

The implementation of artificial intelligence and its future potential

Škavić, Fran

Undergraduate thesis / Završni rad

2019

Degree Grantor / Ustanova koja je dodijelila akademski / stručni stupanj: **University of Zagreb, Faculty of Economics and Business / Sveučilište u Zagrebu, Ekonomski fakultet**

Permanent link / Trajna poveznica: <https://urn.nsk.hr/urn:nbn:hr:148:308959>

Rights / Prava: [In copyright](#) / [Zaštićeno autorskim pravom.](#)

Download date / Datum preuzimanja: **2024-04-24**



Repository / Repozitorij:

[REPEFZG - Digital Repository - Faculty of Economics & Business Zagreb](#)



University of Zagreb
Faculty of Economics and Business

**The implementation of artificial intelligence and its future
potential**

Undergraduate thesis

Fran Škavić, 0066231865

Mentor: Mario Spremić, PhD

Zagreb, September 2019.

STATEMENT ON ACADEMIC INTEGRITY

I hereby declare and confirm with my signature that the final paper dissertation is exclusively the result of my own autonomous work based on my research and literature published, which is seen in the notes and bibliography used. I also declare that no part of the paper submitted has been made in an inappropriate way, whether by plagiarizing or infringing on any third person's copyright. Finally, I declare that no part of the paper submitted has been used for any other paper in another higher education institution, research institution or educational institution.

Student:

In Zagreb, September 2nd
(date)


(signature)

Table of contents:

1. Introduction	1
2. A history of artificial intelligence	2
3. Types of artificial intelligence	8
3.1. Type I:	8
3.1.1. Artificial Narrow Intelligence (ANI)	8
3.1.2. Artificial General Intelligence (AGI)	10
3.1.3. Artificial Superintelligence (ASI)	11
3.2. Type II:	12
3.2.1. Reactive machines	12
3.2.2. Limited memory	13
3.2.3. Theory of mind	13
3.2.4. Self-aware	16
4. Ways of achieving AI	18
4.1. Machine Learning (ML)	18
4.2. Deep Learning (DL)	20
4.2.1. TensorFlow	21
5. Application of artificial intelligence in business	26
5.1. Case studies on the application of AI in finance	26
5.1.1. Banking	26
5.1.2. Trading	29
5.2. Case studies on the application of AI in transportation	32
Case study 1: AI for traffic monitoring and flow prediction	34
Case study 2: AI techniques for vehicle crash prediction	35
5.2.1. Application of AI technologies to traffic crash prediction in practice	37
6. Critical analyses of the current state of AI in banking, trading and transportation	- 47 -
7. Prospects in applied AI	- 53 -
8. Conclusion	- 58 -

1. Introduction

Technology has geometrically progressed during the last few decades, and subsequently artificial intelligence (AI) research reached new peaks. Today, AI plays a sizeable role across many economic sectors throughout the world. Furthermore, AI seems to be here to stay which is evident from the mass digitization of the world. Traditional operations are less able to survive, let alone thrive, in this global market where efficiency of scope and scale is achieved through state-of-the-art technology. With that in mind I set forth to research modern AI applications in business as well as the chronological steps which lead to its current state. My paper is divided into six distinct parts. The first part is concerned with the history of artificial intelligence. It is purposefully called ‘A history of AI’ because it is not the all-encompassing, condensed history, but my selection of, arguably, the most important points in time. Moving on with the second part where I describe the two main types of AI with its subsequent subsections. There I will present our current feats in the field and, hopefully, the future forms it will take. In the third section my attention falls on the ways AI is achieved at the present moment; this includes the two main ways which are Machine Learning and Deep Learning. Furthermore, in the fourth section I will describe some examples of AI in business practice. My attention is mainly focused on examples from the fields of banking and trading in the financial sector, and transportation. I chose the two fields, in part, due to my personal interests, and due to the gravity of AI developments in the fields. In the following, fifth, section I present a critical commentary of the selected case studies and fields in general. In the last section I concerned myself with a commentary of the future prospects in applied AI. In the conclusion I summarized the key points from this paper.

2. A history of artificial intelligence

The beginnings of AI thought span from deep in the past and across fields of science. Philosophers such as R. Descartes or G. W. Leibnitz imagined mechanical men and mechanical reasoning devices respectively. First instances of human complacency had their beginnings when the famous mathematician and philosopher B. Pascal created the mechanical calculator called “Pascaline” in 1642. *“It could only do addition and subtraction, with numbers being entered by manipulating its dials”*¹. Later on, science fiction relied on the possibility of intelligent machines which appealed to the fantasy of intelligent non-humans.

Furthermore, mechanical wonders continued to provoke modern researchers’ curiosity. Perhaps the best-known example of such seemingly autonomous wonder is “The Turk” by W. von Kempelen in 1769. This automaton was a masterpiece of engineering at the time as it featured an intricate array of gears, levers and pulleys. *“Placed on the top of the cabinet was a chessboard. The front of the cabinet consisted of three doors, an opening, and a drawer which could be opened to reveal a red and white ivory chess set”*². Although this machine didn’t actually utilize artificial intelligence, it was thought to have done so by its opponents. Its opponents spanned from Napoleon Bonaparte to Benjamin Franklin, almost all of who lost against the Turk. It is interesting to see just how precise the hand movements were in addition to the scope of movement spanning from its head and eyes to its hands³. At the time, as is partly today, chess was considered to be closely related to the level of intelligence, so seeing a moving mechanical chess player outsmart its prominent opponents sparked a number of theories on its modus operandi, and a bewildered amazement in the process.

The recent history of AI is also multidisciplinary. Notable people with different backgrounds left their marks on its history, including the following people, as Bruce Buchanan wrote in his journal article: *“The inspiration of modern AI thought came from people working in engineering (such as Norbert Wiener’s work on cybernetics, which includes feedback and control), biology (for example, W. Ross Ashby and Warren McCulloch and Walter Pitts’s work on neural networks in simple organisms), experimental psychology (see Newell and Simon [1972]), communication theory (for example, Claude Shannon’s theoretical work), game theory (notably by John Von Neumann and Oskar Morgenstern), mathematics and statistics*

¹ Pascaline – Encyclopedia Britannica. Retrieved June/July, 2019, from <https://www.britannica.com/technology/Pascaline>

² The Turk. Retrieved June/July 2019, from <https://interestingengineering.com/the-turk-fake-automaton-chess-player>

³ The Chess Turk explained. Retrieved June/July, 2019, from <https://youtu.be/0DbJUTsUwZE>

(for example, Irving J. Good), logic and philosophy (for example, Alan Turing, Alonzo Church, and Carl Hempel), and linguistics (such as Noam Chomsky's work on grammar).⁴ However, it wasn't until the furthest half of the 20th century that researchers had enough computing power and programming languages to conduct experiments on the realization of such visions.

A major turning point in AI history was marked by the 1950s paper in the philosophy journal *Mind* where Alan Turing crystalized the idea of programming an intelligent computing device, eventually leading to the imitation game known as Turing's test. In layman terms Turing's test is an imitation game where a human being and a computer are interrogated in such conditions that the interrogator doesn't not know which one is which. Communication is performed over textual messages and if the interrogator does not manage to distinguish them by questioning, the computer would be deemed intelligent⁵. Turing's arguments relied on our own propensity to judge intelligence based on communication capabilities.

In 1956, the work of Allen Newell, J. C. Shaw and Herb Simon was presented at the landmark conference on artificial intelligence which took place in Dartmouth. That conference might as well have engraved the initials "AI" into marble as artificial intelligence got its name then and there. Their presentation revolved around the Logic Theorist (LT) program which startled the world with as it could invent proofs of logic theorems. This feat certainly required the creation and application of programming artificial intelligence as well as creativity and it was deemed as the first program which utilized artificial intelligence. The program was deliberately engineered to mimic the problem-solving skills of humans and it was based on the system of Principia mathematica by A. N. Whitehead and B. Russell⁶. Finally, LT could prove theorems just as well as a talented mathematician which was an astounding success.

Another important example from that era is the checker-playing program by Arthur Samuel in 1956. The program was run on an IBM 701 computer and in 1962 a master checkers player lost a game against its mechanical opponent, however he managed to win the subsequent games⁷. Although simple, the program was inspiring as it learned from its human, and computer opponents. However, computing power and programming languages were still very limited. In

⁴ Buchanan, B. G. (2006). A (Very) Brief History of Artificial Intelligence. *AI Magazine*, 26(4), 56. Retrieved June/July, 2019.

⁵ Alan Turing's scrapbook. Retrieved June/July 2019, from <https://www.turing.org.uk/scrapbook/test.html>

⁶ The development of the first artificial intelligence program. Retrieved June/July 2019, from: <http://www.historyofinformation.com/detail.php?id=742>

⁷ IBM icons of progress. Retrieved June/July 2019, from <https://www.ibm.com/ibm/history/ibm100/us/en/icons/ibm700series/impacts/>

the 1950s and 1960s, some new programming languages such as Lisp, POP and IPL blew more wind in the sails of AI research, but the motionless nature and omnipresent clumsiness of early operating systems as well as the sheer size of programming devices still posed a major problem.

Other examples in the subsequent decade include T. Evans's 1963 thesis on solving analogy problems similar to the ones given on standardized IQ tests, J. Slagle's dissertation program which used heuristics to solve integration problems from freshmen calculus, D. Waterman's 1970 dissertation where he used a production system to play draw poker and another program which learned how to play better.

Meanwhile, two major approaches to AI emerged. Rule-based approach and the learning approach. Proponents of the rule-based approach, which was also called symbolic or expert systems approach, made an effort to teach computers to think according to preset rules based on logic. In a simplified manner, these logical rules are coded in the form of if-then-else, and this approach has worked well for simple games with relatively few decisions. The disadvantage of this system is its reliance on the knowledge of a human expert in a very specialized domain. Therefore, it fails to deliver optimal performance when the scope of possible combinations of choices expands. Due to its reliance on human knowledge, it is sometimes referred to as fake AI and the scientific community is divided on its potential. Maintaining these systems is cumbersome and expensive, and its scope of application is limited due to the inability of expanding its base of knowledge without setting some contradicting rules. On the other hand, the learning or Artificial Neural Network (ANN) approach took to reconstructing the human brain instead of teaching the program human logic. This approach had monumentally more success when it comes to its application in practice. The subsequent ability of the machine to learn, resulted in adaptive intelligence, meaning that knowledge can be altered and rejected as new knowledge is accumulated⁸.

Therefore, engineers created intricate webs of artificial neurons which are fed massive amounts of data such as photos, chess games, go games, sounds, etc. letting the networks learn to identify patterns in the data. The differences between the two approaches can best be portrayed with the modus operandi of an image recognition task. Let's suppose that both methods are used to recognize pictures of cats. The rule-based approach would be to train the algorithm by inputting rules describing a cat. If the image portrays two triangular shapes on top of a circular

⁸ Paragraph based on Tricentis - Artificial intelligence software testing. (n.d.). Retrieved June/July, 2019, from <https://www.tricentis.com/artificial-intelligence-software-testing/ai-approaches-rule-based-testing-vs-learning/#>

shape, the object in the image is probably a cat. On the other hand, the learning approach operates by feeding millions of photos labeled “cat” or “no cat”, letting the program decide which feature are consistently seen across the images.

The neural network approach excels in environments where differences between observed objects hinder logical approaches. This learning approach was very prominent in the beginnings of AI thought during the 1950s and 1960s, it also delivered some impressive results. Yet, 1969 marked the end of an optimistic era for the neural network approach when a group of rule-based researchers convinced others that the neural network approach was very limited in use, in addition to being unreliable. This event plunged AI research into the first of many winters. Mid 1980s sparked a new fire in the implementation of AI with use of the Hidden Markov Model technique, however, it wasn’t long lived as most of 1990s were marked by another AI winter.

Still, near the end of the nineties, IBM’s computer named Deep Blue rekindled the lost vigor in using chess as the game of choice for displaying the full range of artificially intelligent brainpower. A brief history of Deep Blue is best summarized by IBM in their “icons of progress” work: *“In 1985, a graduate student at Carnegie Mellon University, Feng-hsiung Hsu, began working on his dissertation project: a chess playing machine he called ChipTest. A classmate of his, Murray Campbell, worked on the project, too, and in 1989, both were hired to work at IBM Research. There, they continued their work with the help of other computer scientists, including Joe Hoane, Jerry Brody and C. J. Tan. The team named the project Deep Blue. The human chess champion won in 1996 against an earlier version of Deep Blue; the 1997 match was billed as a “rematch.”⁹”* The famous match of 1997 was held at the Equitable Center in New York. Millions of people watched the broadcast glued to their small screens with uncertain expectations regarding the outcome of the match. The first match was won by the chess master, while the second one went to Deep Blue. Three other matches were held and all three ended with a draw, but the last one was claimed by IBM. The story made headlines as a precedent for future dominations of machine over man. In later years Deep Blue served a number of uses, ranging from playing other strategic games to solving complex programs. Its architecture was also used in financial modeling, including risk analysis and data mining, as well as in pharmaceutical uses and biological research. Ultimately, Deep Blue was retired in the Smithsonian museum. However, its legacy lives on through IBM’s latest computer named

⁹ IBM Deep Blue. Retrieved June/July 2019, from:
<https://www.ibm.com/ibm/history/ibm100/us/en/icons/deepblue/>

“Watson” which helps humans detect cancers in medicine, solve simple legal cases as well as perform deep financial analyses quicker and better than its human counterparts.

What then, has happened recently that reignited the lost flame of AI research and implementation, and why does it seem to be omnipresent, both now and seemingly onwards? Well one explanation can be seen in the requirements of AI. Such requirements are twofold. First, AI requires a highly developed, reliable and fast computing power, and second, large amounts of data. The data is used to train the algorithms which is done by feeding them monumental quantities of specific information. This synergy allows for a quick analysis of vast amounts of data. Both of these requirements were in scarce during the last century. Through the years, Moore’s law held true and the development of more powerful processing power has created an exponential leap in hardware capabilities. Additionally, the internet lead to a burst of rich amounts of data, from pictures and videos to purchases. The end result allowed researchers to make use of artificial neural networks with relatively cheap computing power and a plethora of interesting data.

However, ANNs were not yet accurate enough to provide solutions to highly complex problems due to a lack of deeper neural layers. Another big break for AI occurred in the mid 2000s when Geoffrey Hinton discovered how to add and train neurons to the neural networks. This marked the birth of deep learning. These new neural networks were substantially better than the old ones in a variety of tasks. Furthermore, Hinton and his team annihilated the competition in 2012 when they contested in an international computer vision competition.

In recent years, Google made significant breakthroughs in AI. Starting with AlphaGo which beat the best human Go player in the world by a great margin and its successor the new and improved version which beat the old one 100 to 0. Through its feats in machine translation, all the way to its Assistant which is omnipresent in Android devices and Google’s own Home devices. Seemingly every tech company is investing in AI and this trend shows no signs of stopping.

AI today stands on the foundations all these people, and many more, laid across decades. Today, deep learning allows computers to make trading decisions, analyze histological and even drive cars, far better than any human could. Yet, todays AI requires a plethora of relevant data, a sturdy algorithm and a narrow domain with a very concrete goal in mind to work

properly. That is why machine learning, as well as deep learning is known as narrow AI. This is still a far cry from all-purpose general AI technology¹⁰.

¹⁰ LEE, K. (2019). AI SUPERPOWERS: China, silicon valley, and the new world order. S.I.: MARINER BOOKS.

3. Types of artificial intelligence

When it comes to the classification of artificial intelligence types, views become differentiated. Some state that there are seven types of AI, others classify them into two distinct groups. For the purposes of this work I opted for a type one and type two classification. Such AI thought distinguishes two main types of AI based on their functionality. The first type consists of Artificial Narrow Intelligence (ANI), Artificial General Intelligence (AGI) and Artificial Superintelligence (ASI), while the second encompasses reactive machines, limited memory, theory of mind and self-aware AI. The following discourses are compilations of several articles which will hopefully convey an accurate description^{11 12 13 14 15}.




