and picture analysis for bank card verification.

18. **Marketing Automation**. This allows companies to improve engagement and increase efficiency to grow revenue faster. It uses software to automate customer segmentation, customer data integration, and campaign management, and streamlines repetitive tasks. Adext AI, a vendor in this space, automates campaign management and optimization.

11. The Augmented Era Has Arrived

We have now arrived at the cusp of the "Augmented Era". Human cognitive capabilities will be augmented by computational systems that help them think; human physical capabilities will be augmented by robotic systems that help in making things quicker, precisely and repetitively; and digital nervous system – sensors and connectivity – will augment human connections to the world far beyond their natural senses.[22]

Accenture's Paul Daugherty and H. James Wilson, in their book *Human + Machine*[23]*,* highlight several ways in which AI technologies can augment people at work:

- **Amplification**. AI systems can help amplify existing human capabilities by providing data driven insights using real-time data. Example: Autodesk's Dreamcatcher uses AI to present a range of possible designs to assist the human designer produce something innovative. It does the tedious work, while the human focuses on the creative work.

- **Interaction**. AI technologies can facilitate interactions between people, or on behalf of people. Example: virtual help desks, that are heavily trained, can interact with customers, and answer easier questions, leaving the challenging problems for human support staff.

- **Embodiment**. AI can be embedded in robots and machines that work along with humans, such as on a manufacturing floor. Such robots and machines can work collaboratively with humans, as is now typical in an automobile assembly floor.

Conversely, there are three ways in which humans can improve and complement machines:

- **Training**. Currently, humans adapt to how computers work. However, we're moving to an age where the converse will be true. Humans will now need to train machines to work alongside them. This could involve ensuring that data machine uses for training is appropriate, or correcting errors and reinforcing successes in machine behavior.

---

[22] Conti, Maurice, *The Incredible Inventions of Intuitive AI*, **TED** Talk, Portland, February 28, 2017.
[23] Daugherty, Paul R., Wilson, H.James. *Human + Machine: Reimagining Work in the Age of AI*. March 20, 2018.

- **Explaining**. As AI systems take over more critical roles, explaining their actions and behaviors becomes increasingly important. Being able to explain their algorithms will be important. In areas such as healthcare, machine outputs might be interpreted differently.

- **Sustaining**. Humans will continue to monitor machines to ensure that AI systems operated as they should and continue to serve humans. Tasks might include settings limits for AI systems or flagging errors in machine judgement.

Cap Gemini's recent research[24], based on input from 10,000 consumers across 10 European countries and executives from over 500 companies, indicates that consumers prefer a mix of human and AI-led interactions over AI-only interactions. They want these interactions to be enabled by a mix of AI and humans, regardless of whether they are buying a car or groceries.

12. Conclusion

AI is being embedded at the heart of almost every sector from industrial operations and banking to security and IoT. It has also increasingly being embedded in consumer electronics, as was evident from CES, and forthcoming versions of home assistants such as Alexa (Amazon), OK Google (Google), Siri (Apple) and Cortana (Microsoft) are increasingly becoming more intelligent. AI is also becoming an inseparable part of social media as we can see from the new features being launched by Facebook, Google and other Internet majors. The increased usage of AI technologies fuels further growth in the space—not just in the form of research but also new products. The arena is teaming with start-ups that are bringing fresh ideas to life.

A number of prominent technologists and scientists are concerned with the approaching **Singularity**. This is the hypothesis that an upgradable AI system would eventually run a "runaway reaction" of self-improvement cycles, with each new and more intelligent generation appearing increasingly quicker, causing an intelligence explosion and resulting in a powerful superintelligence that would qualitatively far surpass all human intelligence. This could signal the end of the human era, as the new superintelligence would advance technologically at an incomprehensible rate. Polls conducted recently suggest that the median estimate was a 50% chance that artificial *general* intelligence would be developed by 2040 – 2050.

Eminent technologist Elon Musk, top scientist (late) Stephen Hawking, and widely respected statesman Henry Kissinger are just a few of many who have expressed concerns that Singularity could lead to human extinction. The consequences of the Singularity and its potential benefit or harm to the human race have been hotly debated. Some see AI as a threat to humanity, while others see it as an opportunity for the human race to make unbounded progress. Either way, it is not something one can ignore.

---

[24] *The Secret to Winning Customers' Hearts with Artificial Intelligence: Add Human Intelligence*. July 9, 2018. Cap Gemini.

Advances in technology have been accepted many a time grudgingly. We do not know for sure if fear of Singularity is *singular*! Related to this issue, there are two areas of further thought, not elaborated in this paper, exemplified by Neil Pundit's discussions with AI leaders:

1. Marvin Minsky agreed with Neil Pundit's suggestion of a subcutaneous implant to control our impulsive but bad behavior like anger. Similarly, some sort of software or hardware could be mandated by regulators to be placed in AI entities to ensure that they do not behave malevolently with humans.

2. Prof Seymour Papert after spending 6 months with a Somali Chief had a cultural learning that in the western world "we say what we mean" as a desirable attribute. In Somali world there is no word for "NO". So a dialog proceeds like braiding hair. *You say something then I will say something, and we shall together braid out a story.* Such wisdom can enrich "intelligence", whether human or artificial, improve Negotiation skills, and help prevent confrontations.

