A Strategic Perspective on the Commercialization of Artificial Intelligence: A socio-technical analysis

by

Siddhartha Ray Barua

B.E (Information Science & Engineering), MSRIT (2007)
PGDIM (Industrial Management), NITIE (2012)

Submitted to the System Design & Management Program in partial fulfillment of the requirements for the degree of

Master of Science in Engineering and Management

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2019

Signature redacted

Author ....

System Design and Management Program
Signature redacted May 14, 2019

Certified by ....

Donna H. Rhodes
Principal Research Scientist, Sociotechnical Systems Research Center
Director, Systems Engineering Advancement Research Initiative, MIT
Signature redacted

Accepted by ....

Joan Rubin
Executive Director, System Design & Management Program

This page is intentionally left blank

A Strategic Perspective on the Commercialization of Artificial Intelligence: A socio-technical analysis

by

Siddhartha Ray Barua

Submitted to the System Design & Management Program in partial fulfillment of the requirements for the degree of Master of Science in Engineering and Management at the Massachusetts Institute of Technology, June 2019

ABSTRACT

Many companies are increasing their focus on Artificial Intelligence as they incorporate Machine Learning and Cognitive technologies into their current offerings. Industries ranging from healthcare, pharmaceuticals, finance, automotive, retail, manufacturing and so many others are all trying to deploy and scale enterprise AI systems while reducing their risk. Companies regularly struggle with finding appropriate and applicable use cases around Artificial Intelligence and Machine Learning projects.

The field of Artificial Intelligence has a rich set of literature for modeling of technical systems that implement Machine Learning and Deep Learning methods. This thesis attempts to connect the literature for business and technology and for evolution and adoption of technology to the emergent properties of Artificial Intelligence systems. The aim of this research is to identify high and low value market segments and use cases within the industries, prognosticate the evolution of different AI technologies and begin to outline the implications of commercialization of such technologies for various stakeholders. This thesis also provides a framework to better prepare business owners to commercialize Artificial Intelligence technologies to satisfy their strategic goals.

Thesis Supervisor: Donna H. Rhodes
Title: Principal Research Scientist, Sociotechnical Systems Research Center
Director, Systems Engineering Advancement Research Initiative, MIT

This page is intentionally left blank

ACKNOWLEDGEMENTS

The MIT experience has been transformational and awe-inspiring for me. Throughout my time here, the professors and researchers I have had the opportunity to interact with challenged and stretched my thinking. I am grateful to the Executive Director, Joan Rubin for admitting me into a program that has been very flexible in allowing me to explore the many opportunities offered by the school. Many thanks to professors Bryan Moser, Edward Crawley, Bruce Cameron and Eric Rebentisch for all the hard work that went into preparing the core class.

I express my sincere gratitude to my thesis advisor Dr. Donna Rhodes for providing guidance throughout my thesis writing process. She helped me shape my idea, provided extensive research material and gave periodic feedback.

I would also like to thank the Subject Matter Experts in Artificial Intelligence in Microsoft and Google, who shared their thoughts, experiences and resources to make this thesis a success.

There are no words to express the gratitude I feel towards my father, mother, sister and brother for always believing in me and serving as a pillar of support at difficult times. With immense sense of accountability, I acknowledge the pain and suffering my wife has endured in helping me get through these two years. She not only sacrificed so much but has also been a source of inspiration by following in my footstep and joining me as a graduate student in this esteemed institution. For their unwavering support, I dedicate this degree to my family.

I am proud to be a part of a unique SDM cohort at MIT. It was always great to hear a diversity of views on any specific topic. These folks made this journey so beautiful and enjoyable, and for that I am thankful.

This page is intentionally left blank

Table of Contents

LIST OF FIGURES

LIST OF TABLES

This page is intentionally left blank

CHAPTER 1: INTRODUCTION

Artificial Intelligence (AI) technology is on the cusp of going mainstream as a new engine of growth of business and society. The idea that computer programs can perform functions typically associated with the human mind, has gone from speculation to present-day reality. Thanks to advances in data collection and aggregation, algorithms, and processing power, computer scientists have achieved significant breakthroughs in artificial intelligence. Maps that compute the optimal driving route for a journey based on current traffic, movies suggestion from a video streaming service based on the movies the user watches, smart song suggestion from a music streaming service, friend recommendation from a social networking platform, product recommendations based on purchase history, and email spam filters - all are powered by AI.

With the advent of autonomous vehicles, digital personal assistants that can anticipate user needs and computerized health diagnosis, AI is realizing its potential to change people's lives for the better. The human development community is looking at ways to use AI-enabled innovations to achieve key Sustainable Development Goals (SGDs), including low-cost medical diagnosis expert systems, more efficient capital markets, use of automated drones in public emergencies, and to improve the efficiency and effectiveness of operations. Thus, by providing new information and improving decision-making through data-driven strategies, AI could potentially help to solve some of the complex global challenges of the 21st century, from climate change and resource utilization to the impact of population growth and healthcare issues.