Table 1 Types of artificial intelligence according to Kumar, C.

3.1. Type I:

3.1.1. Artificial Narrow Intelligence (ANI)

Artificial Narrow Intelligence is also called weak artificial intelligence and it focuses on one very specific, narrow task. All existing examples of AI to date fall into this category, including

¹¹ Hintze, A. Understanding the four types of AI (2016). Retrieved from: <https://theconversation.com/understanding-the-four-types-of-ai-from-reactive-robots-to-self-aware-beings-67616>. Retrieved June/July 2019

¹² Reynoso, R. (2019). 4 Main types of artificial intelligence. Retrieved from: <https://learn.g2.com/types-of-artificial-intelligence>. Retrieved June/July 2019

¹³ Joshi, N. 7 Types of artificial intelligence. Retrieved from: <https://www.forbes.com/sites/cognitiveworld/2019/06/19/7-types-of-artificial-intelligence/#96809db233ee>. Retrieved June/July 2019

¹⁴ Kumar, C. (2018, August 31). Artificial Intelligence: Definition, Types, Examples, Technologies. Retrieved June/July, 2019, from <https://medium.com/@chethankumargn/artificial-intelligence-definition-types-examples-technologies-962ea75c7b9b>

¹⁵ Jajal, T.D. (2018). Distinguishing between Narrow AI, General AI, and Super AI. <https://medium.com/@tjajal/distinguishing-between-narrow-ai-general-ai-and-super-ai-a4bc44172e22>. Retrieved June/July 2019

the most complicated and capable machines which employ deep learning and machine learning. Due to their narrow focus, they can be used to perform solely those tasks. However, they are able to operate autonomously within that field, and much better than their human counterparts. As seen later in the text regarding the second type of AI, modern examples include reactive machines, limited memory machines, and rudimentary examples of theory of mind machines. Theory of mind machines are not yet existent, but the selected examples incorporate certain characteristic features.

From our pocket companion “Siri” through Google’s AlphaGo and coming to the most modern examples of autonomous vehicles, we can observe examples of narrow AI. They attend to tasks in real time by subsequently including learned information from a specific data-set. Unlike General or Superintelligent AI, Narrow AI is not conscious or sentient. This is best observed in the example of Siri. When we ask her abstract questions such as the meaning of life or with regards to a personal problem, her answer is vague and strictly defined as on the web. Questions about the weather or calendar events are easily answered and fall within her intelligence domain by synthesizing the available data.

This is not said to undermine the marvelous results Narrow AI performed in the recent years and the monumental feats of human innovation and intelligence required to create such machines. Such narrow intelligence systems are able to process data and achieve certain goals significantly quicker than humans, as well as with more accuracy in the majority of cases. These systems are crucial for improving the quality of life and advancing human kind. A good example is IBM’s Watson. IBM Watson is used in a variety of fields and for different, specialized, tasks. For example, in healthcare it utilizes optical image recognition to detect cancer and keeps track of mountains of relevant data. Human fatigue errors are avoided¹⁶ and medical costs are decreasing.

In addition, narrow AI helps us relieve the mundane, tedious, and repetitive tasks. An example of such a program is “Sighthound”, previously “Vitamin D”. It uses AI to detect the brand and model of cars in the street as well as their license plate numbers which allows easy tracking and filtering. This search for a proverbial needle in the haystack is brought down to entering filters in a program and seeing instant results¹⁷. In addition to recognizing vehicles, the program

¹⁶ Artificial Intelligence in medicine. Retrieved from: <https://www.ibm.com/watson-health/learn/artificial-intelligence-medicine>. Retrieved June/July 2019

¹⁷ Sighthound. Retrieved from: <https://www.sighthound.com/products/sighthound-video>. Retrieved June/July 2019

can easily identify people or even recognize them by their facial characteristics. Additionally, the program can identify the age, sex, ethnicity and mood of a person based on their facial features. Another key functionality is computer redaction which removes personally identifiable information automatically, such as faces, license plate numbers, or any manually identified data.




Figure 1 Automobile recognition using Sighthound




Figure 2 Age, gender, race, and mood recognition in humans using Sighthound

3.1.2. Artificial General Intelligence (AGI)

Machines possessing Artificial General Intelligence have not yet been created, but they would embody the ability to learn, perceive, understand and function like a human being. Such systems would be able to independently build multiple competencies and form connections across different domains. This could make machines just as capable as humans with the addition of eliminating human fatigue and error. These machines would also be sentient, conscious and driven by emotion. Examples of this kind of AI are limited to the human imagination in sci-fi movies.

While current AI machines already process information faster than any human can, our ability to strategize, think abstractly and tap our deep thoughts and memories in order to create abstract ideas is still impossible to replicate in machine automation. Similarly, AGI is expected to have reasoning, problem solving, planning and learning capabilities as well as their own imagination and creativity. Machines employing AGI would be required to pass the Turing Test which, briefly explained, means that if a person is in textual correspondence with another entity on the other side of the line, the human will not be able to differentiate an AGI machine from another human.

3.1.3. Artificial Superintelligence (ASI)

Perhaps the best definition of Artificial Superintelligence (ASI) is given by professor Bostrom from the University of Oxford as he provides a simple, yet precise definition of the term. “Superintelligence” is taken to be an intellect which is much smarter than the best human brains in practically every field, spanning from scientific creativity, through general wisdom and even social skills. This leaves open how such a superintelligence should be implemented; It could be a digital computer, a cluster of networked computers, culture cortical tissue or some other solution. It also leaves open the question of its sentience and subjective experiences¹⁸.

The development of ASI is considered to be the pinnacle of AI research as it would become by far the most capable and advanced form of intelligence on Earth. It would be able to replicate the complex multi-disciplinary intelligence of human beings. Furthermore, it would eventually be exceedingly better at everything it performs due to a plethora of memory, faster data processing and decision-making capabilities. Development of AGI and ASI creates the possibility of a phenomenon called “technological singularity”. Meaning if the machines become exponentially better than human at everything they do, and their evolution is progressive, what would people do for work? Additionally, if such machines have sentience as well as superiority over humans, what is stopping them from promoting their own growth and survival over ours?¹⁹ However, if futurist Ray Kurzweil is right, we will be able to coexist with AI in a world where such machines only reinforce human capabilities²⁰.

¹⁸ Bostrom, N. (2006). How long before superintelligence? *Linguistic and Philosophical Investigations*, Vol. 5, No. 1, pp. 11-13. Retrieved June/July 2019.

¹⁹ Pandya, J. (2019). The troubling trajectory of technological singularity. Retrieved from: <https://www.forbes.com/sites/cognitiveworld/2019/02/10/the-troubling-trajectory-of-technological-singularity/#7038d3496711>. Retrieved: June/July 2019

²⁰ Kurzweil, R. (2017). AI Will Not Displace Humans, It’s Going to Enhance Us. Retrieved from: <https://futurism.com/ray-kurzweil-ai-displace-humans-going-enhance>. Retrieved June/July 2019

Kurzweil also predicts that we would be able to multiply our intelligence billionfold by linking our neocortex wirelessly to a synthetic one in the cloud by 2045. This would also allow wireless connections among humans. Today it is hard to imagine how our world will look like in the future, especially if ASI becomes a reality. However, such hard-pressing questions are still not relevant as the current state of AI is, although very impressive and immensely helpful, in its rudimentary phase. The fact that we are still scratching the surface of AI leaves much to the imagination and loads of excitement for our future progress.

3.2. Type II

3.2.1. Reactive machines

Moving on with the second type of classification, the first concept are reactive machines. They are the most basic types of artificial intelligence systems, meaning that they cannot form memories nor can they utilize past experiences when it comes to present tasks and decisions. Hence their name “reactive” as they only react to currently existing situations. Such machines have no concept of their environment so their application is limited to the simple task they were programmed for. Their specific characteristic is that no matter where or when they are used, their operation will always be as programmed, without growth or change.

An example of such reactive machines is the aforementioned IBM Deep Blue and its chess performance. It was programmed to identify a chess board and its pieces, understanding the pieces’ function. This allowed it to make predictions about its moves and the possible moves its opponent could make. Deep Blue became the first computerized program to have defeated a human. However, it has no concept of past or future and only comprehends and acts according to its programming, solely on the chessboard.

Furthermore, Google’s AlphaGo program based on the artificial neural network approach, also performs a limited scope of operations with no cognition about its environment. Although its system beats even the best Go players in the world, its specific scope of application is limited to playing such game. However, one important area where AlphaGo excels is that it learns from its past plays and applies them to its subsequent matches. On the other hand, AlphaGo is not an example of reactive machines in its entirety as the neural network approach of combining programmed data and learned experimental data falls within limited memory type. Still, even though such machines cannot interactively participate in the general environment as we image them to do at some, not so distant, point in the future, they are reliable. Reliability is not something to be frowned upon in certain AI systems as it ensures a trustworthy application in

practice. A well-programmed autonomous vehicle will behave safer than if it learns from human drivers and their follies.

3.2.2. Limited memory

Limited memory systems have the ability to use historical data to create informed future decisions in conjunction with pre-programmed data. Nearly all existing applications of AI today are based on the limited memory concept. Deep learning is a staple of today's AI endeavors and it operates by feeding the computer giant quantities of information which is then analyzed and implemented in creating knowledge. For example, image recognition AI is trained by being shown thousands of labeled pictures which allows the program to label the objects it scans. The subsequent images the AI encounters are then labeled by the program based on its "learning experience" with increasing accuracy.

These past experiences are not stored for a long time, they are fleeting inputs of information. An example of limited memory machines in practice are autonomous vehicles. To observe and understand how to effectively drive among human-operated vehicles, autonomous vehicles read their environment to detect patterns and changes in external factors so it can adapt to them. Such vehicles are able to keep track of cars in its line of vision as well as the cyclists and pedestrians. Previously, such feats took up to 100 seconds, but with the recent improvements in hardware and software this figure has dramatically decreased.

3.2.3. Theory of mind

Theory of mind is the dividing point between the machines we currently use and concepts of the future. The definition states that: *"An individual has a theory of mind if he imputes mental states to himself, and others. A system of inferences of this kind is properly viewed as a theory because such states are not directly observable, and the system can be used to make predictions about the behavior of others."*²¹ In other words, such machines would have the ability to represent mental states of humans which includes their beliefs, intentions and desires.

This was, and still is, crucial in the interactions of humans and the creation of societies. Without the ability to understand each other's feelings and intentions as well as utilizing each other's knowledge of the environment, humanity as we know it would not exist. If an AI machine is ever to walk among us, it will have to know that each human requires a different approach and has its own intentions and motives. Additionally, such information would have to be retained

²¹ Premack D., Woodruff G. Does the chimpanzee have a theory of mind? (1978). THE BEHAVIORAL AND BRAIN SCIENCES, 4, 515. Retrieved June/July, 2019.

and improved upon in subsequent interactions. Creating such machines also requires prompt responses to the rapid shifts of behavior in humans.

Researchers have also created their own iterations of theory of mind in practice and the most recent one was created in 2018 by a group of Google's engineers in their work called "Machine Theory of Mind²²". They designed a Theory of Mind neural network called ToMnet. It uses meta-learning to build models of agents it encounters. Furthermore, it can recognize that others can hold false beliefs about the world, the classic "Sally-Anne" test. The authors distinguished general theory of mind and agent-specific theory of mind. The learned weights of the network encapsulating predictions of common behavior of all agents and specific formed observations about a single agent at test time respectively. Their research is only the starting point in AI evolution based on theory of mind, but it is a very inspiring one. In the future, ToMnet should be able to introduce gentle inductive biases in its model of agents, and know how to draw from their own experiences in order to inform models of others.

Two popular examples which include certain elements from the theory of mind include Kismet and Sophia, formed in 2000 and 2016 respectively. Kismet is an expressive sociable humanoid robot developed by Professor Cynthia Breazeal and it is capable of recognizing human facial signals as well as replicating the said emotions by moving its lips, ears, eyes, eyebrows and eyelids. It is equipped with cameras inside its "eyes", microphones on each side of the head, speech synthesizer and several motors controlling the tilt and orientation of its head, and the movement of its lips, ears, eyes, eyebrows and eyelids.

Its system architecture consists of six subsystems. A low-level feature extraction system, a high-level perception system, an attention system, motivation system, behavior system, and motor system. The low-level feature extraction system extracts sensor-based features from its environment, while the high-level perceptual system condenses these features into percepts that influence its behavior, motivation, and motor processes. The attention system determines what is the most prominent and relevant stimulus of the environment at any time in order to adjust its behavior. The motivation system regulates the robot's state of "wellbeing" in the form of homeostatic regulation processes. The behavior system's task is to arbitrate between competing behaviors. The winning behavior defines the current task of the robot. It has many behaviors

²² Rabinowitz, N. C., Perbet, F., Song, H. F., Zhang, C., Eslami, S. A., & Botvinick, M. (2018). Machine Theory of Mind. (1-21). Retrieved June/July, 2019.

in its repertoire, and several motivations, so its goals differ over time. The motor system carries out the correct actions goals via the output modalities to achieve its goals²³.

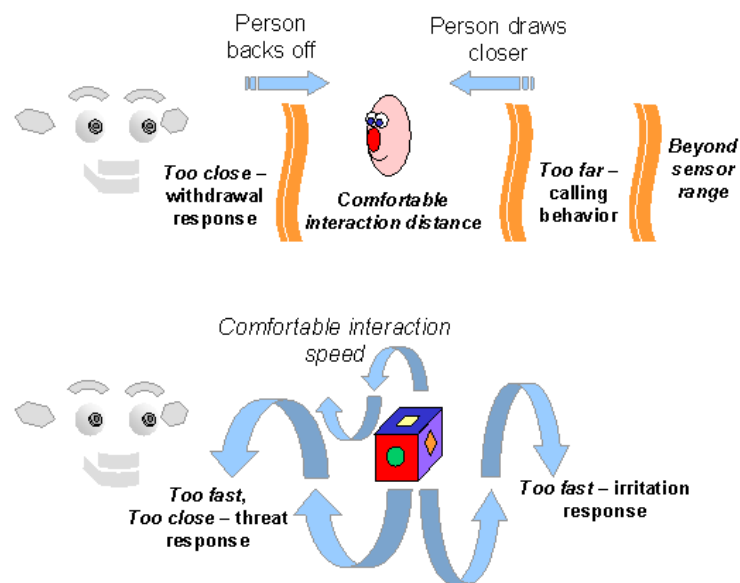


Figure 3 Kismet's socializing responses (AI MIT)

On the other hand, Sophia the robot is an advanced humanoid bot created by Hanson Robotics. Her physical likeness to a human being and her advanced image recognition allows her to interact with humans by responding with appropriate facial expressions. She appeared in a plethora of conventions and TV shows. In 2019 a TV surreality show starring Sophia will be produced, which will focus on Sophia's life, adventures, experiences and her evolution into a super intelligent, benevolent being²⁴. Sophia can walk over various terrain; she can recognize human faces and emotional expressions as well as various hand gestures. It is interesting to see how Sophia's voice and responses changed over the years through her interviews. The most recent appearances seem very impressive as her movements and responses appear very fluid and on queue. Aside from Sophia, Hanson Robotics also created a full-size humanoid robot called Han, as well as Little Sophia which is, at the time of writing, being offered on the crowdfunding campaign Indiegogo. However, Sophia isn't fully autonomous, she can operate in full AI mode, but usually AI is intermingled with human generated responses. Therefore, she is a form of "hybrid human- AI intelligence."