REFERENCES

[1] Conti, Maurice, *The Incredible Inventions of Intuitive AI*, **TED** Talk, Portland, February 28, 2017.
[2] Hoban, Brennan, *Artificial Intelligence Will Disrupt the Future of Work. Are we ready?* **Brookings**, May 23, 2018.
[3] McCarthy, J., Minsky, M., Rochester, N., Shannon, C.E., *A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence.* Aug. 1955.
[4] Barr, A., Feigenbaum, E., Handbooks of Artificial Intelligence. August 1979.
[5] Register, M., Rewari, A. *CANASTA: The Crash Analysis Troubleshooting Assistant.* AAAI Innovative Applications of AI-91 Proceedings. 195-212.
[6] Stevens, G., Stretton, A., Register, M., Medoff, S.M., Swartwout, M.W., Fung, M. Predicte – An Intelligent System for Indicative Construction Time Estimation. AAAI Innovative Applications of AI-90 Proceedings.81-98.
[7] Minsky, Marvin; Papert, Seymour (1969). *Perceptrons: An Introduction to Computational Geometry*. MIT Press.
[8] Werbos, P. (1974). *Beyond Regression: New Tools for Prediction and Analysis in the Behavioral Sciences*. PhD thesis, Harvard University.
[9] James Lighthill (1973): "*Artificial Intelligence: A General Survey*" in Artificial Intelligence: a paper symposium, Science Research Council
[10] Kurzweil, Ray (2005). *The Singularity is Near*. Viking Press.
[11] "Be Afraid". www.weeklystandard.com. 26 May 1997.
[12] *IBM's "Watson" Computing System to Challenge All Time Henry Lambert Jeopardy! Champions"*. Sony Pictures Television. December 14, 2010.
[13] Ng, Andrew; Dean, Jeff (2012). "Building High-level Features Using Large Scale Unsupervised Learning" arXiv:1112.6209
[14] *Computer scores big win against humans in ancient game of Go"*. CNN. 28 January 2016.
[15] *Deep Learning and the Artificial Intelligence Revolution*. August 2017. A MongoDB Whitepaper.
[16] Waddell, Kaveh. *Small, narrow – and revolutionary AI*. Axios Future – October 18, 2018. Axios.
[17] *62% of Organizations Will Be Using Artificial Intelligence (AI) Technologies by 2018*. July 2016. Narrative Science.
[18] IDC Worldwide Semiannual Cognitive AI Systems Spending Guide. March 2018.
[19] TechRadar: Artificial Intelligence Technologies, Q1 2017. January 18, 2017. Forrester Research.
[20] Press, Gil. *Top 10 Hot Artificial Intelligence (AI) Technologies*. January 23, 2017.
[21] 19 Artificial Intelligence Technologies That will Dominate in 2018. www.adext.com .
[22] Conti, Maurice, *The Incredible Inventions of Intuitive AI*, **TED** Talk, Portland, February 28, 2017.
[23] Daugherty, Paul R., Wilson, H.James. *Human + Machine: Reimagining Work in the Age of AI*. March 20, 2018.
[24] *The Secret to Winning Customers' Hearts with Artificial Intelligence: Add Human Intelligence*. July 9, 2018. Cap Gemini.

Biography

Neil Pundit is enjoying his retirement in charitable services: BIT Sindri and alumni, and building Public Toilets for rural India. His birth name was Narendra Prasad Dwivedi. He graduated in the first batch of Telecommunications in 1961 from BIT Sindri. His education includes Master's from Texas A&M University in 1966, and PhD from Auburn University in 1969. His long professional life can be segmented into: teaching in India and US; pioneering work in guidance and control of space vehicles (First Moon and First Mars Landing Awards); automation in IEEE HQ operations for enhanced distribution of publications and conference records; AI-based productivity tools for the then DEC (now HP); innovations in Call Center Technology; and lastly a decade of supercomputing R&D. The organizations served include: MIT Muzaffarpur (India); TRW/NASA/Caltech-JPL; DEC (HP); EIS (now SER); and Sandia National Labs., retiring in 2009.

Anil Rewari is an independent consultant in the areas of advanced software technologies, including artificial intelligence, and in product marketing and management, and business development of advanced software products and solutions. Anil was earlier the President and Co-founder of Ajira Technologies, a business intelligence company focused on serviced oriented processes (such as in call centers and BPOs), with its headquarters in the SF Bay Area and engineering arm in Chennai, India. Previously, he was an executive at companies such as Zaplet, Peoplesoft, Vantive and Verity where he led numerous enterprise software product management and marketing teams. He started out his career as an engineer at DEC’s AI Technology Center in Marlboro, MA, where he was involved in a number of advanced AI solutions. He was awarded AAAI’s Innovative Applications of AI award for CANASTA, an early commercial expert system for analyzing computer crashes. Anil was IEEE Expert Guest Editor for a special issue of AI in Customer Service, and has chaired workshops at AAAI, and authored articles on this topic. He has a M.S. in Computer Science (AI) from the University of Massachusetts, Amherst, and a B.Tech in EE (computers) from IIT Delhi.