Because they can dramatically boost productivity, AI technologies may have a disruptive impact on the global economy and on the workforce. Hundreds of millions of workers could be affected, and jobs made up of routine work activities and predictable, programmable tasks will be particularly vulnerable. AI capabilities have exciting and far-reaching potential to enhance human welfare by improving health care, the environment, security, and education. At the same time, AI also raises complex ethical, legal and security questions surrounding issues such as privacy, discrimination, liability and regulation.

1.1. MOTIVATION

The Artificial Intelligence market is starting to solidify around real-world applications with the pace of change being faster than it ever has been before, as startups and technology providers rush to create platforms and targeted niche solutions for solving specific enterprise problems. Artificial Intelligence is already key to how consumer internet companies operate today, allowing them to roll out hyper-personalized services by following an 'AI first' strategy. The rest of the market in the enterprise, government and non-profit sectors is still catching up on adopting AI and has yet to fully understand its value, including the breadth and depth of use cases, the technology choices surrounding AI, and the implementation strategies.

The business world is just beginning to harness these technologies and their benefits. Many organizations have deployed AI in a specific function but are yet to embed AI into multiple business units or function. Indeed, many organizations still lack the foundational practices to create value from AI at scale. One critical factor of using AI effectively is an organization's progress on transforming the core parts of its business through digitization. Another foundational challenge with AI is finding skilled people to implement it effectively. Many organizations still lack the foundational practices to create value from AI at scale. They do not have a clear strategy for sourcing the data that AI requires, nor they understand where their AI opportunities lie.

The literature review outlines the overview of Artificial Intelligence and looks at some of the seminal work that contributed to the field of AI. These observations underscore the need to provide a broad understanding of Artificial Intelligence, as well as a framework with which decision makers can use to adopt AI technologies to generate value at scale.

1.2. RESEARCH OBJECTIVES

This thesis lays out an overview of Artificial Intelligence, its benefits and challenges from both technology and business perspectives. Next, the thesis aims to provide a framework to guide organizations in commercializing various AI technologies and integrating AI technologies into their enterprises. This framework takes a holistic approach to architecting AI into the organization, by considering various technological, environmental and organizational aspects as well as stakeholders needs.

Four different research objectives are defined for this thesis:

A. Review the genesis of Artificial Intelligence and the overview of the field of AI.
B. Provide a technology perspective of AI by presenting an overview of the definition and concepts, landscapes and techniques.
C. Provide a business perspective of how industries have been commercializing AI and highlight market drivers and barriers and investigate the opportunities, challenges and trends in AI across the industry.
D. Formulate an AI Commercialization Framework that will guide organizations in commercializing and adopting Artificial Intelligence technologies for their enterprises.

1.3. SYSTEM PROBLEM STATEMENT

To achieve the above aims and objectives, the key research question was designed using the To-By-Using framework that articulates the System Problem Statement as defined in the book System Architecture by Crawley, Cameron and Selva [1.1]. The SPS of this thesis has been summarized in Table 1.

Table 1: System Problem Statement

To	Support organizations in the commercialization and implementation of Artificial Intelligence technologies
By	Providing an understanding of AI in the technology and business context and creating an AI Commercialization framework
Using	Literature Review, Technology and Market Analysis and System Architecting Principles

1.4. OUTLINE OF THESIS

This thesis consists of eight chapters that describe the background, literature review, research methodology, technology landscape, commercial perspective, ethical, legal & societal perspective, framework development, and conclusion to the above research questions. This section briefly describes the content of these chapters.

Chapter 2 reviews existing literature from journals, magazines, books and internet to gain an overview of the field of Artificial Intelligence.

Chapter 3 outlines the research methodology used in this thesis and explains the rationale for the approach.

Chapter 4 provides an overview of Artificial Intelligence from the Technical perspective through definitions, concepts and examples. The aim is to have a clear understanding of the fundamentals of various approaches to Artificial Intelligence such as Machine Learning and Deep Learning with a brief glance at the technical details and theories behind them.

Chapter 5 provides an overview of Artificial Intelligence from the Commercial perspective. Through survey of industries and cutting-edge research work, this chapter highlights potential opportunities to industrialize AI technologies.

Chapter 6 looks at the ethical, legal and societal implications of adopting Artificial Intelligence at scale across industries.

Chapter 7 proposes a framework to support organization in commercializing and implementing Artificial Intelligence technologies to gain strategic advantage.

Chapter 8 enumerates the summary of findings of this research study and recommends area for future work.

CHAPTER 2: LITERATURE REVIEW

To execute the approach and objective of this thesis, there will be some implied knowledge originating from several relevant subjects. This thesis will be drawing mostly from knowledge, frameworks and approaches from several fields.

The term Artificial Intelligence was first coined by John McCarthy in 1956 [2.1] when he organized the first academic conference on the subject on the campus of Dartmouth College. During the 1960's the first article of AI was published under the name "Computers and Thought" [2.2] and the computer mouse was invented by Doug Engelbart in his research lab at Stanford Research Institute [2.3]. The first ever international conference about AI was held in 1969 in Washington DC [2.4]. The 1970's witnessed a significant progress in the field of Robotics. A robot called Freddy [2.5] was built that was capable of assembling wooden models using vision technology. The INTERNIST-1 program [2.6] which was a broad-based computer-assisted diagnostic tool was also invented in the 1970's. In the 1980's neural networks with backpropagation became widely used [2.7] and "The Society of Mind" a theoretical description of how the mind actually works, was published by Marvin Minsky [2.8]. In the 1990's AI took significant developments in areas such as machine learning, virtual reality and in games. A chess program called Deep Blue created by IBM [2.9] defeated the best human chess player in the world in 1997. The decade also saw significant progress in Robotics and AI when the first autonomous robotics machinery system "Sojourner" [2.10] was deployed on the surface of Mars by NASA.