²³ AI MIT (2000). Retrieved from: <http://www.ai.mit.edu/projects/sociable/regulating-interaction.html>. Retrieved June/July 2019

²⁴ Hanson Robotics „Being Sophia“. Retrieved from: <https://www.hansonrobotics.com/being-sophia/>. Retrieved June/July 2019

Additionally, the creator of Sophia, David Hanson, has a history of making incredibly realistic figures which he honed when he worked as an Imagineer in a Disney theme park²⁵. It is exactly the realism of Sophia's sculpting and animatronics that makes her so special and interesting, the underlying AI is not very innovative, as is the entirety of her features. When dissected, she offers facial recognition capabilities and a chatbot engine combined with mechanism for rudimentary walking. The problem is that her answers seem to be scripted more often than not, where keywords trigger a response. Therein lies the problem, Sophia is an excellent show robot to showcase the current state of AI capabilities in a humanoid form, however she is not yet an incarnation of theory of mind, let alone self-aware AI.

3.2.4. Self-aware

Finally, the last step in the evolution of AI is considered to be a sentient, self-aware machine. Currently, this stage is only hypothetical and we are decades, if not centuries away from such technology. Such a machine would function as a human being with an evolved computing system akin to a human brain, if not more advanced, with underlying self-awareness. Mastering the creation of such a machine is, and always will be, the final goal of AI research. Such AI will be able to understand and evoke emotions in their counterparties, in addition to having their own emotions, needs, desires and beliefs. The peril of such AI is that thanks to their reasoning capacities, they will have their own self-interest as well as the possible need for self-preservation, possibly at the expense of their human counterparts.

Another domain which causes significant disputes in the AI community is how can such sentience be objectively measured and determined. Today's AI systems already have a wide repertoire of available answers on what consciousness is and it can be easily programmed into the machine's vocabulary. Additionally, the AI's ability to research vast quantities of data come in handy for filling it up with scientific literature. Therefore, a scientist could program the AI system to respond to certain keywords and questions with the programmed answers or compilations of such. This could give off an appearance of sentience without the underlying ability. Such research was performed by Professor Bishop in the most recent of such studies with inconclusive results regarding the measurement method²⁶.

²⁵ Sharkley N. A show robot or dangerous platform to mislead? (2018). Retrieved from: <https://www.forbes.com/sites/noelsharkey/2018/11/17/mama-mia-its-sophia-a-show-robot-or-dangerous-platform-to-mislead/#4e587af7ac99>. Retrieved June/July 2019

²⁶ Bishop, J.M. (2018). Is anyone home? A way to find out if AI has become self-aware. *Frontiers in robotics and AI*. Retrieved from: <https://www.frontiersin.org/articles/10.3389/frobt.2018.00017/full#B1>. Retrieved June/July 2019.

In conclusion, we have begun exploring the Theory of Mind stage of AI thought and self-aware computing seems to lie distant at the horizon of our discoveries. However, with the rapid pace of technology development I am sure that it is closer than we imagine and we have to be ready to accept and integrate it into our society in a mutually beneficial way for both us as humans and our mechanical counterparts. Moving on to ways of achieving AI, outlined below are some of the methods which are used in today's AI feats.

4. Ways of achieving AI

4.1. Machine Learning (ML)

The past decade proved to be very fruitful in creating and utilizing networked and mobile computing devices thanks to the prevalence of capable, portable and ubiquitous smartphones and other more powerful computing devices. Additionally, internet of things (IOT) created a significant impact, spreading to almost every device making its way into minute components of different stages in the production process. Its definition sheds more light on the concept: *“Internet of things is a computing concept that describes how physical objects are connected to the internet and are able to identify themselves to other devices through RFID and QR technology, among others.”*²⁷ Modern examples include Sony and its digitalized gaming consoles and TVs, wirelessly controllable Nest thermostats and GE’s use of sensors in medical products among others.

With this proverbial sea of information comes the question of analyzing and using data. That is where Big Data comes in. The definition of big data is as follows: *“big data refers to advanced data analysis and quick knowledge from huge amounts of different datasets including both structured and unstructured data. This technology represents the analytical digital platform characterized by the three Vs: Volume, Variety and Velocity.”*²⁸ These massive quantities of data are then sent to the cloud. Cloud computing is a model which enables ubiquitous, convenient and on-demand data access over the network. It requires networks, servers, loads of storage, applications and services to operate. The end result is flexible, reconfigurable, efficient and affordable operation.

Today, this vast availability of data has introduced the concept of machine learning as a subcategory to and a prerequisite for AI. Machine learning implies coding computers to behave like a human brain instead of teaching them what we know. It provides computers with access to big data and allows them to extract important features in order to solve complicated problems²⁹. From this, an important question arises: how can we create computer systems which automatically improve through experience?

Well the answer to the question emerged as a predominant common denominator in the process of machine learning spanning from computer vision, through speech recognition and natural

²⁷ Spremić, M., Ph.D. (2017). Enterprise information systems in digital economy. Zagreb, Pg.7

²⁸ Ibid, Pg.7

²⁹ Abduljabbar R, Dia H, Liyanage S, Bagloee SA. Applications of Artificial Intelligence in Transport: An Overview. Sustainability. 2019; 11(1):189; Page 2

language processing to robot control. Such methods are known as **Supervised Learning Systems**. In essence, such training consists of showing examples of input-output behavior through the use of massive quantities of highly specific and relevant labeled data. The computer then finds out the differences between the examples and uses them as benchmarks in assessing further examples. This approach is used in deep learning. It utilizes billions of parameters which are then trained on large collections of images, speech samples or videos. Professors Jordan and Mitchell in their work on machine learning gave the example of using a vast number of historical credit card purchases. Some were fraudulent and others were legitimate, they were labeled accordingly. It was up to the system to learn from the examples and eventually form its own verdicts³⁰. This simple binary classification problem where the possible combinations are limited to “fraud” and “legitimate” are not the only type of labeling as there are also multilabel classifications with the inclusion of ranking problems and general structured prediction problems. Such methods are used in speech recognition with simultaneous labeling of words in sentences.

Another method is called **Unsupervised Learning** and it involves the analysis of unlabeled data with the assumption about the structural properties of data. The machine is presented with a cluster of unsorted information which it gradually sorts according to their underlying similarities, features or patterns. It can be performed by clustering the large data sample into smaller groups of similar data, or it can be analyzed via association, where the machine aims to discover rules describing large portions of data. An example of association AI in retail is the “frequently bought together” list. Such a recommendation engine uses real-time concepts such as collaborative filtering, content-based filtering or hybrid recommendation systems to analyse the specific purchases and link products together³¹. This kind of AI is ubiquitous for almost every e-tailer, click-and-mortar or even brick-and-mortar shop such as Ikea where the frequently bought together products are placed close to one another.

The third method is called **Reinforcement Learning**. It refers to goal-oriented algorithms which learn by attaining a complex objective or maximize a particular dimension, for example maximizing points won in a game. The starting point is a blank slate, similar to a child, and its actions are incentivized via the “carrot and stick” method to achieve a certain goal through a

³⁰ Jordan, M. I., & Mitchell, T. M. (2015). Machine learning: Trends, perspectives, and prospects. *Science*, 349(6245), 255-260. Retrieved June/July, 2019.

³¹ Deb, K. (2019). Retrieved from: <https://medium.com/datadriveninvestor/product-recommendation-using-association-rule-mining-for-a-grocery-store-7e7feb6cd0f9>. Retrieved June/July 2019.

model of behavior. In psychology, such actions are called reinforcement learning, hence the name. This method was utilized in combination with deep learning to train Google's AlphaGo AI system which eventually beat the world's best Go player³². Reinforcement learning is quite similar to supervised learning, but with sparse feedback.

All in all, machine learning is still a young field with a plethora of unexplored or underexplored research opportunities. Its learning methodology is based on interdisciplinary breakthroughs combining knowledge from the fields of sociology, psychology, engineering, and others to lay the foundations of artificial intelligence. However, as mentioned before, machine learning is focused on highly specific learning endeavors contrary to the broad and all-encompassing efforts achieved by the human brain. Furthermore, there are some disputes regarding the collection, use and ownership of personal data which is collected for the purpose of analysis.

In recent years a plethora of new user data is being collected, processed and stored by the some of the titans of information technology for economic profit, in fact, companies like Google and Facebook thrive on selling user data to fund its free products. The question is how much is too much shared personal data? A good example of users benefiting from the incorporation of different technologies in delivering a message is demonstrated in the following research paper: *"By combining location data from online sources (e.g., location data from cell phones, from credit-card transactions at retail outlets, and from security cameras in public places and private buildings) with online medical data (e.g., emergency room admissions), it would be feasible today to implement a simple system which would telephone individuals immediately if a person they were in close contact with yesterday was just admitted to the emergency room with an infectious disease, alerting them to the symptoms they should watch for and precautions they should take."*³³ Such considerations lead us to believe in the positive, transformative effect of such technology in the 21st century and beyond.

4.2. Deep Learning (DL)

Deep learning is one of the types of machine learning, a technique that enables computer systems to improve themselves through experience and data. Authors Goodfellow, Bengio and Courville argue that machine learning is the only viable option to create capable AI machines which are capable of operating in a complex, real-world environment. Furthermore, deep learning, as a particular kind of machine learning seems to be the best fit between great power

³² Op. cit. LEE, K. (2019). AI SUPERPOWERS (1-6)

³³ Op. Cit. Jordan, M. I., & Mitchell, T. M. (2015). Machine learning: Trends, perspectives, and prospects, 255-260.

and enough flexibility. It is capable of representing the world in a form of nested hierarchies of concepts. Meaning that each concept is defined in relation to simpler ones which make the whole³⁴. These operations are performed through a multitude of visible and invisible computational layers. They are called such as the variables are readily observable, while the others are extracted from the subjects in an increasingly abstract fashion.

The key to its *modus operandi* lies in Stochastic Gradient Descent (SGD). It consists of presenting the input vector a couple of examples, calculating the output and errors, calculating the average gradient and adjusting the weights. This process is repeated for a large quantity of small sets of examples and it gives a noisy estimate of average gradient across all examples. To calculate the gradient of the function, researchers employed backpropagation techniques and it proved to be successful. The combination results in multilayer neural network which have several visible variables which interact with a number of hidden variables through an interrelated net which eventually creates the output sigmoid.

Deep learning is the most used mode of AI application today. Its potential is best seen on the examples of computer vision, speech recognition and machine translation which will be explained in more detail within the TensorFlow application examples. TensorFlow is one of the most widespread platforms for applying AI in practice. Other, most notable, deep learning platforms include Caffe by Berkeley AI researchers and H2O.ai. Caffe is a deep learning platform made with speed, expression and modularity as the guiding lights on the AI development path³⁵. H2O is an open source AI platform with a mission of democratizing intelligence³⁶. Coming back to TensorFlow, the following chapter is dedicated to its brief history and noteworthy examples.

4.2.1. TensorFlow

TensorFlow is an open-source deep learning system which is applied in heterogeneous environments. It uses dataflow graphs in presenting its computational operations. The nodes of its dataflow graph are mapped across devices including CPUs, GPUs and ASICs, known as Tensor Processing Units (TPUs) after which TensorFlow got its name. The benefit of this architecture is increased flexibility for the developer, allowing for novel optimizations.

³⁴ Goodfellow, I., Bengio, Y., & Courville, A. (2017). Deep learning. Cambridge, MA: MIT Press. (6-8)

³⁵ Caffe website. Retrieved from: <http://caffe.berkeleyvision.org/>

³⁶ H2O.ai website. Retrieved from: <https://www.h2o.ai/>

TensorFlow is based on many years of experience from the previous first-gen system called DistBelief³⁷.

DistBelief was developed to train very large deep networks, based on neural networks and layered graphical models. It allowed the user to define computation at each node and layer of the models well as for all the messages which were passed during upward and downward phases of computation. It also supported multiple machine computation which came in handy for delegating asks between machines which worked in unison. The largest models supported an efficient use of 32 machines with average CPU utilization of 16 cores as well as 512 CPU cores for training of a single large neural network. DistBelief performed several crucial experiments pertaining to speech and image recognition, image recognition aided ImageNet as their experiment provided an improved platform for analyzing images. Finally, DistBelief created a fertile ground for TensorFlow as it showed that new large-scale training methods use a cluster of machines to train deep networks significantly faster than a GPU³⁸.

TensorFlow is designed to be much more flexible than its predecessor while still retaining the ability to perform massive machine learning workloads. Furthermore, during the years, engineers added and upgraded the features of DistBelief for the new and improved version. For example, TensorFlow supports advanced graphing dataflow of primitive operators, meaning that represents individual mathematical operators as nodes in the graph. Next, it supports deferred execution. The process of the platform has two phases, the first one defines the program as a dataflow graph in which placeholders represent input data and variables showing the state. While the second one executes the optimized version of the program. This deferral of execution allows the program to be executed after the entire program becomes available. Finally, the common abstraction for heterogenous accelerators sets the minimum device specifications which allow for an optimal user experience when implementing the AI system in practice. The common abstractions lead to the creation of a custom Tensor Processing Unit (TPU) which allowed for Google's feats in applied AI such as: Google Street View, Inbox Smart Reply and voice search to name a few³⁹.

³⁷ Abady, M., Barham, P., Chen, J., Davis, A., et. Al. (2016) Tensorflow: A System for Large-Scale Machine Learning. 12th USENIX Symposium on OSDI. Retrieved: June/July 2019

³⁸ Paragraph based on Dean, J., Corrado, G. S., Monga, R., Chen, K., Devin, M., Le, Q. V., . . . Ng, A. Z. (2012). Large Scale Distributed Deep Networks. 1-9. Retrieved June/July, 2019.

³⁹ Paragraph based on Jouppi, N. (2016). Google Custom Search. *Google Cloud Blog*. Retrieved June/July, 2019, from <https://cloud.google.com/blog/products/gcp/google-supercharges-machine-learning-tasks-with-custom-chip>.

Engineers at Google created a new and improved, second generation TPU. The previous knowledge and software as well as hardware for machine learning was, although vast and exponentially improved over the years, still too weak to perform large scale computing. Due to this, Google's engineers created their own hardware and software which they used internally, as well as externally to greatly scale machine learning training and inference. As a result, training a large-scale translation model used to take a full day and 32 of the best GPUs currently available to train the algorithm, now it takes just one afternoon and an eighth of a TPU. Google has opened its TPUs to offer Cloud TPUs where customers can connect to virtual machines of all shapes and sizes to create their own custom platform for their specific machine learning AI needs⁴⁰.

Thanks to TensorFlow, from Fall 2016 the quality of machine translation for English-French, English-Chinese and English-Japanese language pairs, rose by a steep margin. The deep neural network trained with word embeddings turned the translation efforts from awkward and clumsy attempts to near professional standard levels⁴¹. Aside from machine translation, Google relies on TensorFlow to offer smart clipping tools for Android devices as well as image recognition through its Google Glass application. Furthermore, Google's engineers utilized AI in its Android keyboard for recognizing handwriting and turning it into typed letters⁴². Voice recognition is another notable example as Google Assistant relies heavily on voice communication. It allows the user to ask it different questions and delegate simpler tasks such as writing notes, setting calendar appointments and even booking tables in a restaurant. Google's efforts in the AI frontier are further exacerbated by the fact that they offer other users to use their TensorFlow platform and outsource their computing needs. In order to educate people on how to use the platform, Udacity, Coursera, Deeplearning.ai and Fast.ai websites created or hosted teaching materials which will train the next generations of AI users.

In the three years of its existence TensorFlow has created value for innumerable researchers, practitioners and companies who downloaded it over 41 million times⁴³. The intermodality of

⁴⁰ Paragraph based on Dean, J., & Holzle, U. (2017). Build and train machine learning models on our new Google Cloud TPUs. *Google Cloud Blog*. Retrieved June/July, 2019, from <https://blog.google/products/google-cloud/google-cloud-offer-tpus-machine-learning/>.

⁴¹ Heller, M. (2019). Deep learning explained. Retrieved from: <https://www.infoworld.com/article/3397142/deep-learning-explained.html>. Retrieved June/July 2019

⁴² Roberts, C., Leichenauer, S. (2019) Retrieved from: <https://ai.googleblog.com/search/label/TensorFlow>. Retrieved June/July 2019

⁴³ Alcober, F., Gupta, S. Retrieved from: <https://medium.com/tensorflow/recap-of-the-2019-tensorflow-dev-summit-1b5ede42da8d>. Retrieved June/July 2019

programming languages also benefits users as the line between mobile, desktop, cloud and IOT becomes very blurred. Therefore, it is not surprising that some of the biggest companies in the world rely on TensorFlow to tackle difficult problems.

Twitter-ranking tweets with TensorFlow⁴⁴

The mission of Twitter is to keep users informed with relevant content. However, the tweets are originally presented in a reverse chronological order and the community became more connected. This made it impossible for users to be kept informed about the most important posts their followees made. For that reason, Twitter partnered up with Google in creating a “Ranked Timeline” which was designed to show the most relevant posts at the top of the user’s timeline.

Ranking is performed by a relevance model which tailor fits the importance of the Tweet to a specific user. Some of the parameters which were used are: the recency of the Tweet as well as the use of media such as images or videos, real-time interactions in terms of retweets and likes, as well as the user’s past interactions with the author.

The engineers at Twitter have used another platform called Lua Torch, but it required its own ecosystem and a specialized programming language. While TensorFlow offered support for many other programming languages. Additionally, Lua required separate programming for different layers of specific “yaml” files, while TensorFlow allows modeling deep neural nets in a fast and flexible fashion.

A crucial differentiation in ranking Tweets as opposed to image classification is their sparse nature. It is truly astounding to see the promptness with which the program analyses the plethora of Tweets and their individual relevance to each specific user. It is in that regard that machine and deep learning platforms excel over manpower by an incredible margin.

GE healthcare-intelligent scanning using deep learning for MRI⁴⁵

In essence, GE healthcare used TensorFlow to train a neural network which would identify specific anatomy during a brain resonance imaging (MRI) exam. This would improve the speed and consistency of MRI scanning. A specific feature of an MRI is its ability to differentiate gray and white matter as well as the complete perfusion, diffusion and blood flow of a specific area. However, due to its advanced abilities, the operators must carefully plan the scans as the

⁴⁴ Zhuang, Y., Thiagarajan, A., Sweeney, T. (2019). Ranking Tweets with TensorFlow. Retrieved from: <https://medium.com/tensorflow/ranking-tweets-with-tensorflow-932d449b7c4>. Retrieved June/July 2019

⁴⁵ Polzin, J.A. (2019). Intelligent Scanning Using Deep Learning for MRI

quality and consistency of positioning and orientation of the MRI slices relied on the skill and experience of the operator. The process tends to be very time-consuming and difficult.

A solution to this problem revealed itself in the development of a deep learning-based framework for intelligent MRI slice placement. Thanks to the Convolutional Neural Network approach used by TensorFlow, GE healthcare managed to utilize deep learning which is able to determine plane orientations automatically in around 3.5 seconds on a high-performance CPU. The training and testing data for deep learning came from more than 1,300 subjects through various GE scanner models.

The subsequent LocalizerIQ-Net was trained using a total of 29,000 images and tested on over 700 images. The classification accuracy was an astounding 99.2%. A true test of its performance was when a patient held his head in three very different positions and the program automatically and very precisely created a consistent brain image.




Figure 4 Consistent brain image orientation in three various head rotations (Polzin, J.A.)

The end result is a 40% to 60% reduction in required time as well as fewer errors and an improved accuracy. Furthermore, thanks to its successes in MRI scans of the brain, GE healthcare plans to include knee and spine diagnostics in the future. Additionally, other manual tasks in the overall procedures are to be automated in the near future allowing for a completely automated MRI machine.

5. Application of artificial intelligence in business

5.1. Case studies on the application of AI in finance

Due to the numerical nature of finance, AI has a great base for unleashing its full potential seeing as it excels in recording, categorizing and reporting about large numbers of data. For that reason, AI technology has become highly commoditized. Some of the platforms include Google and its TensorFlow, MXNet supported by Amazon Web Services (AWS) and h2o, to name a few. H2o being a general-purpose platform on which different deep learning systems can be used such as Caffe, Torch and Theano. These aforementioned platforms are used in a plethora of industries across companies as they provide a scalable and customizable way to geometrically boost productivity and achieve more than was ever imaginable. Some companies even use their own platforms which are tailor fit for their needs.

Furthermore, today's digital world requires very low friction and immediate response. People and companies alike are seeking out ways to build stronger commerce relations which can be scaled up promptly if needed. A good example of changes due to digital influences on finance is the idea of a bank.

5.1.1. Banking

A bank used to be a physical location where individuals stored their hard-earned money to keep it safe and eventually earning some interest on savings. But today, banks are more connected to their users than ever. It is expected to be able to access all of the services a bank offers in an instant from any corner of the world with the use of mobile banking. Additionally, physical currency is being used less and less as cash payments shift to digital ones via card, phone or even fingerprint payments which require a constant digital connection between a Point Of Sale (POS) device and the user's bank. Keeping up with the trends of digitalization, most banks employ AI to help them provide accurate, fast, reliable and personalized services to their customers. In the following pages my focus will be directed at investigating the most important examples in use today.

As previously mentioned, AI's ability to scan through heaps of data makes it a perfect candidate for analyzing and extracting information from important documents. Such an example is seen at the JPMorgan Chase bank where computer engineers created a form of AI called COiN.

They call it contractual intelligence and it is able to process 12,000 credit agreements in seconds, relative to the usual 360,000 hours⁴⁶.

Lloyds bank, on the other hand, has a different approach to AI. Their approach is very interesting because it focuses on two key features which are a holistic approach and augmentation instead of pure automation. In their case, AI performs the time consuming and/or mundane tasks while learning from the operations and interactions. Employees then have instant access to specific quantities of data pertaining to their area of interest. Director of digital development at Lloyds, Marc Lien, said that the bank is about: *“pairing brilliant people in our business with increasingly smart technology to deliver great things.”*⁴⁷ The cornerstone of AI at Lloyds are chatbots. The smart assistants resolve customer queries, or if they are unable to do so, pass them along to human operators. However, the chatbots are able to see the way human operators resolved the query and they learn from that, so the scope of operation becomes greater. Furthermore, these chatbots are also deployed internally where they have access to the entire corpus of information, from client managers to telephone operators from whose interactions the AI draws information which can then be accessed by employees to draw relevant information pertaining to their specific case.

Moreover, Bank of America created the first AI driven virtual assistant intended for wide public use. The innovation is named “Erica” and is a combination of the latest AI technology, predictive analytics and natural language processing. As any other AI, Erica is able to learn client’s behaviors and adapt to them offering tailor fit solutions. Clients can communicate with her through text, voice or screen interactions as she understands intent and context of a conversation. Its knowledge base consists of 200,000 preset customer questions with the ever-expanding corpus of knowledge due to machine learning algorithms. The current scope of operation includes searching the database of past transactions, informing the customer of their bank balance and due payments, calculating credit scores, payment processing capabilities, locking and unlocking debit cards, navigating the user to the specified bank or ATM and other innovative capabilities⁴⁸. The two major benefits of AI assistants such as Erica is 0-24hr service

⁴⁶ JPMorgan Chase Annual Report 2016. Retrieved from: <https://www.jpmorganchase.com/corporate/investor-relations/document/2016-annualreport.pdf>. P.49. Retrieved June/July 2019

⁴⁷ Mortleman, J., (2017). Retrieved from: <https://www.i-cio.com/management/insight/item/banking-on-an-ai-future>. Retrieved: June/July 2019

⁴⁸ Bank of America newsroom (2018). Retrieved from: <https://newsroom.bankofamerica.com/press-releases/consumer-banking/bank-america-delivers-first-widely-available-ai-driven-virtual>. Retrieved June/July 2019

availability as well as personalization on a broad scale. Deep learning is utilized to help customers manage their daily finances in a more effective and prompter way as well as for optimizing their budget and cash flow whilst offering suggestions for saving money.

Similar case of AI utilization is seen in the case of NatWest bank and their “digital human” named Cora. Deployed in 2017, Cora is a text-based chat bot with which customers can interact by on the website’s online help page. It can give answers to more than 200 banking queries and at the time of writing (2018) it already has more than 100,000 conversations a month⁴⁹. Furthermore, the new Cora version allows customers to have a two-way verbal communication through their smart devices 24 hours a day. The hardware at the customer’s end of the line is quite rudimentary, requiring only auditory and visual sensors along with a digital input device such as a keyboard or touch-screen. This is done in an effort to expand the customer base by not posing limitations on hardware requirements. However, the crown jewel of the project is the collaboration with a New Zealand company Soul Machines, which gave Cora a super realistic appearance along with refined facial gestures. This was done in an effort of bringing AI closer to customers and hopefully gain the trust of technophobic users which prefer face-to-face communication. Additionally, such interactions aid communication with persons suffering from certain handicaps, such as blindness.

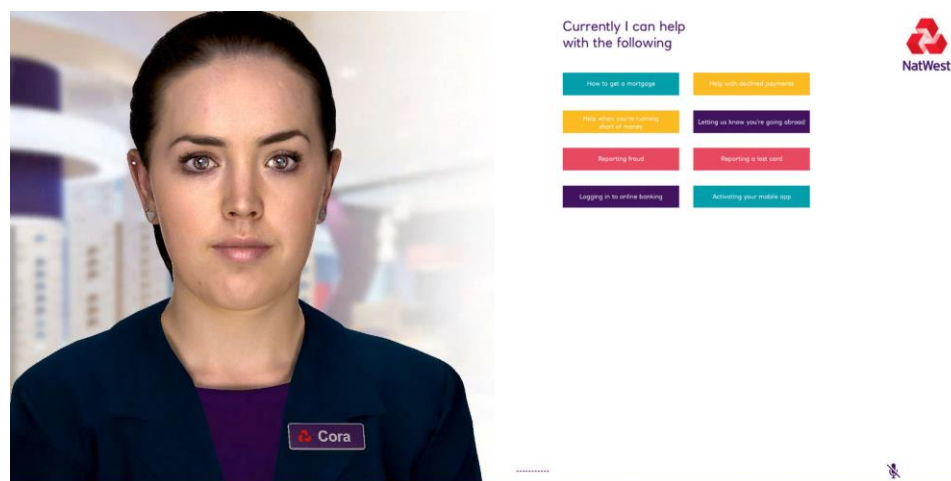


Figure 5 AI assistant Cora from NatWest bank

⁴⁹ NatWest begins testing AI driven ‘digital human’ in banking first. (2018). Retrieved from: <https://www.rbs.com/rbs/news/2018/02/natwest-begins-testing-ai-driven-digital-human-in-banking-first.html>. Retrieved June/July 2019

Case study 1: Transforming banking support with an AI-powered virtual assistant⁵⁰

Similar to the previous examples, IBM Watson put its own staple on customer service. The challenge posed by Creval Sistemi e Servizi (CSS) was to reduce the need for human effort regarding internal service-support operations. Seeing that CSS is a bank which takes pride in its superb customer service and satisfaction, this feat required meticulous attention to detail in utilizing AI. An example of their dedication to customer satisfaction is their lead time for mortgage inquiries. The bank average in Italy is 30 to 60 days, CSS managed to cut it to 15 days. Furthermore, CSS has invested in cutting-edge technology over the years, including IBM Business Process Management systems (BPM) as well as in AI with IBM's Watson Analytics technology which makes use of big data in their CRM systems. The new IBM Watson virtual assistant is a result of a likely marriage between the bank's passion for technology and their drive to provide excellent customer service. This gave birth to "Alfredo", the AI virtual assistant. Alfredo is currently the single point of contact for branch users to request help from the bank's service desk. The assistant utilizes IBM Natural Language Understanding services, IBM Watson Assistant solutions and is run on IBM's Cloud platform. CSS trained Alfredo on 14 knowledge domains of the bank. He provides around 1,000 autonomous responses without human intervention, however, when human help is needed, he directs the inquiries via skype to the appropriate service department. The service is available 0-24 hours a day. The final result was an astonishing reduction of calls to back-office personnel by 80% as well as a 40% reduction in required efforts from such personnel. Most importantly, 92% of users exclaimed that Alfred was a positive change for their bank. CEO of CSS Pizzicoli summarized the importance of Alfred in one key sentence which I believe is the staple behind most AI ventures today: *"The cognitive system can reduce human effort and eliminate or reduce tasks that are repetitive and without value; this allows our people to be focused on the very high-value customer questions. For me, that's really powerful."* Due to the resounding success of the project, CSS plans to make Alfredo available to its retail banking customers in 2019.

5.1.2. Trading

Trading is the number one sector with the greatest financial value added in AI. Fleets of machines manage millions upon millions of Dollars, Euros, Yen and other currencies in

⁵⁰ IBM Watson case study-Transforming banking support with an AI-powered virtual assistant. Retrieved from: <https://www.ibm.com/case-studies/crevalsistemieservizi>. Retrieved June/July 2019

markets all over the world. It is precisely the lack of emotion in AI that makes it a perfect candidate for managing money. Human investors tend to be slaves of emotions which gives way for mistakes. For example, investors are prone to misinterpreting the true value of companies they are fond of. As well as keeping shares of such companies for longer, even if the market shows a lack of potential for their growth, eventually leading to further losses.

However, there is always an inherent risk when handing off control to algorithms, either through some initial coding error or falsely learned data. Such occurrences are very expensive when massive amounts of money are at stake. One such example includes an algorithm gone rogue in 2012. The algorithm began buying and selling massive quantities of batched stocks. The blizzard of erratic orders in the minutes after the trading started was first noticed by NYSE officials who then reported the issue to Knight group. These kinds of situations are scary if they subsist for over a minute, five minutes are considered to be the absolute worst nightmare with irreparable consequences. This occurrence lasted a good 30 minutes. Everybody from the SEC through NYSE and the Knight group was shocked that the problem persisted that long. The problem was in the lack of a turn-off switch. This mistake cost Knight a possibly fatal 440 million dollars⁵¹. Even at the time of writing other high-speed trading groups such as Aite Group had several built-in signals with an integrated automatic kill switch that immediately stops trading.

Now all of this happened because technology changed Wall Street beyond recognition. Orders are executed in a millionth of a second and financial firms battle with their sophisticated algorithms over fractions of cents. However, with High Frequency Trading (HFT), these fractions mean millions of dollars. HFT operates by buying a stock, for example priced 1\$ and selling it for 1,0001\$⁵². Seems insignificant, but when repeated 10.000 times per second throughout the day, the profits start to build up. Such algorithms constantly move in and out of positions for tiny profits and end the trading day owning nothing. As a result of this widespread use of technology, the average stock holding period plummeted from eight years, fifty years ago, to the current average of five days. Additionally, HFT accounts for over fifty percent of US trading.

⁵¹ Ahmed, A., Protess, B. Trading program ran amok, with no 'Off' switch. (2012). Retrieved from: https://dealbook.nytimes.com/2012/08/03/trading-program-ran-amok-with-no-off-switch/?_r=0&mtref=www.entefy.com&gwh=2995327C21A7C54B7D7583F1968CD11A&gwt=pay. Retrieved June/July 2019

⁵² Baumann, N. (2013). Too fast to fail. Retrieved from: <https://www.motherjones.com/politics/2013/02/high-frequency-trading-danger-risk-wall-street/>. Retrieved June/July 2019

Another key ingredient for modern successful trading utilizing algorithms is low latency. Latency is the time it takes to execute a financial transaction over a network connection. Tech companies are launching the lowest-latency link yet, it will be positioned between Illinois and New Jersey. The 733-mile chain of microwave towers will hurdle data in 8.5 milliseconds round trip. That is because a trader with a faster internet connection can sell at higher prices and buy at lower ones because of the timing advantage. It is estimated a broker loses out to around four million dollars in revenues per millisecond if his electronic trading platform lags behind its competition.⁵³ It is for that reason modern trading companies fight over the closest spaces around the stock exchanges, or for access to the various optical routes dotted around the world. Access to such routes costs up to 300,000 USD a month, but considering the profits it leads to, such investments are pennies on a dollar.