2.1. ROOTS OF ARTIFICIAL INTELLIGENCE

2.1.1. PHILOSOPHY

Philosophy has contributed a lot to the genesis of Artificial Intelligence. George Luger did research work on unifying the diverse branches of AI though its theoretical foundations.

According to him, the philosophical foundations of AI started with Aristotle [2.11], as his philosophical work formed the basis for modern science. The great philosopher scientist Galileo, used mathematics as a tool to challenge our understanding of the world. The work of Rene Descartes [2.12] on the theory of the mind was also very influential to AI:

- It established the separation of the body from the mind. This forms the basis of the methodology of AI – mental processes have an independent existence and follow their own laws.
- Once it become established that the mind and body were separated, it became necessary to find innovative ways to connect the two.

According to Rene, "from a rational point of view, the external world can be clearly reconstructed through the rules of mathematics. From an empirical point of view, knowledge can be explained through an introspective but empirical psychology, and not through clear and distinct ideas" [2.12]. This knowledge plays a significant role in the development of AI structures and programs. Therefore, the philosophical foundation of AI regarded thinking as a form of computation.

Gottfried Wilhelm Leibnitz was one of the first thinkers to take the materialist position which holds that the mind operates by ordinary physical processes [2.13]. This has the implication that mental processes can potentially be carried out by machines. Russell and Norvig based their philosophical analysis of intelligence options [2.14] on the notion that intelligence is a manifestation of rational action and that an intelligent agent takes the best action in a given situation. Philosophy gave rise to this idea which formed the basis of Artificial Intelligence by equating the behavior of the mind to that of a machine. The mind operates on knowledge encoded in some internal language, and that thought can be used to choose what actions to take.

2.1.2. MATHEMATICS AND LOGIC

Charles Babbage and Ada Lovelace's work focused on the "development of formal logic", [2.15] wherein the patterns of algebraic relationships were treated as entities that could be studied, resulting in the creation of a formal language for thought. George Boole introduced Boolean Algebra [2.16] in 1847, which was a formal language for making logical inference. Boole's operations of "AND", "OR" and "NOT" have remained the basis for all operations in formal logic.

Kurt Godel showed in 1931 that there are limits to what logic can do. His 'Incompleteness Theorem' [2.17] showed that in any formal logic powerful enough to describe the properties of natural numbers, there are true statements whose truth cannot be established by any algorithm. Mathematics has been used to manipulate statements of logical certainty as well as probabilistic statements, in addition to laying the groundwork for computation and algorithms. The field of economics has, by formalizing the problem of decision-making to maximize outcome, furthered the contribution of mathematics to the field of AI.

2.1.3. COMPUTATION

Once thinking came to be seen as a form of computation, the next steps were to formalize and mechanize it. William Jevon invented the Logic Machine in 1869, that could handle Boolean Algebra and Venn Diagrams. The Logic Machine [2.18] was able to solve logical problems faster than human beings.

In 1950, English Mathematician Alan Turing, who is often referred to as the father of modern computer science, published a paper entitled "Computing Machinery and Intelligence" [2.19] which was foundational to the field of AI. The paper itself began by posing the simple question, "Can machines think?". Turing then went on to propose a method for evaluating whether machines can think, which came to be known as the Turing

Test. The test, or "Imitation Game" as it was called in the paper, was put forth as a simple test that could be used to prove that machines could think.

The original game upon which Turing's idea was based required a man, a woman and an interrogator. The goal was for the interrogator to identify which of the participants was a man and which was a woman. Since the interrogator would be able to identify the gender of the respondent by their voice and handwriting, the answers to the interrogator's questions would be type written or repeated by an intermediary. For the Turing Test, one of those two participants would be replaced by a machine and the goal of the interrogator would not be to identify the gender of the participants, but which is human and which is a machine. The Turing test assumes that a computer that is indistinguishable from an intelligent human exhibit intelligent behaviour. The central idea that machines can think was a key to defining the field of AI because it provided a method to assess whether or not a machine can think.

2.1.4. PSYCHOLOGY AND COGNITIVE SCIENCE

Language is an important part of human intelligence. Much of the early work on knowledge representation was tied to language and informed by research into linguistics. Modern Psychology, Cognitive Psychology or Cognitive Science is the science which studies how the mind operates, how we behave, and how our brains process information.

Psychology helps us understand the decision-making process of the human brain, how it processes different scenarios to arrive at a particular decision. It is natural for us to try to mimic the way human brain lead to intelligent behaviour in our quest to build artificial intelligent systems. Conversely, it makes sense to explore the properties of artificial systems such as algorithmic models to test our hypotheses concerning human systems.