With the following factors of ubiquitous digitalization, low latency and ever-improving software, AI software was just a matter of time. Today, AI funds are common in the investing world, and not just for HFT, although its impact on high frequency trading is far from insignificant.

A suitable start into the implementation of AI in trading is its use in high frequency trading. Starting with the basics, the stock market works on the principle of the order book. An order book contains the top few bid and ask prices for certain commodities or stocks. The individual or machine places orders which are either orders to buy a number of stocks at a specified price or current price (bid orders), or one places a order to sell a number of stocks at a certain price (ask orders). Then the orders are matched and orders go through. The problem is, by the time a human reads the market trend and places the bids, market conditions have changed. For that reason, automated algorithmic trading which utilizes AI can reap a plethora of benefits and profits. Most HFT programs are based on the Support Vector Machine (SVM) algorithm which analyses data for classification and regression analysis. It is utilized because of its nimbleness and precision.

Case study 2: Deep reinforcement learnig in high frequency trading⁵⁴

Deep learning applications are being developed with some very interesting examples already in use, however such programs are mostly not utilized in HFT, although one recent study

⁵³ Shalom, N. (2019). Retrieved from: <https://www.gigaspace.com/blog/amazon-found-every-100ms-of-latency-cost-them-1-in-sales/>. Retrieved June/July 2019

⁵⁴ Ganesh, P., Rakheja, P. (January 2019). Deep reinforcement learnig in high frequency trading. CODS-COMAD 2019. (1-6). Retrieved June/July 2019

managed to shine some light on the potential of deep learning in HFT. This research stated that the two most important challenges are: the complexity of microsecond sensitive live trading which brings forth the question of expensive computation power, and the tremendous amount of fine granularity of data which entails vast quantities of historical data. Their goal was to create a holistic deep learning program which would go beyond merely exploiting the order book's bid and ask prices. The proposed program would rely on past trading data as well as current market data which is analyzed by a deep reinforcement learning model. The model would offer its predictions on the price movements, followed by placing orders in the market according to the prediction.

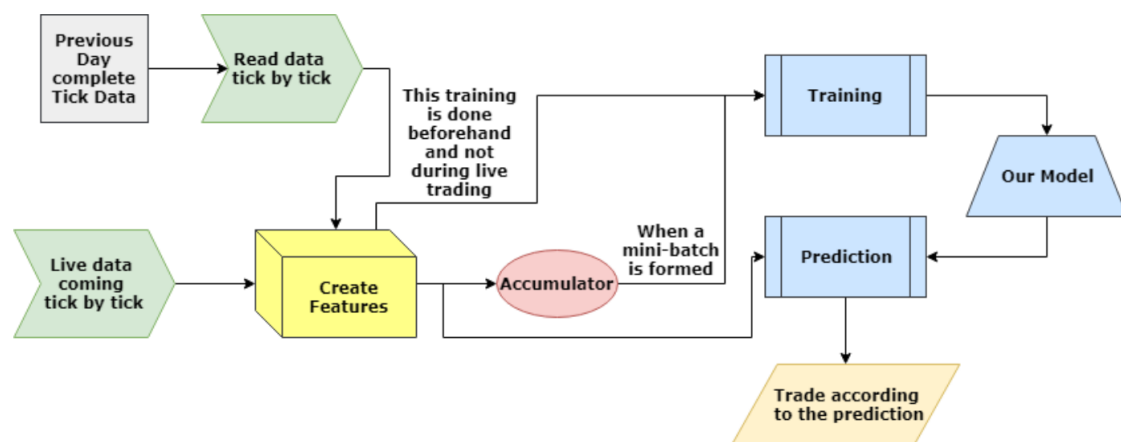


Figure 6 Pipeline model proposed by the authors (Ganesh, Rakheja 2019)

The dataset of the model consisted of 200 filtered securities with an average activity of around 70,000 ticks a day, in total, the model managed 14 million data points every single day as well as a whole week's worth of similar data to work with. In The end, the model performed with around 70% accuracy and opened the door for future research of the fit between AI and HFT.

5.2. Case studies on the application of AI in transportation

By now, artificial intelligence has a plethora of uses in transportation. Furthermore, McKinsey dubbed it the second most prominent sector with high value-added potential in its 2018 research⁵⁵. When it comes to transportation, artificial intelligence has a lot to offer. For example, AI can be used to assist drivers and improve the existing road infrastructure in order to create a base for safer transport. These systems are jointly used to gather data which is sent to the server where it is analyzed and reported back to the user's devices. Additionally, the

⁵⁵ McKinsey Global Institute. Notes from the AI Frontier. Insights from Hundreds of Use Cases; McKinsey Global Institute: San Francisco, CA, USA, 2018; p. 36

general technological trends of improvement, streamlining and digitalization of modern processes aid these quests as more and more researchers set to tackle these problems. Besides, global transportation requires even greater integration and optimization which creates a challenge to form an agile and flexible global system which will stand the test of time.

If we take logistics as an example, we can observe that AI is set to change its course by shifting the importance of a human workforce to a mechanical one. The operational requirements become vaster and more complex, creating a gap called the *knowledge accumulation gap*⁵⁶.

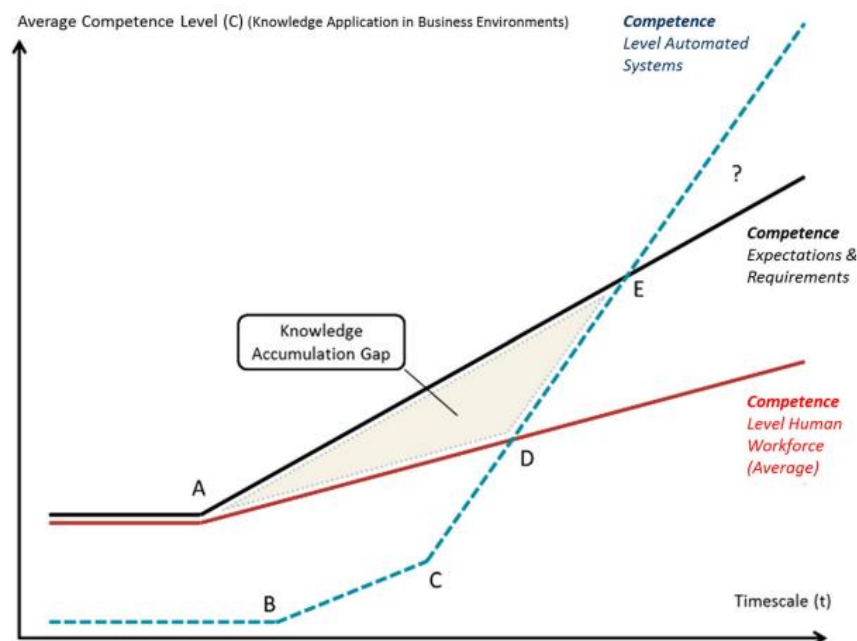


Figure 7 The knowledge accumulation gap (Klumpp 2018, 227)

This knowledge accumulation gap represents a disproportion in the competence level of an Average Human Workforce versus the required competence level through time. This gap is becoming bridged and surpassed by Automated Systems over time. The beginning of the timescale at point A is taken to be the Industrial revolution where the requirements for the workforce began their constant rising trend due to the rising expectations and complexities. These increased expectations are in part due to the required skills that had to possess such as driving, production and warehousing processes and similar, as well as the legal requirements posed by the Government regarding training in workplace safety, handling of hazardous goods and such similar knowledge. Such knowledge requires longer hours of education and employee

⁵⁶ Klumpp, M. (2018). Automation and artificial intelligence in business logistics systems: Human reactions and collaboration requirements. *International Journal of Logistics Research and Applications*, 21, 3rd ser., 224-242

exertion which is not the case for computers as they can store, categories and use vast quantities of information. Yet, computers today are reaching point D and some are very close to point E.

Depending on the sector and application of AI, the results are polarizing between superiority of man and machine. For example, artificial intelligence excels at gathering, organizing and using vast quantities of data. This feature is especially important in the study of traffic flow where AI has the possibility to reduce carbon emissions, alleviate traffic congestions and advance traffic operation efficiency. Such traffic flow predictions depend on historical and real-time data collected from sources such as radars, loops, cameras, social media and the Global Positioning System. However, the majority of traffic flow prediction systems still rely on shallow models.

Case study 1: AI for traffic monitoring and flow prediction⁵⁷

A group of researchers in California decided to create a deep architecture model which would take advantage of the rich amounts of data. Their research was performed in the stacked autoencoder (SAE) model, which was used to learn generic traffic flow features and was trained in a layerwise greedy manner. Training was the most difficult phase as the straightforward approach to training was notorious for inferior results.

Therefore, the researchers had to use an unsupervised layerwise greedy approach to pretrain the algorithm, from that point on it was just a matter of tuning the parameters. This deep architecture model was applied to data collected from the Caltrans Performance Measurement System (PeMS) database for calibration. Onwards, traffic data was collected every 30 seconds from over 15 000 individual detectors, deployed across highways in California. The collected data was aggregated in 5-minute intervals for each detector and traffic flow was collected during weekdays in January through March of 2013. The first two months were the training sets and the third was the testing set.

The results showed that the algorithm performed best in conditions of medium to high traffic flow as low traffic caused a bigger relative error. Nevertheless, this model was created with heavy traffic flow in mind. When such conditions were met, the algorithm predicted 15 minutes of flow with 93% accuracy. Moreover, its predictions 60 minutes into the future were still above a very promising 90% success rate.

⁵⁷ Paragraph based on Yisheng, L., Yanjie, D., Wenwen, K., Zhengxi, L., & Fei-Yue, W. (april 2015). Traffic Flow Prediction With Big Data: A Deep Learning Approach. *IEEE TRANSACTIONS ON INTELLIGENT TRANSPORTATION SYSTEMS*, 16(2), 1-9. Retrieved June/July, 2019.


Case study 2: AI techniques for vehicle crash prediction⁵⁸

Vehicle crash prediction is another domain in transportation where AI can help save lives. For that reason, another group of researchers decided to investigate the subject matter further. Their research revolves around the application of AI systems which would help diminish, if not eliminate, certain types of traffic accidents. The significance of this particular research lies in its attempt to encapsulate a vast array of literature related to AI based techniques for accident prediction, driver identification, unsafe driving pattern detection and the future trends of these domains.

The World Health Organization reported that the average annual fatality rate is 18 per 100 000 persons in the world. It varies geographically from country to country regarding the severity and extent of the injuries; however, the figure is still alarming. Due to the loss of precious lives and property, traffic accidents are a major and growing concern around the world. These accidents are caused by a variety of factors such as the weather conditions, attention of the driver (the use of cellphones in traffic), traffic congestions and similar. Due to the differentiated nature of causes, the methods of traffic accident prediction vary accordingly.

Selection of a specific technique for detection relies on the type, size and format of recorded data.

⁵⁸ Paragraph based on Halim, Z., Kalsoom, R., Bashir, S., & Abbas, G. (october 2016). Artificial intelligence techniques for driving safety and vehicle crash prediction. *Artificial Intelligence Review*, 46(3), 351-387. Retrieved June/July, 2019.



and their complete integration would achieve the greatest success in predicting, and subsequently preventing, traffic accidents.

- **Generation of early warnings:** At first sight it is very similar to accident protection, however, they are very different as early warning generation depends on the subsequent actions of the driver.
- **Modeling individual features of the driver:** This feature could be used for personalized autonomous driving as the individual driver style would be analyzed and modeled.
- **Accident identification:** These systems would be applied when the accident already occurs. The data would be collected from sensors in the vehicle. Its application would be to alert the emergency services about the accident.
- **Prediction of driver suitability:** Prediction would be made according to past data which would be collected, analyzed and stored. It could be used in police stations where the fit between the vehicle and driver would be determined.
- **Postmortem analysis:** Study of the circumstances and reasons which lead to the accident would allow for preventive future action. In this area machine learning (ML) would play a sizeable role.

5.2.1. Application of AI technologies to traffic crash prediction in practice

There are different techniques through which artificial intelligence is being applied in practice, some of the most notable ones are outlined below in separate research examples.

- **Genetic Algorithm (GA)** is an evolutionary approach inspired by the Darwin theory. It applies genetic operations of mutation as well as crossover to find the best solution. The process occurs in stages of populations which evolve using the rules of the problem domain. Subsequent populations will integrate the knowledge of previous populations to adapt more effectively to the environment.

The application of GA on traffic accident prediction is describe in the following paper produced by two prominent scholars. Their work is based on the use of a customized multi-objective genetic algorithm (MGOA), i.e. non-dominated sorting genetic algorithm (NSGA-II). Five crucial steps were followed in the proposed method starting with preprocessing, capturing preferences, creating training and test sets, applying NSGA-II and evaluating rules. Additionally, human preferences of users were captures,

including the weight of comprehensibility and the features of traffic accident instances which were interesting to the users. The final proposed method was used to evaluate 14,211 accidents on rural and urban roads of Teheran for the period of five years (2008-2013), though after the testing phase the number declined to 12,625 due to instances of statistical noise. As previously mentioned, GA uses generations of populations, each one greater than the previous one due to the evolutionary process. To train the program, the developers had to create several predetermined conditions in order to successfully steer AI which they successfully did. The final results showed a 4.5% increase in the accuracy metric⁵⁹.

- **Genetic Programming (GP)** genetically breeds computer programs to solve problems, meaning that it performs similarly to GA but the solutions are shown on a tree-like structure. It has two major advantages over GA, mainly in the ability to generate better solutions without an underlying perspective and in the removal of the black box effect.

Such research using Genetic Programming was done by Mexican scholars who developed a system which gathered driving data from cellphones of the drivers. The sensors and systems in modern mobile devices are a great base for such research. In order to validate the accuracy of their model, human observers had to evaluate driving performance on a scale from 1 (very unsafe) to 10 (very safe). Following tests were performed on a scale from 1 to 4 in order to have a fine and coarse dataset. GP was very suitable for this research as it uses an evolutionary search to derive small programs and models. It is also used to solve a variety of machine learning tasks with the most common being the symbolic regression model (SRM). SRM represents the relationship between input variables and dependent output variable. To complement GP, a combination of NeuroEvaluation of Argumentation Topologies algorithm (NEAT) and GP was used to create the neatGP model. This model preserved a diverse population of individuals by taking speciation techniques and standard fitness sharing. Finally, a dataset of 200 road trips was collected where human observers graded the trips, with the final scores being an average of all scores combined. Taking into account average speed, distance of the

⁵⁹ Paragraph based on Hashmienejad, S. H., & Hasheminejad, S. M. (2017). Traffic accident severity prediction using a novel multi-objective genetic algorithm. *International Journal of Crashworthiness*, 22(4), 425-440. Retrieved June/July, 2019.

trip, number of lane changes, abrupt steering and sudden stops, the program managed to successfully predict the human observer score in most cases⁶⁰.

- **Artificial Neural Network (ANN)** is used to model complex relationship between input and output, finding patterns in data. It is developed in three phases consisting of modeling, training and testing. Preparation of data and adaption of learning laws are performed in the training phase, while accuracy and performance evaluation are performed in the testing phase. In the past decade ANNs were used in a number of studies.

One such study was performed by a group of Korean researchers who used the new Two-way Probe Car System (TPCS) measuring link travel speeds in South Korea. TCPS is a means of collecting roadway condition information and transmitting it to the users. This allows on board navigation systems to make better decisions. The study was performed with taxis equipped with on-board equipment (OBE) which communicated with roadside equipment (RSE) placed at major signaled intersections and arterial roads as shown in the figure below.