Many sub-fields of Artificial Intelligence are simultaneously building models of how the human system operates. Artificial Intelligence systems such as Deep Neural Networks mimics the working of the human brain are being deployed for solving real world problems.

2.1.5. BIOLOGY AND NEUROSCIENCE

Neuroscience is a specialized department of Biology focused on the study of the anatomy and physiology of the human brain, including structures, neurons and molecules. The computer scientist Nils Nilsson believes that aspects of biology and "life" in general have provided important clues about intelligence [2.20]. This includes principles relating to neurons and the workings of the human brain, psychology and cognitive science, evolution, development & maturation and bionics.

According to Nils, "our brains, which give rise to our intelligence, are made up of tens of billions of neurons, each connected to hundreds or thousands of other neurons. Each neuron is a simple processing device. However, large networks of neurons are extremely powerful computational devices that can learn how best to operate" [2.20].

Neuroscience, in discovering that the human brain can be said to be similar to computers in some ways, provided the institutive basis for AI. The field of "Deep Neural Networks" attempts to build biologically inspired Artificial Neural Networks based on simplified networks of artificial neurons, in an attempt to mimic the human brain network. The aim is to build powerful Artificial Intelligence systems that can behave or surpass the behaviour of the human brains. Neural networks work at a sub-symbolic level, whereas much of conscious human reasoning appears to operate at a symbolic level. Artificial neural networks perform well at many simple tasks, and provide good models of many human abilities. However, there are many tasks that they are not so good at, and other approaches seem more promising in those areas.

2.1.6. ENGINEERING

Engineering has made a more direct contribution to the field of Artificial Intelligence by being the tool used to create machines on which AI applications are allowed to run. Particular facets of the field that have made this possible include:

- Automata, Sensing and Feedback
- Statistics and probability
- Computational Thinking
- Digital Computer

CHAPTER 3: RESEARCH METHODOLOGY

This chapter describes the research methodology used in this thesis. Since this thesis encompassed both the social and technical domains, it combines both social science as well as scientific research methods to understand and generalize the trends observed in Artificial Intelligence. This thesis also used company surveys and deep industry analysis to formulate the framework for adoption and commercialization of Artificial Intelligence technologies.

3.1. RESEARCH APPROACH

The research takes an inductive method of reasoning as opposed to a deductive approach. The current research in the area of Artificial Intelligence tends to be more applied rather than fundamental theories. In inductive inference, we go from the specific to the general. We make many observations, discern a pattern, make a generalization, and infer an explanation. Inductive reasoning allows for the synthesis of the literature review and the conclusions from the company surveys, which forms the likeliest possible explanation for such sets. Inductive reasoning help to provide a broad-based explanation of the current and future trends in Artificial Intelligence in both scientific and business domains and address the set of research questions set forth in this thesis.

The potential of commercializing Artificial Intelligence is analyzed from the business perspective through survey of market research papers and business reports. In understanding the research findings and observations, the research paradigm chosen is that of an interpretative perspective because the data that is collected is a presentation of other people's construction about what they are up to, influenced by the social interactions between them.

3.2. RESEARCH DESIGN AND METHODS

Exploratory Research is conducted as the thesis does not aim to provide conclusive and final answers to the research questions. This research does not intend to provide conclusive evidence, but rather to explore the research topics with varying levels of depth. This is due in part to the incomplete body of knowledge of all Artificial Intelligence techniques and commercialization potential of these technologies. Conclusive research was not chosen as the overall design because it is difficult to generalize the verifiable insights given the dynamic field of Artificial Intelligence.

Primary data is collected from some of the leading companies working in the field of Artificial Intelligence. Any proprietary data is sanitized and normalized so that the data does not skew the conclusions. Secondary data is collected from document analysis of available literature from scientific journals, scholarly articles, books, internet and experts in the field. The procedure for document analysis is to read these literatures and document and describe statistics relevant to the research questions.

3.3. LIMITATIONS

A comprehensive survey of the literature and companies is not likely given that Artificial Intelligence is a broad, diverse and complex field that is changing as the research is conducted. A different approach could be to organize future research into the different taxonomies of AI research.

The choice of secondary data collection is also limited because companies do not articulate their strategy or share proprietary information. The research takes an interpretative approach to reason their strategies. Interviews were conducted with selected companies so that better insights can be drawn into their strategies of commercializing AI technologies.

Lastly, the proposed framework for evaluating Artificial Intelligence technologies for commercialization is yet to be tested in a real-world scenario despite the availability of data that points towards the applicability of the framework. Further research is anticipated for testing the viability of such frameworks in real business cases which will fine tune these guiding principles.

CHAPTER 4: ARTIFICIAL INTELLIGENCE – TECHNOLOGY PERSPECTIVE

The concept of Artificial Intelligence was first established in 1956. Artificial Intelligence is considered to be an emerging interdisciplinary subject involving theory, methodology, technology and system applications that integrates cybernetics, informatics, computer science, mathematical logic, neurophysiology and other disciplines and is used to simulate, extend and stretch human intelligence. The ultimate goal of Artificial Intelligence is to enable computers to think and act as capable as humans and eventually, to enhance human intelligence. Artificial Intelligence is particularly apt at sorting through data, finding patterns and making predictions.