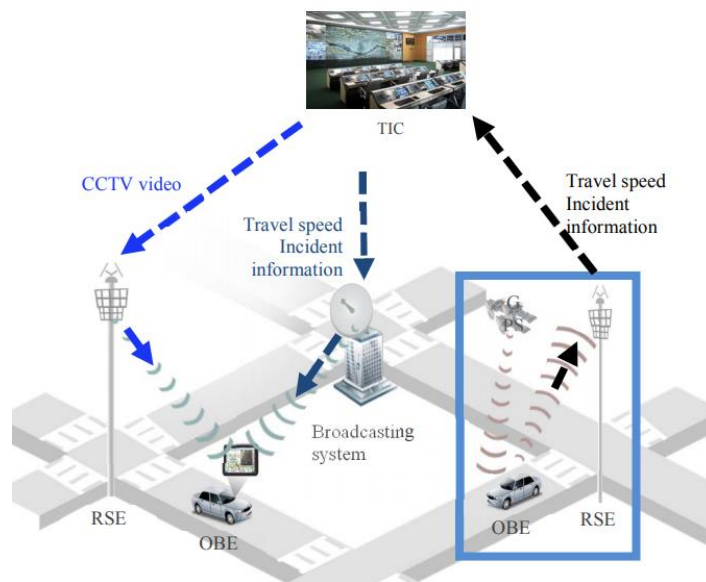


Figure 9 Schematic view of the TCPS process (Ki, et al. 2018)

⁶⁰ Paragraph based on López, R., Gurrola, L. G., Trujillo, L., Prieto, O., Ramírez, G., Posada, A., . . . Méndez, L. (2018). How Am I Driving? Using Genetic Programming to Generate Scoring Functions for Urban Driving Behavior. *Mathematical and Computational Applications*, 23(2), 1-13. Retrieved June/July, 2019.

In the paper, a new algorithm is suggested to detect traffic incidents using the TCPS system. This algorithm is based on the ANN approach as it is capable of learning, self-adaptation and fault tolerance. The ANN performs on a three-layer model where an incident is detected due to the usual creation of upstream congestion and reduced flow in the downstream direction. Creating a high velocity difference between two stations. This experiment was conducted at five sites across Seoul. Each database was split in two parts with 42 percent of randomly selected data serving as training data. A number of false incident logs was discovered in the testing phase and eradicated for the future phases. Finally, out of 40 incidents, 29 were detected, however 11 were not detected. Creating a final accuracy of 72.5% which is a marked increase from previous detection model results performed in other areas of the world⁶¹.

- **Principal Component Analysis (PCA) and Hidden Markov Model (HMM):** PCA performs as an unsupervised feature extraction model used to select smaller numbers of artificial variables within a large number of observations. They are called principal components. HMM, on the other hand, is a statistical model where process states are hidden, however, the output state is visible.

Two Indian scholars applied the PCA and HMM techniques in creating an application which analyses the odds of crashing a vehicle. This application combined user data as well as real-time driving data which was analyzed to create a judgment of user's fit for operating a vehicle. The authors decided to create their own system because the existing VEDAS and ADAS systems were too expensive to implement on a broad scale, and as modern smartphones already have an integrated camera, accelerometer and GPS, they make for a perfect computational device. User data consists of: gender, age, validity of license, potential disabilities, employment status and position, experience and vehicle type. They are supplemented by real-time driving data and are subsequently sent to the server for analysis.

⁶¹ Paragraph based on Ki, Y., Heo, N., Choi, J., Ahn, G., & Park, K. (2018). An incident detection algorithm using artificial neural networks and traffic information. *2018 Cybernetics & Informatics (K&I)*, 1-5. Retrieved June/July, 2019.

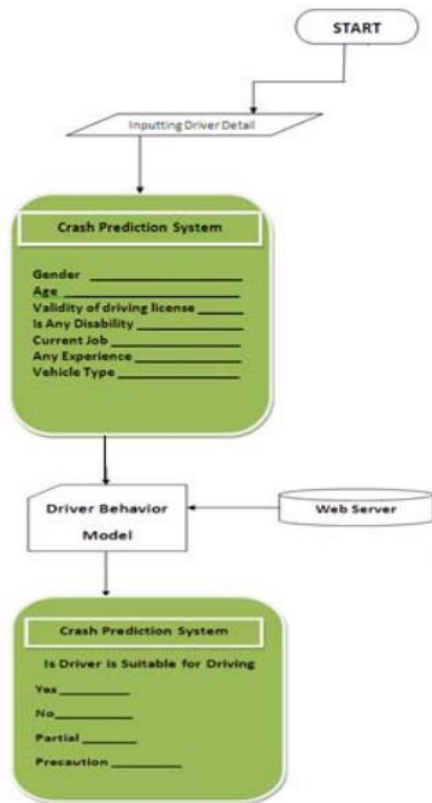


Figure 10 Block diagram of the Crash Prediction System (Singh, et al. 2012)

Crash prediction system is the user interface with which the user can interact. It takes inputs from the user and produces the output in a graphical model using HMM. PCA emphasizes the relationships between age, gender and vehicle type, to the crash variables. HMM then finds its use in categorizing the sequences of data by taking input of one or two dimension made by PCA. When all the variables are collected and analyzed, the application generates its verdict according to the following algorithm⁶²:

INPUT		
All Values of	If between 30 to	For all values of
Age, Kms, Health	45,	Health
and Vision	Then score is -2	If 0,
	Else, -0.5	Then score is -2
OUTPUT	For all values of	If 5,
Completely Fit	Kms	Then score is 0.5.
Partially Fit	If between 1000	If 10,
Not Fit	to 2000,	Then score is 2.
START	Then score is 0.5.	For all values of
Variable Score;	If between 2000	Vision,
For all values of	to 5000,	If -0.25,
Age	Then score is 1.5.	Then score is 0.5.
If between 18 to	If above 5000	If 0.25,
30,	Then score is 2.5.	187
Then score is 1	Else, -0.5	Then score is 1.

⁶² Paragraph based on Singh, G. R., & Dongre, S. S. (2012). Crash Prediction System for Mobile Device on Android by Using Data Stream Mining Techniques. *2012 Sixth Asia Modelling Symposium*, 185-189. Retrieved June/July, 2019.

If -0.75,	If Score ≤ 0 ,	If Score ≥ 5 ,
Then score is -2.	Then Not Fit.	Then Completely
END FOR.	If Score > 0 and	Fit.
Total Score =	Score ≤ 4	STOP
Age+ Kms+	Then Partially	
Health + Vision.	Fit.	

- Fuzzy Logic** represents possible decisions in natural language which can express our decisions in terms of words. This allows computers to think and reason like a human being. The process is constructed from a fuzzifier, interference engine, rule base and defuzzifier. The fuzzifier maps contain rules provided by the programmer.

The practical use of such system is shown in a 2008 study performed by several authors where they combined the inputs of several sources and devices such as an accelerometer, OBD reader and a camera to create a base for computing hazardous driving.




Figure 11 Hazardous driving detection system (Imkamon, et al. 2008)

All data from the vehicle was collected and recorded on a laptop and three passengers were also asked to assess how safe they felt on a scale from one to three every five minutes in order to calibrate the system. After the calibration was completed, the driver

was asked to drive at different levels of risk after which the results from passengers and the fuzzy logic system were compared⁶³. The results were as follows:

Time interval	Output	
	Average passenger value	System output
1	1	1
2	1,5	1,5
3	2	1,5
4	1,5	1
5	1	1
6	1	1

- **Temporal Difference Learning (TD)** is a supervised learning technique used for reinforcement- based learning measuring the expected rewards. It achieves its goal in a succession of time steps, hence the name temporal difference learning.

An example of TD in use is the seen in the work of several authors where they collected multi- sensor data through the STISIM driving simulator. Data included a time stamp, distance, lane position, acceleration due to throttle and negative acceleration due to braking, velocity, steering, throttle and brake input and so on. 36 drivers were driving for about 20 minutes, having between one and three accidents per driving course. Each time one driver is selected for testing and the others for training, creating an average performance. Afterwards, the features were extracted. Mean, max, min and variance were extracted for each dimension and a weave was created. The weave is a frequency of the vehicle's oscillation in the lane, detecting a drowsiness of a driver to some extent⁶⁴.

⁶³ Paragraph based on Imkamon, T., Saensom, P., Tangamchit, P., & Pongpaibool, P. (2008). Detection of hazardous driving behavior using fuzzy logic. *5th International Conference on Electrical Engineering/Electronics, Computer, Telecommunications and Information Technology*, 1-4. Retrieved June/July, 2019.

⁶⁴ Paragraph based on Ning, H., Xu, W., Zhou, Y., Gong, Y., & Huang, T. (December 2008). Temporal Difference Learning to Detect Unsafe System States. *19th International Conference on Pattern Recognition*, 3-4. Retrieved June/July, 2019.

Although the main areas of focus in this paper are banking trading and transport, I will mention some of the notable uses of AI in business from other sectors as well in the table below:

Agriculture	Autonomous harvest robots Crop and soil monitoring via drones Predictive analysis of environmental impacts on crop yields
Automotive	Driver assist (emergency braking, blind-spot monitoring, steering correction) Driverless trucks and cars (Mercedes-Benz, Tesla, Waymo) Predictive maintenance and OTA updates Driver identification and monitoring
Banking	Chatbots Fraud prevention Personalized insights Analyzing loan applications Cryptocurrencies
Defense	Advanced cybersecurity Autonomous vehicles (drones, ATV's, demining vehicles) Target recognition Combat simulation and training
Energy	Forecasting energy use Energy efficiency monitoring
Healthcare	Analysis of patohistological samples Intelligent MRI scanning Administrative task automation
High tech	IoT applications Big data analysis Virtual assistants Cloud computing Blockchain technology
Logistics and transportation	Robotized warehouses (autonomous transport robots, robotic sorting and packaging) Efficient route planning Autonomous trucks (human supervision is mandatory by law)

	Delivery drones
Manufacturing	Machine vision quality checks Maintenance prediction Smart fuel-valve adjustment Cutting waste Post-production AI-driven support
Media and entertainment	Personalized marketing Search optimization AI-driven movie set creation
Public and social sectors	Social media surveillance (Las Vegas food department- restaurant review analysis via Twitter) Integrated CCTV monitoring (Beijing, London) Virtual call center assistants
Retail	Virtual sales and CRM assistance Analysis of frequently bought together products Sales analysis and prediction
Telecommunications	Robotic Process Automation (RPA Network optimization) Preventive maintenance Virtual assistants
Tourism	Predictive systems in trip planning Recommendation systems Sales and cost optimization Real-time translation Chatbots
Trading	High-Frequency Trading (HFT) Individual and mutual stock analysis Market analysis and prediction

Table 2 Present-day examples of AI implementation in business

The use of AI in these sectors, among others, resulted in decreasing costs in the long-run, more efficient processes, and better performance. These applications are just the tip of the iceberg as more advanced forms of AI are looming on the horizon, the future is brimming with intricate

examples of massive interconnected AI systems, and some of these uses will be presented in the last chapter of this paper.

6. Critical analyses of the current state of AI in banking, trading and transportation

AI applications are all-pervading far and wide in today's business environment. Companies are constantly searching more innovative ways to utilize modern technology in a way which speeds up their processes, cuts long-term expenses and gives their company the leading edge. Though this phenomenon is present throughout sectors, my work focuses on examples from the fields of finance and transportation, respectively. I chose the two domains because of the high propensity of AI utilization within these sectors.

Starting with the **banking sector**, more and more banks utilize AI in their operations. As previously mentioned in the paper, examples in use are divided between supplementing human performance and substituting people. The type of utilization depends on the nature of the task. AI excels in collecting, organizing, analyzing and presenting vast quantities of data, as well as, performing tedious, time-consuming tasks so such actions are readily completed solely by AI. On the other hand, tasks involving advanced human interactions, including empathy and critical thinking, are still safe from AI automation at this point in time.

The current state of AI in banking spans through the following sectors:

- **Customer support:** Banks use AI for their always on customer support services through chatbots, some notable examples have been laid out in the paper so far. In my opinion, this is a great use of AI which will hopefully result in a general positive acceptance of new technology in our AI driven era. Technophobia is still a major negative influence in the world today, and it isn't showing signs of slowing down due to some well-founded privacy concerns. Yet, AI is here to stay, and it can make life easier for us if we use it to our advantage. With the advancement of its capabilities, bank customers will be able to handle critical situations such as having their cards stolen or having made a wrong payment in a prompt manner through their smartphones around the clock. Additionally, chatbots learn from past interactions with the bank's clients so future interactions can be predicted and the commonly used functionalities will be more easily accessible. Furthermore, the bank will be able to tailor fit its products to each specific client.
- **Analysis of loan applications:** The area of loan applications is among the first sectors in banking where AI will substitute humans. That is because of the infamous ability of AI to analyze massive quantities of data in a short amount of time through pre-set

parameters to give a verdict. Besides, unlike humans, they are not prone to being emotionally attached to certain options and clients thus eliminating cognitive bias (at least until sentient AI systems begin running companies and bank loan approvals). However, some similar bias may be present due to a fault in programming as the pre-set parameters can be manipulated upon in the application if they are known. Still, another advantage of AI is the fact that such a system usually takes a holistic approach where the client's past credit score and transactions are part of the database from which the system draws an informed decision.

- **Fraud prevention:** Fraud preventions with AI takes on multiple functions. In the USA, authenticity of checks is being monitored through an AI system from the company Cognizant. Their AI system utilizes Optical Character Recognition (OCR) to scan checks, process data and analyze signatures which allows for an incredible 70 millisecond evaluation time. Already, more than 20 million dollars was saved through the work of Cognizant's AI system. This figure will climb exponentially as the algorithm expands its current database with more advanced examples of check fraud. Furthermore, other AI systems battle with credit card fraud, which in the UK alone caused over 2 billion pounds worth of damage in the past year.
- **Cutting costs and quickening processes:** What most AI feats in the business world boil down to is gaining an advantage in the marketplace through cutting costs, building efficiencies, and most often, a combination of the two. The banking world is no exception to the rule. The figure below shows the full extent of revenues and savings that banks could reap if they utilize AI from 2019 to 2023. Front office use would bring around 200 billion dollars while middle office uses in fraud prevention would save around 220 billion dollars. Last, but not least, back office uses would bring around 30 billion dollars in savings. Banks which already utilize some form of AI such as Bank of America with their chatbot Erica, already reap the rewards in the marketplace.

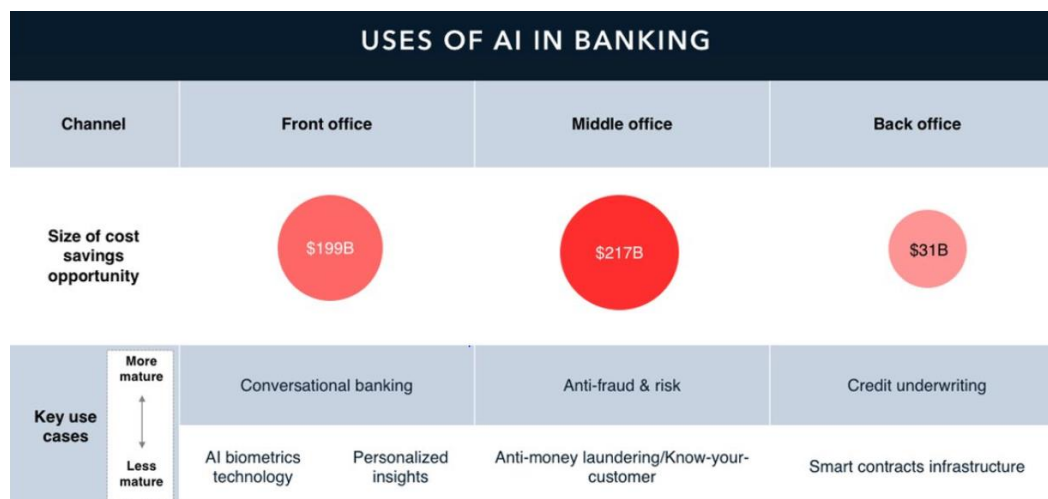


Figure 12 Cost savings in banks due to AI utilisation (Digalaki, E. 2019)⁶⁵

In my opinion, these applications of AI in banking are very well grounded and form a stable foundation for future expansion. The use of virtual assistants provides a dose of promptness and connectedness which younger generations of user's praise. These two characteristics are closely connected with its rapport. In fact, I think that the two words which best describe today's personal and working life are NOW and DIGITAL. Any person who is not instantly available through their digital devices, is a very rare breed. This outlook has created the same pressure for companies. Big and small companies alike have their own Twitter, Facebook, Instagram, and other social media accounts which are frequently updated. Comments and complaints are dealt with in a matter of minutes, if not seconds. Therefore, working hours are a matter of the past, and banks must adapt as well. Bank's vital services are implicitly required to be available 0-24 hours every day of the year, even during the weekend, and even during holidays. This creates a large pressure on the bank's employees, which is now mitigated through a precise use of AI in customer services. Additionally, the time-consuming work of fraud prevention is also lifted from the backs of specialists onto the monumental back of AI systems.

However, not everything is so rosy. A valid concern is the risk of vulnerable servers and systems. Most AI systems which are used across companies today are developed by numerous IT firms and operated on an even smaller number of notable servers. Thankfully, the servers are extra ordinarily capable to handle large flows of traffic in addition to their excellent resilience against cyberattacks. Still, the question of software oligopoly creates a set of globalized benefits and problems alike. For example, companies usually rely on a number of IT companies to develop AI systems for their specific use in order to gain a leading edge,

⁶⁵ Digalaki, E. (2019) AI in banking: AI could be a near 450 billion dollar opportunity for banks <https://www.businessinsider.com/the-ai-in-banking-report-2019-6>. Retrieved: June/July 2019

however, other competing companies can buy the same services for the same problems which creates a new plateau, leveling the playing field. On the other hand, relying on turnkey solutions and broadening the customer base of IT companies creates a larger pool of resources from which the provider can learn and improve upon. Additionally, the price of such turnkey solutions is lower than if an AI system is built from the ground up every time it is utilized by another client.

Furthermore, use of AI in banking raises red flags when it comes to privacy. Virtual assistants and loan approval systems alike, have access to large databases of personal and vulnerable transaction information. The benefit of tailor made offers and targeted assistance comes at a price some are just not willing to pay, and that is privacy. However, privacy is almost an impossible commodity to afford nowadays, as apps, sites, companies and even whole countries are gathering more and more sensitive data about their users, client and citizens to create more enticing, if not addictive, services. Therefore, the data gathered and used by banks has to remain private, without exposure to outside parties, while not forgetting the banks fiduciary responsibility.

With **trading**, on the other hand, the main concern is not privacy, but safety. When funds from individual and corporate investors are pooled in investment companies' funds, investees expect a degree of safety. Of course, they are taking risks, but their risk is calculated with regards to the profits they hope to gain, if they invest in riskier funds, they expect greater profits. However, autonomous AI driven funds can add a higher degree of risk due to the possibility of malfunctions. Such malfunctions have occurred in the past with HFT software and as a result of them, investing companies and their clients lost sizeable amounts of money. Modern algorithms came a long way since their initially vulnerable state, but more refined malfunctions and greater cyberthreats are still a valid concern in today's AI investing scene. If some similar disaster does occur, the ever-present question which spans through all autonomous AI developments remains, and that is who will be responsible?

In addition to such threats, another more likely negative scenario is the proverbial race to the bottom which is occurring with trading companies. A great deal of traders pay a large premium to be ever so closer to the stock exchange. Additionally, huge sums of money are invested in optic cables which provide a marked improvement in internet speeds. This all boils down to miniscule profits on individual transactions, however on a yearly basis these huge investments in better infrastructure still pay off by a big margin. Nonetheless, as more and more companies

invest in such infrastructure, and as they fight for the closest positions near trading hubs, profits become smaller and smaller, eventually leading to diseconomies of investing. Coupled with investments in faster and better algorithms for HFT as well as autonomous AI investing, the only survivors in the market will be companies with the latest and greatest software. Although not for long, as such competition is predicted to lead trading into a sort of singularity equilibrium where profits become slimmer and slimmer as more traders utilize algorithmic trading.

Finally, we come to the subject of **transportation**. Various instances of AI implementation occurred in recent years. The proposed concepts from examples I included in the paper have partly come true. For example, in the second case study researchers suggested that through AI utilization, vehicles could identify their drivers. A similar undertaking has been performed by a Croatian company Rimac automobili where their Concept Two car is able to scan the user's face and therefore unlock the car without a need for classical key fobs. However, unlike vehicle predictions in the examples above where drivers would be identified through their driving patterns, here the AI system relies on image processing through a detailed 3D image scanning.

Furthermore, the same case study details AI's role in accident identification, modern car systems already have an SOS call feature where the vehicle detects if an accident occurred using in-vehicle motion sensors to call emergency services, giving them the coordinates to the crash site. Furthermore, modern vehicles are already equipped with radar systems and cameras which utilize AI in traffic sign recognition as well as in distance control. The vehicle is able to display warning messages if it predicts that another vehicle is too close in front of it. In some cases, the system can autonomously apply brakes to slow down the vehicle, or even stop it.

Moreover, a general trend in transportation is autonomous driving. Companies like Tesla and some others have already put into use modes for autonomous driving, however human presence is still required, and in some cases, it is necessarily implemented to overrule actions. There are several stages of autonomous driving and the most advanced ones are already developed and tested on closed circuits. Mercedes Benz has already developed autonomous trucks which are being tested on highways with their human chaperones, other companies such as Volvo are also developing such vehicles in a prompt fashion. Volvo has already developed a small unmanned truck for use in closed off areas for transport. Companies like Uber, Domino's pizza and Waymo all have their own developments under way for an unmanned future. Many of such cars are already on the road with their human supervisors.

In my opinion such uses are well justified and help to pave the way for a safer future on the road. Human drivers and their follies are to blame for astoundingly large numbers of road accidents and the related human casualties due to carelessness, tiredness, drunkenness, and general inconsideration. For example, if there were no human drivers, there would be no traffic stalls, AI driven vehicles would always maintain equal distance, starting and stopping in a uniform manner. A system where driverless cars make up the entirety of traffic, and a redefined infrastructure is the basis for such operation, would be infinitely more reliable and safer than the present human operated systems. However, I am not so fond of mixed systems where drivers are intermingled with autonomous vehicles on existing roads. I believe that human irrationality and other follies would create a chaotic environment for autonomous vehicles to operate, perhaps leading to some accidents in the process. Such accidents have already happened, autonomous vehicles have run over some people and have been involved in certain crashes. Of course, the number is miniscule, and the fault was very rarely due to the vehicle itself and mostly to the other participants, but the media vilified those companies and vehicles far more than any driver would ever be. Therein lies the folly, if such an accident does occur, who is to blame and who takes the responsibility? That is the main dispute in transportation related AI.

7. Prospects in applied AI

The speed of future AI implementation in **transportation** is uncertain due to the main constraints imposed by legislative bounds. Today's AI is capable of autonomous operation when it comes to certain applications such as driving, however, due to liability issues in case of accidents, its autonomous application is not possible. Legislation is bound to change in the coming years, with some initial progress being made today as well. Numbers of robots and artificially intelligent machines in business have been rising exponentially. Most manufacturing across the world is now robotized and mostly automated as such use transfers liability on the backs of companies.

For example, certain Amazon warehouses are highly automatized. In fact, its warehouse near Denver airport is their crown jewel. Small robotic devices whizz around the warehouse with high stacks of products on top of them from one part of the warehouse to another. Humans are taught special procedures for interaction and their role is shrinking drastically. Today, they are mainly in charge of packaging, with all the heavy lifting and whizzing around performed by their robotic counterparts⁶⁶. Amazon's success lies in its capital endowment and tech savviness because they created their own robots and algorithms. Other companies worldwide are also joining the bandwagon of warehouse automation, but with bought know-how and hardware. One such developer is Boston Dynamics which created several robots, and one of their latest creations is designed to be used in warehouses. It is called the 'rolling handle' robot due to its arm-like extension which handles boxes using special suction technology⁶⁷. Amazon has also pioneered drone delivery with its fleet of automated drones.

Companies from different sectors are relying on autonomous vehicles to provide their services in a more efficient way. They span from Domino's pizza delivery cars through Waymo's personal transport concepts and all the way to autonomous ride sharing. In fact, Uber bought the artificial intelligence group Geometric Intelligence in 2016 to form the core of its proprietary research⁶⁸. Carmakers are no exception when it comes to the development of autonomous vehicles, Tesla, Mercedes-Benz and several other manufacturers already have some sort of "autonomous" driving mode enabled, but they are not yet near level 4 autonomous

⁶⁶ Simon, M. (2019). Inside the Amazon warehouse where humans and machines become one. Retrieved from: <https://www.wired.com/story/amazon-warehouse-robots>. Retrieved: June/July 2019

⁶⁷ Simon, M. (2019). What Boston Dynamic's rolling handle robot really means. Retrieved from: <https://www.wired.com/story/what-boston-dynamics-rolling-handle-robot-really-means>. Retrieved: June/July 2019

⁶⁸ Uber steps up efforts on artificial intelligence. (2016). Retrieved from: <https://search.proquest.com/docview/1845952802?accountid=132154>. Retrieved: June/July 2019

driving. Mercedes-Benz and Volvo have pioneered level 4 autonomous trucks, Renault and others have also developed their versions of level 4 equipped passenger cars. For example, Renault's concept version includes an AI headset which is capable of creating specific sceneries for different driving moods⁶⁹.

Essentially, the only barrier to autonomous driving at this present moment is the regulatory system and the main question of liability for possible accidents caused by such vehicles. If such an accident does occur, who will be liable for it? The owner of the vehicle, the manufacturer or the program developer? Furthermore, if such questions become regulated, the next concern is for the state of jobs. The vast majority of lorry drivers as well as bus, taxi and delivery drivers will eventually be out of work as level 4 and subsequently more advanced autonomous vehicles will take their place.

When it comes to **banking and trading**, increasing numbers of jobs are being threatened by artificially intelligent machines today. Both sectors are data driven, which creates an unfair playing ground as AI can leverage much more data in a prompt manner compared to their human counterparts. Furthermore, as more advanced types of AI become a reality, almost all jobs will become obsolete in a manner where humans are simply less efficient than machines. Superintelligent AI paints both a positive and grim picture for humans. On the one hand, such intelligence will allow us to tap into more advanced knowledge and economies of scale, but it will also create grounds for our possible obsolescence. Banks are set to become automated in a progressively larger degree in the coming years, Deloitte proposes that most banks will become smart banks where they engage with their customers through AI, their analytics are driven by AI as well and their internal operations are supplemented by AI⁷⁰. Trading is no different as the number of AI driven and AI operated funds increases in number as well as with more advanced versions of HFT algorithms. However, the end state of trading is uncertain due to the competitiveness of rivaling HFT businesses. Furthermore, if and when higher level AI is developed and put into trading practice, who is the sentient program going to trade for? Why should it promote our goals if it can trade for its own account? Is a rogue AI program going to sabotage a trade on purpose?

⁶⁹ Testing the world's smartest car. Retrieved from: <https://www.youtube.com/watch?v=l3ELVACR2VY>. Retrieved: June/July 2019

⁷⁰ Banking on the future: Vision 2020. Retrieved from: <https://www2.deloitte.com/content/dam/Deloitte/in/Documents/financial-services/in-fs-deloitte-banking-colloquium-thoughtpaper-cii.pdf>. Retrieved: June/July 2019

Other notable examples of prospects in applied AI are listed in the table below:

Agriculture	Farm to table autonomous harvest and processing plants Nanorobotic pest control Satellite and past/present data integration and analysis
Automotive	Driverless cars and trucks on public roads Automated maintenance cycles Autonomous service stations Targeted OTA control of rogue vehicles
Banking	Human-like automated tellers Massive reduction in human staff Bank services tailored for each customer
Defense	Autonomous weapons Unmanned fleets of warfare vehicles Cyber warfare taking central stage
Energy	Predicting renewable energy needs and optimizing delivery Smart, integrated power gridlines
Healthcare	Robotic operators and dentists AI-driven drug discovery Virtual nurses and pain management
High tech	Predictive instead of reactive business models Use of agile sprints to decrease time to market Always-improving product development models Wide implementation of Virtual Reality (VR) for business and personal use
Logistics and transportation	Improvement of supply chain reliability via blockchain technology Smart roads including sensors and IoT technology Self-driving cars and trucks Streamlining the supply chain with vehicle telematics
Manufacturing	Widespread and directed automation Matching production with supply demands Generative design Digital twins

Media and entertainment	Deep and precise individual client profiles for precise personalization Holographic live performances AI generated movie stars
Public and social sectors	Containing disease spreading and fighting crime through integrated monitoring Providing targeted scholarships based on student's probability to graduate Real-time monetary adjustment with the help of AI monitoring and predictions Focused tax audits
Retail	Virtual shopping experiences AI-powered pattern learning on a massive scale Tailored shopping experiences
Telecommunications	Smart 5G networks Self-healing and self-learning grids
Tourism	Maximum impact journey planning Seamless, practically invisible AI border security Anytime anywhere narrative travel guides
Trading	Autonomous AI fund management Automated personal finances with an emphasis on maximizing value Possible end of High Frequency Trading (HFT) due to drastically diminished returns

Table 3 Future examples of AI implementation in business

Some of these uses are closer, and others are farther on the horizon, but the great majority of them is certain to be implemented. They will carry efficiencies of scale and scope with them, as well as savings in the long run. In the more distant future, however, the possibility of sentient AI is going to drastically change all of these sectors and will bring some other considerations with it as well.

If such machines become a reality and they truly become better than us in every way, what purpose will people serve? Today, most of the jobs at risk from AI automation are lower level, repetitive ones as well as some dangerous ones, and in most cases, people are supplementing

the machines. But in the future all jobs will become threatened. If that does occur, some valid questions will be raised, and a clear answer is not yet in sight. Many authors stake claim on their differentiating outlooks regarding such matters as the following: Will our population continue to rise in number, or will it drop in a slow or fast manner? Will jobs be a thing of the past where all countries have universal income for their citizens and machines perform the vast majority of work? Who will receive the royalties when AI begins to create works of art such as paintings or songs or symphonies? Which ethical standards will be implemented in developing sentient AI? The answers to these questions are still far from answered, but one thing is sure and that is further development of new awe-inspiring feats in technology.

8. Conclusion

In conclusion, the path which lead to modern artificial intelligence (AI) thought, and subsequently, implementation was long and varied. Spanning through centuries of intermingled work from some of the greatest minds. Its recent development since the 1950's endured periods of so-called winters and booms to become the ubiquitous staple of our modern business and personal lives. Today, machine learning (ML) and deep learning (DL) form the basis of our AI applications. Deep, intricate neural networks analyze massive quantities of data each day, working around the clock, to aid our quests for faster and more precise decision-making, whether it is in banking, trading, transportation, or medicine. These systems have become better than us in areas such as image processing. Which allows AI systems such as IBM Watson to detect certain types of cancer more promptly, as well as more accurately than us. Furthermore, due to their nature, such systems aren't prone to cognitive bias or fatigue errors which makes them ideal candidates for performing tasks in trading, loan applications, and a variety of other tasks. Such a task is driving. Autonomous vehicles have already reached the level 4 stage where they are able to safely maneuver through traffic on their own, however the legislative system is dated so despite the technological capabilities, autonomous vehicles are not yet able to be implemented without human chaperones. This is the bottleneck which is prone to stump future AI implementation in practice from different sectors. Questions of liability are great concerns as the autonomous vehicles cannot be punished if they cause an accident, a person will have to be fined or sentenced, but which person, the driver, the producer, or the programmer? It is due to this grey moral area that autonomous AI is advancing slowly when it comes to real-world interaction. Thankfully though, companies can choose to implement such autonomous AI on their premises and in their operations if they are ready to take on the liability. Therefore, tech-savvy modern companies such as Amazon, Tesla and Google, to name a few, already have fleets of autonomous systems spanning from warehouse robots, through AI driving support systems and all the way to autonomous vehicles, respectively. Legislation is bound to change, and technology is bound to progress. Yet, as technology progresses and more advanced forms of AI become a reality, humanity will be presented with some difficult questions which will impact the future of our personal and professional lives.