Russell and Norvig provide a set of possible answers to the "What is AI?" question [2.14]. These answers all assume that Artificial Intelligence should be defined in terms of its goals. The answers all fall under a quartet of types placed along two dimensions. One dimension is whether the goal is to match human performance, or instead, ideal rationality. The other dimension is whether the goal is to build systems that reason/think, or rather systems that act. The situation is summed up in Table 2:

Table 2: Four Possible Goals for AI

	Human-Based	Ideal Rationality
Reasoning Based	Systems that think like humans.	Systems that think rationally.
Behavior Based	Systems that act like humans.	Systems that act rationally.

According to Russell and Norvig [2.14], AI was categorized in 4 buckets based on their capabilities:

Acting Humanly: The relevant test for this category is the Turing Test. These machines would possess the capabilities of natural language processing to enable it to communicate successfully in English, knowledge representation to store what it knows or hears, automated reasoning to use the stored information to answer questions and to draw new conclusions, machine learning to adapt to new circumstances and to detect and extrapolate patterns, computer vision to perceive objects and robotics to manipulate objects and move about.

Thinking Humanly: To create machines that mimic human thinking, we must understand how humans think. This is possible through introspection, psychological experimentation and brain imaging. Computer science and human psychology are brought together via cognitive science, which ultimately allows for the creation of what we now term as AI.

Thinking Rationally: This 'laws of thought' approach stresses the importance of logic in the computational process. This refers to acting in a manner to achieve one's goals, given one's beliefs. AI is viewed as the study and construction of rational agents.

Acting Rationally: AI is responsible for the creation of computer agents, who are expected to be rational agents and achieve the best possible outcome. This is not the same as 'Thinking Rationally' because not all best outcomes are the result of logical inferences. This approach has two advantages – it is more generic in nature and more amenable to scientific development.

4.1. STAGES OF AI RESEARCH

Research on Artificial Intelligence has undergone three main developmental stages, as illustrated in Figure 1 [4.1].

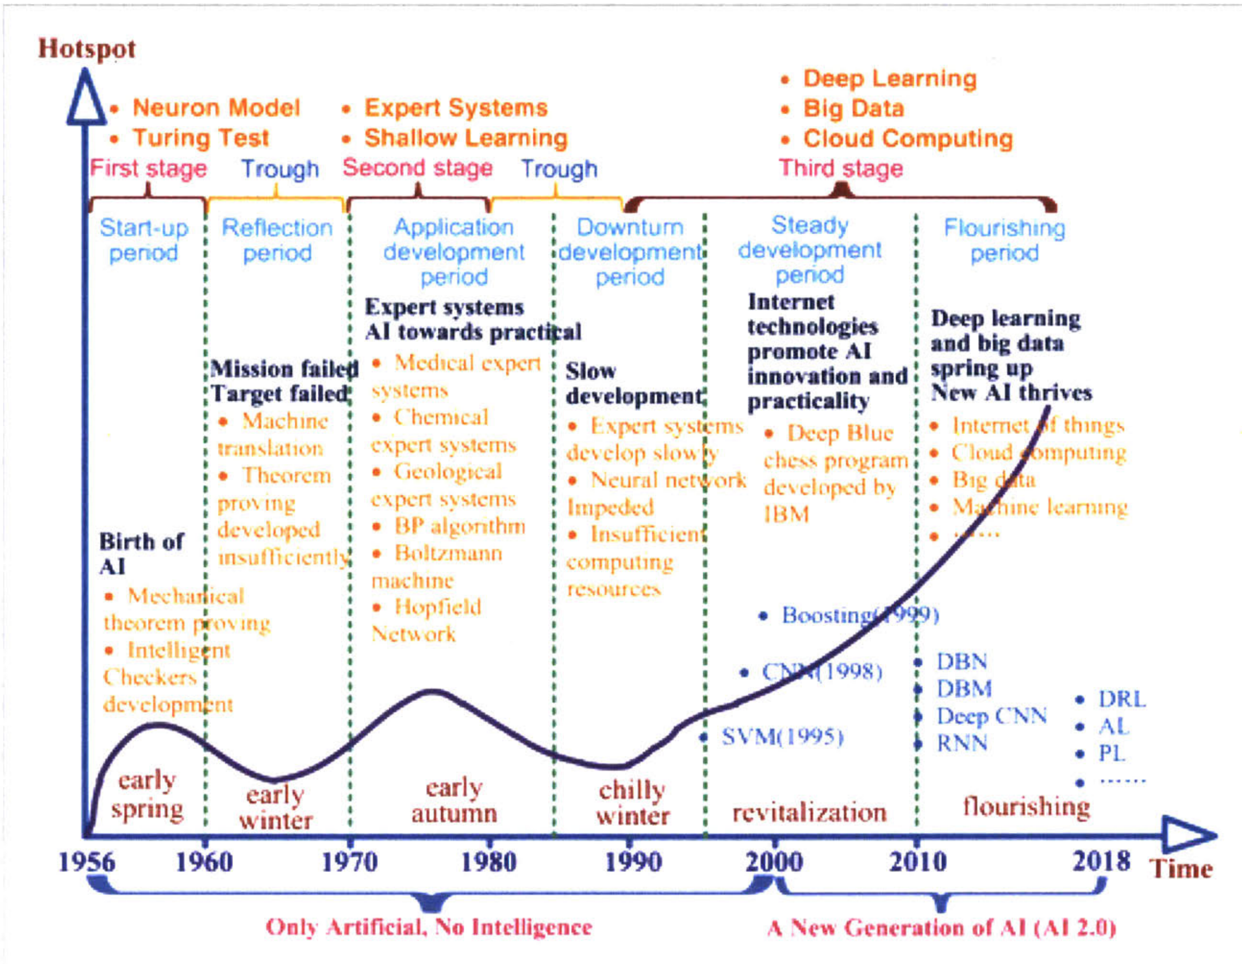

Figure 1: Process of Artificial Intelligence Evolution

(Source: https://onlinelibrary.wiley.com/doi/full/10.1002/er.4333#er4333-fig-0002)

The first stage of Artificial Intelligence research lasted from the mid 1950s to the 1960s. At this stage, the concept of AI was proposed, and the main focus was on machine translation of logical reasoning. The symbolism developed rapidly and Expert Systems and Knowledge Engineering formed a bulk of mainstream research. After 1960, AI research experienced several ups and downs. Gradually the realization dawned that logical reasoning ability alone was insufficient to reach the level of machine intelligence.

The second stage of AI research lasted from 1970s to 1980s. At this stage, scholars considered whether knowledge could be summarized and then handed over to a computer system. Scientist began to program the knowledge of human experts into computer systems to produce "Expert Systems" that could be used to solve actual problems [4.2]. However, due to high development costs, such investigations were slow to develop. Thus, AI development entered another trough. At this stage, researchers hoped to model and mine the implicit knowledge in datasets through Machine Learning techniques. Therefore, the mainstream research of AI entered the third stage, which continues today.

The third stage of AI research began in the 1990s. The Machine Learning discipline emerged to solve the bottlenecks in knowledge acquisition. In May 1997, IBM's Deep Blue computer defeated a chess master Garry Kasparov for the first time, sparking a rush into AI research [4.3]. After entering the 21st century, the advances in Deep Learning and Big Data technology brought about a new peak in AI research, called a new generation of AI. The field of Artificial Intelligence today mainly investigates the theory and method of how to use computers to analyze data and can be considered as the theory and methods for intelligent data analysis.

In 2006, AI began to enter the era of cognitive intelligence that emphasizes the accumulation of big data, innovations in theoretical algorithms, improvements in computing power and autonomous learning. In 2016, Google's AlphaGo [4.4] defeated Sedol Lee with a score of 4:1, indicating a major breakthrough in AI 2.0 after 60 years. In May 2017, AlphaGo defeated Jie Ke, who was ranked first in the world, by a score of 3:0. Currently, AlphaGo has evolved through several successive versions, including AlphaGo Master, AlphaGo Zero and AlphaZero, which can be trained faster, execute faster and have achieved increasing Go proficiency. AlphaZero does not require any artificial features or examples of human Go moves.

4.2. CATEGORIES OF ARTIFICIAL INTELLIGENCE

Artificial Intelligence can be broken down into three main categories:

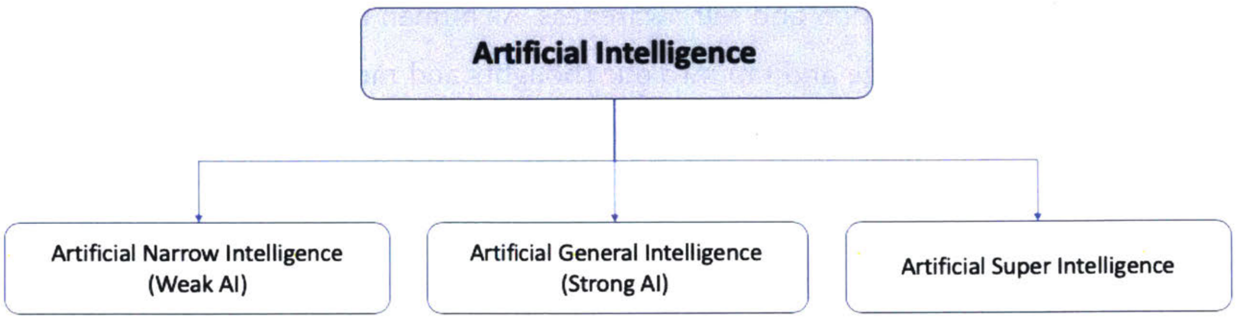

Figure 2: Categories of Artificial Intelligence

Artificial Narrow Intelligence (ANI): Narrow AI or Weak AI is the only form of Artificial Intelligence that humanity has achieved so far. ANI is programmed to perform a single task such as playing chess, checking the weather, making purchase suggestions, weather forecasts and sales predictions [4.5]. ANI systems can attend to task in real-time, but they pull information from a specific data-set. Narrow AI operates within a pre-determined, pre-defined range, even if it appears to be much more sophisticated than that. Every sort of machine intelligence that surrounds us today is Narrow AI. Computer Vision, Natural Language Processing are still at the current stage Narrow AI. Google Assistant, Google Translate, Siri are examples of Narrow AI. Self-driving car technology is still considered a type of Narrow AI, or more precisely, a coordination of several narrow AIs.