Literature:

1. Abady, M., Barham, P., Chen, J., Davis, A., et. Al. (2016) Tensorflow: A System for Large-Scale Machine Learning. 12th USENIX Symposium on OSDI. Retrieved: June/July 2019
2. Abduljabbar R, Dia H, Liyanage S, Bagloee SA. Applications of Artificial Intelligence in Transport: An Overview. Sustainability. 2019; 11(1):189; Page 2
3. Ahmed, A., Protes, B. Trading program ran amok, with no 'Off' switch. (2012). Retrieved from: https://dealbook.nytimes.com/2012/08/03/trading-program-ran-amok-with-no-off-switch/?_r=0&mtref=www.entefy.com&gwh=2995327C21A7C54B7D7583F1968CD11A&gwt=pay. Retrieved June/July 2019
4. AI MIT (2000). Retrieved from: <http://www.ai.mit.edu/projects/sociable/regulating-interaction.html>. Retrieved June/July 2019
5. Alan Turing's scrapbook. Retrieved June/July 2019, from <https://www.turing.org.uk/scrapbook/test.html>
6. Alcober, F., Gupta, S. Retrieved from: <https://medium.com/tensorflow/recap-of-the-2019-tensorflow-dev-summit-1b5ede42da8d>. Retrieved June/July 2019
7. Artificial Intelligence in medicine. Retrieved from: <https://www.ibm.com/watson-health/learn/artificial-intelligence-medicine>. Retrieved June/July 2019
8. Bank of America newsroom (2018). Retrieved from: <https://newsroom.bankofamerica.com/press-releases/consumer-banking/bank-america-delivers-first-widely-available-ai-driven-virtual>. Retrieved June/July 2019
9. Banking on the future: Vision 2020. Retrieved from: <https://www2.deloitte.com/content/dam/Deloitte/in/Documents/financial-services/in-fs-deloitte-banking-colloquium-thoughtpaper-cii.pdf>. Retrieved: June/July 2019
10. Baumann, N. (2013). Too fast to fail. Retrieved from: <https://www.motherjones.com/politics/2013/02/high-frequency-trading-danger-risk-wall-street/>. Retrieved June/July 2019
11. Bishop, J.M. (2018). Is anyone home? A way to find out if AI has become self-aware. Frontiers in robotics and AI. Retrieved from: <https://www.frontiersin.org/articles/10.3389/frobt.2018.00017/full#B1>. Retrieved June/July 2019.
12. Bostrom, N. (2006). How long before superintelligence? *Linguistic and Philosophical Investigations*, Vol. 5, No. 1, pp. 11-13. Retrieved June/July 2019.
13. Buchanan, B. G. (2006). A (Very) Brief History of Artificial Intelligence. AI Magazine, 26(4), 56. Retrieved June/July, 2019.
14. Caffe website. Retrieved from: <http://caffe.berkeleyvision.org/>
15. Deb, K. (2019). Retrieved from: <https://medium.com/datadriveninvestor/product-recommendation-using-association-rule-mining-for-a-grocery-store-7e7feb6cd0f9>. Retrieved June/July 2019.
16. Digalaki, E. (2019) AI in banking: AI could be a near 450 billion dollar opportunity for banks <https://www.businessinsider.com/the-ai-in-banking-report-2019-6>. Retrieved: June/July 2019
17. Ganesh, P., Rakheja, P. (January 2019). Deep reinforcement learnig in high frequency trading. CODS-COMAD 2019. (1-6). Retrieved June/July 2019
18. Goodfellow, I., Bengio, Y., & Courville, A. (2017). Deep learning. Cambridge, MA: MIT Press. (6-8)
19. H2O.ai website. Retrieved from: <https://www.h2o.ai/>

20. Hanson Robotics „Being Sophia“. Retrieved from:
<https://www.hansonrobotics.com/being-sophia/>. Retrieved June/July 2019
21. Heller, M. (2019). Deep learning explained. Retrieved from:
<https://www.infoworld.com/article/3397142/deep-learning-explained.html>. Retrieved June/July 2019
22. Hintze, A. Understanding the four types of AI (2016). Retrieved from:
<https://theconversation.com/understanding-the-four-types-of-ai-from-reactive-robots-to-self-aware-beings-67616>. Retrieved June/July 2019
23. IBM Deep Blue. Retrieved June/July 2019, from:
<https://www.ibm.com/ibm/history/ibm100/us/en/icons/deepblue/>
24. IBM icons of progress. Retrieved June/July 2019, from
<https://www.ibm.com/ibm/history/ibm100/us/en/icons/ibm700series/impacts/>
25. IBM Watson case study-Transforming banking support with an AI-powered virtual assistant. Retrieved from: <https://www.ibm.com/case-studies/crevalsistemieservizi>. Retrieved June/July 2019
26. Jajal, T.D. (2018). Distinguishing between Narrow AI, General AI, and Super AI.
<https://medium.com/@tjajal/distinguishing-between-narrow-ai-general-ai-and-super-ai-a4bc44172e22>. Retrieved June/July 2019
27. Jordan, M. I., & Mitchell, T. M. (2015). Machine learning: Trends, perspectives, and prospects. *Science*, 349(6245), 255-260. Retrieved June/July, 2019.
28. Joshi, N. 7 Types of artificial intelligence. Retrieved from:
<https://www.forbes.com/sites/cognitiveworld/2019/06/19/7-types-of-artificial-intelligence/#96809db233ee>. Retrieved June/July 2019
29. JPMorgan Chase Annual Report 2016. Retrieved from:
<https://www.jpmorganchase.com/corporate/investor-relations/document/2016-annualreport.pdf>. P.49. Retrieved June/July 2019
30. Klumpp, M. (2018). Automation and artificial intelligence in business logistics systems: Human reactions and collaboration requirements. *International Journal of Logistics Research and Applications*, 21, 3rd ser., 224-242
31. Kumar, C. (2018, August 31). Artificial Intelligence: Definition, Types, Examples, Technologies. Retrieved June/July, 2019, from
<https://medium.com/@chethankumargn/artificial-intelligence-definition-types-examples-technologies-962ea75c7b9b>
32. Kurzweil, R. (2017). AI Will Not Displace Humans, It's Going to Enhance Us. Retrieved from: <https://futurism.com/ray-kurzweil-ai-displace-humans-going-enhance>. Retrieved June/July 2019
33. LEE, K. (2019). AI SUPERPOWERS: China, silicon valley, and the new world order. S.I.: MARINER BOOKS.
34. McKinsey Global Institute. Notes from the AI Frontier. Insights from Hundreds of Use Cases; McKinsey Global Institute: San Francisco, CA, USA, 2018; p. 36
35. Mortleman, J., (2017). Retrieved from: <https://www.icio.com/management/insight/item/banking-on-an-ai-future>. Retrieved: June/July 2019
36. NatWest begins testing AI driven ‘digital human’ in banking first. (2018). Retrieved from: <https://www.rbs.com/rbs/news/2018/02/natwest-begins-testing-ai-driven-digital-human-in-banking-first.html>. Retrieved June/July 2019
37. Pandya, J. (2019). The troubling trajectory of technological singularity. Retrieved from: <https://www.forbes.com/sites/cognitiveworld/2019/02/10/the-troubling-trajectory-of-technological-singularity/#7038d3496711>. Retrieved: June/July 2019
38. Paragraph based on Dean, J., & Holzle, U. (2017). Build and train machine learning models on our new Google Cloud TPUs. *Google Cloud Blog*. Retrieved June/July,

- 2019, from <https://blog.google/products/google-cloud/google-cloud-offer-tpus-machine-learning/>.
39. Paragraph based on Dean, J., Corrado, G. S., Monga, R., Chen, K., Devin, M., Le, Q. V., . . . Ng, A. Z. (2012). Large Scale Distributed Deep Networks. 1-9. Retrieved June/July, 2019.
 40. Paragraph based on Halim, Z., Kalsoom, R., Bashir, S., & Abbas, G. (october 2016). Artificial intelligence techniques for driving safety and vehicle crash prediction. *Artificial Intelligence Review*, 46(3), 351-387. Retrieved June/July, 2019.
 41. Paragraph based on Hashmienejad, S. H., & Hasheminejad, S. M. (2017). Traffic accident severity prediction using a novel multi-objective genetic algorithm. *International Journal of Crashworthiness*, 22(4), 425-440. Retrieved June/July, 2019.
 42. Paragraph based on Imkamon, T., Saensom, P., Tangamchit, P., & Pongpaibool, P. (2008). Detection of hazardous driving behavior using fuzzy logic. *5th International Conference on Electrical Engineering/Electronics, Computer, Telecommunications and Information Technology*, 1-4. Retrieved June/July, 2019.
 43. Paragraph based on Jouppi, N. (2016). Google Custom Search. *Google Cloud Blog*. Retrieved June/July, 2019, from <https://cloud.google.com/blog/products/gcp/google-supercharges-machine-learning-tasks-with-custom-chip>.
 44. Paragraph based on Ki, Y., Heo, N., Choi, J., Ahn, G., & Park, K. (2018). An incident detection algorithm using artificial neural networks and traffic information. *2018 Cybernetics & Informatics (K&I)*, 1-5. Retrieved June/July, 2019.
 45. Paragraph based on López, R., Gurrola, L. G., Trujillo, L., Prieto, O., Ramírez, G., Posada, A., . . . Méndez, L. (2018). How Am I Driving? Using Genetic Programming to Generate Scoring Functions for Urban Driving Behavior. *Mathematical and Computational Applications*, 23(2), 1-13. Retrieved June/July, 2019.
 46. Paragraph based on Ning, H., Xu, W., Zhou, Y., Gong, Y., & Huang, T. (December 2008). Temporal Difference Learning to Detect Unsafe System States. *19th International Conference on Pattern Recognition*, 3-4. Retrieved June/July, 2019.
 47. Paragraph based on Singh, G. R., & Dongre, S. S. (2012). Crash Prediction System for Mobile Device on Android by Using Data Stream Mining Techniques. *2012 Sixth Asia Modelling Symposium*, 185-189. Retrieved June/July, 2019.
 48. Paragraph based on Tricentis - Artificial intelligence software testing. (n.d.). Retrieved June/July, 2019, from <https://www.tricentis.com/artificial-intelligence-software-testing/ai-approaches-rule-based-testing-vs-learning/#>
 49. Paragraph based on Yisheng, L., Yanjie, D., Wenwen, K., Zhengxi, L., & Fei-Yue, W. (april 2015). Traffic Flow Prediction With Big Data: A Deep Learning Approach. *IEEE TRANSACTIONS ON INTELLIGENT TRANSPORTATION SYSTEMS*, 16(2), 1-9. Retrieved June/July, 2019.
 50. Pascaline – Encyclopedia Britannica. Retrieved June/July, 2019, from <https://www.britannica.com/technology/Pascaline>
 51. Polzin, J.A. (2019). Intelligent Scanning Using Deep Learning for MRI
 52. Premack D., Woodruff G. Does the chimpanzee have a theory of mind? (1978). *THE BEHAVIORAL AND BRAIN SCIENCES*, 4, 515. Retrieved June/July, 2019.
 53. Rabinowitz, N. C., Perbet, F., Song, H. F., Zhang, C., Eslami, S. A., & Botvinick, M. (2018). Machine Theory of Mind. (1-21). Retrieved June/July, 2019.
 54. Reynoso, R. (2019). 4 Main types of artificial intelligence. Retrieved from: <https://learn.g2.com/types-of-artificial-intelligence>. Retrieved June/July 2019
 55. Roberts, C., Leichenauer, S. (2019) Retrieved from: <https://ai.googleblog.com/search/label/TensorFlow>. Retrieved June/July 2019

56. Shalom, N. (2019). Retrieved from: <https://www.gigaspace.com/blog/amazon-found-every-100ms-of-latency-cost-them-1-in-sales/>. Retrieved June/July 2019
57. Sharkley N. A show robot or dangerous platform to mislead? (2018). Retrieved from: <https://www.forbes.com/sites/noelsharkey/2018/11/17/mama-mia-its-sophia-a-show-robot-or-dangerous-platform-to-mislead/#4e587af7ac99>. Retrieved June/July 2019
58. Sighthound. Retrieved from: <https://www.sighthound.com/products/sighthound-video>. Retrieved June/July 2019
59. Simon, M. (2019). Inside the Amazon warehouse where humans and machines become one. Retrieved from: <https://www.wired.com/story/amazon-warehouse-robots>. Retrieved: June/July 2019
60. Simon, M. (2019). What Boston Dynamic's rolling handle robot really means. Retrieved from: <https://www.wired.com/story/what-boston-dynamics-rolling-handle-robot-really-means>. Retrieved: June/July 2019
61. Spremić, M., Ph.D. (2017). Enterprise information systems in digital economy. Zagreb, Pg.7
62. Testing the world's smartest car. Retrieved from: <https://www.youtube.com/watch?v=l3ELVACR2VY>. Retrieved: June/July 2019
63. The Chess Turk explained. Retrieved June/July 2019, from <https://youtu.be/0DbJUTsUwZE>
64. The development of the first artificial intelligence program. Retrieved June/July 2019, from: <http://www.historyofinformation.com/detail.php?id=742>
65. The Turk. Retrieved June/July 2019, from <https://interestingengineering.com/the-turk-fake-automaton-chess-player>
66. Uber steps up efforts on artificial intelligence. (2016). Retrieved from: <https://search.proquest.com/docview/1845952802?accountid=132154>. Retrieved: June/July 2019
67. Zhuang, Y., Thiagarajan, A., Sweeney, T. (2019). Ranking Tweets with TensorFlow. Retrieved from: <https://medium.com/tensorflow/ranking-tweets-with-tensorflow-932d449b7c4>. Retrieved June/July 2019

List of tables:

Table 1 Types of artificial intelligence according to Kumar, C.	8
Table 2 Present-day examples of AI implementation in business	- 45 -
Table 3 Future examples of AI implementation in business.....	- 56 -

Table of figures:

Figure 1 Automobile recognition using Sighthound.....	10
Figure 2 Age, gender, race, and mood recognition in humans using Sighthound	10
Figure 3 Kismet's socializing responses (AI MIT)	15
Figure 4 Consistent brain image orientation in three various head rotations (Polzin, J.A.)	25
Figure 5 AI assistant Cora from NatWest bank	28
Figure 6 Pipeline model proposed by the authors (Ganesh, Rakheja 2019).....	32
Figure 7 The knowledge accumulation gap (Klumpp 2018, 227)	33
Figure 8 Type of gathered data in traffic (Halim Z, et al. 2016, 353)	36
Figure 9 Schematic view of the TCPS process (Ki, et al. 2018)	39
Figure 10 Block diagram of the Crash Prediction System (Singh, et al. 2012).....	41
Figure 11 Hazardous driving detection system (Imkamon, et al. 2008).....	- 42 -
Figure 12 Cost savings in banks due to AI utilisation (Digalaki, E. 2019)	- 49 -