Narrow AI systems are able to process data and complete tasks at a significantly quicker pace than any human can, which has enabled us to improve our overall productivity and efficiency. From automating mundane tasks to ordering groceries online to crunch incredible amount of data to produce results, Narrow AI has made our lives significantly better. ANI systems act as the building blocks of more intelligent AI that we might encounter in the near future.

Artificial General Intelligence (AGI): General AI or Strong AI refers to machine that exhibit human intelligence. It is the type of AI that can understand and reason its environment as a human would [4.5]. These would represent systems that are conscious, sentient and driven by emotion and self-awareness. As human beings, we have the ability to think abstractly, strategize and tap into our thoughts and memories to make informed decision or come up with creative ideas. This is something very difficult to replicate in machines.

Some say we might see General AI before the turn of the century. AGI is expected to be able to reason, solve problems, make judgements under uncertainty, plan, learn, integrate prior knowledge in decision-making. AGI systems will be innovative, imaginative and creative.

Artificial Super Intelligence (ASI): According to University of Oxford scholar and AI expert Nick Bostrom [4.6], "when AI becomes much smarter than the best human brains in practically every field, including scientific creativity, problem solving, general wisdom and social skills, we would have achieved Artificial Super Intelligence".

This is the type of AI that many people are worried about, and the type of AI that people like Elon Musk and Stephen Hawking think will lead to the potential end of humanity. Others believe that the smarter AI gets, the better humans will become at saving the environment, curing diseases, explore the universe and at understanding ourselves.

4.3. TYPES OF ARTIFICIAL INTELLIGENCE

Artificial Intelligence is the overarching discipline that covers anything related to making machines smart. Whenever a machine completes tasks based on a set of stipulated rules that solve problems, such an intelligent behaviour is called Artificial Intelligence.

Machine Learning is a subset of AI. That is, all Machine Learning counts as AI, but not all AI counts as ML. Rules Engines, Expert Systems and Knowledge Graphs could all be described as AI, but none of them are Machine Learning. One aspect that separates ML

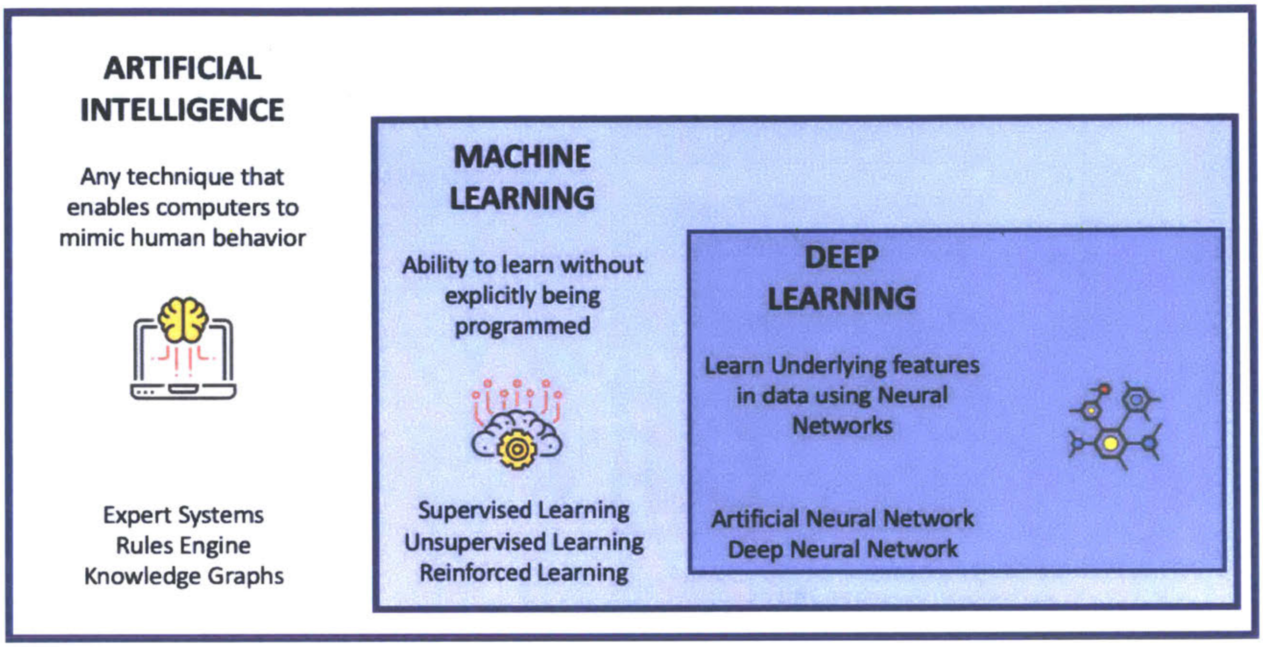

Figure 3: Types of Artificial Intelligence

from Knowledge Graphs and Expert Systems is its ability to modify itself when exposed to more data. Machine Learning is dynamic and doesn't not require human intervention to make certain changes.

Deep Learning is a subset of Machine Learning. DL algorithms are inspired by the information processing patterns found in the human brain. Just like we use our brains to identify patterns and classify various types of information, deep learning algorithms can be taught to accomplish the same tasks for machines. Deep Learning uses multi-layered Artificial Neural Networks to deliver high accuracy in tasks such as object detection, speech recognition and language translation.

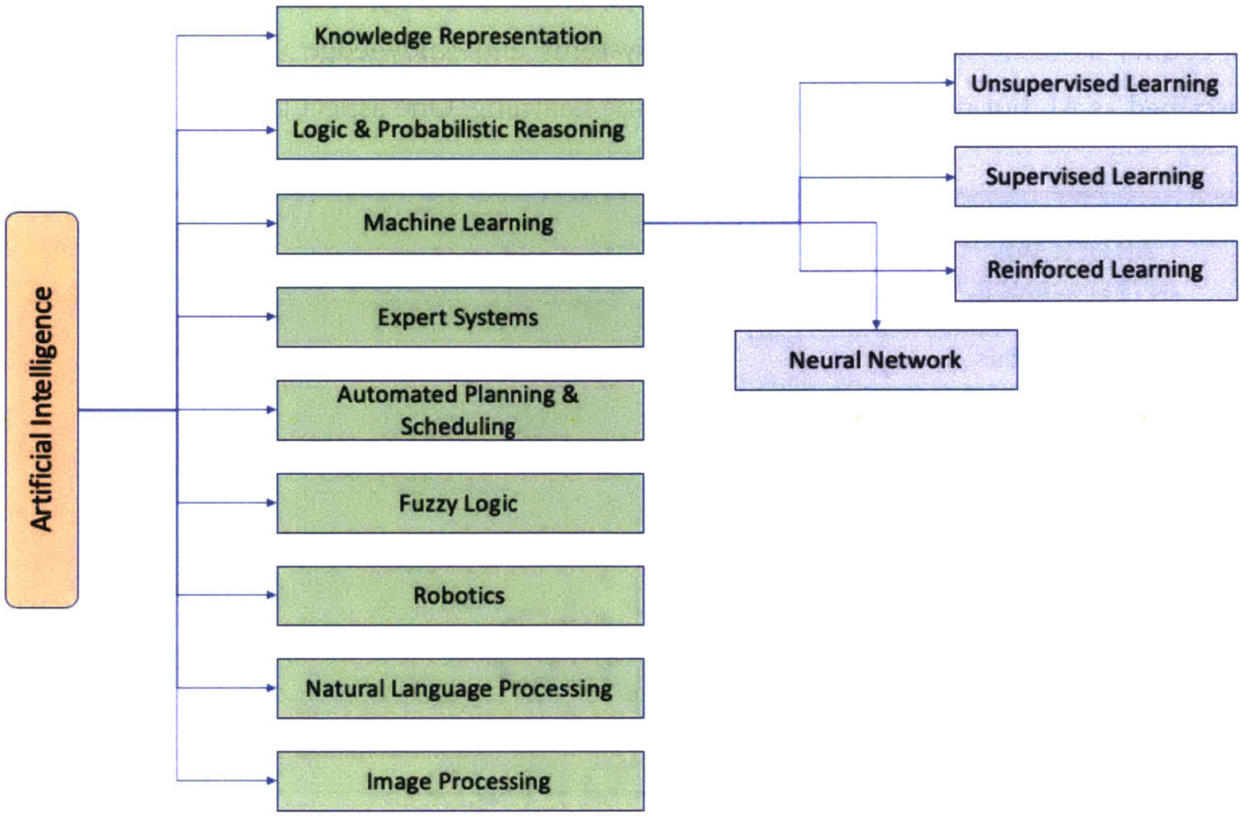

Figure 4: Sub-Fields of Artificial Intelligence Research

Knowledge Representation: It is a field of Artificial Intelligence dedicated to representing information about the world in a form that a computer system can utilize to solve complex tasks [4.7]. The justification for knowledge representation is that conventional procedural code is not the best formalism to use to solve complex problems. Knowledge representation makes complex software easier to define and maintain than procedural code. Knowledge bases are used to model application domains and to facilitate access to stored information.

Logical & Probabilistic Reasoning Systems: It is a system that generates conclusions from available knowledge using logical techniques such as deduction and induction as well as probabilistic techniques such as deductive logic to exploit structure of formal argument.

Reasoning systems come in two modes: interactive and batch processing [4.8]. Interactive systems interface with the user to ask clarifying questions or otherwise allow the user to guide the reasoning process. Batch systems take in all the available information at once and generate the best answer possible without user feedback or guidance.

Expert Systems: An Expert System is a computer software that can solve a narrowly defined set of problems using information and reasoning techniques normally associated with a human expert. It could also be viewed as a computer system that performs at or near the level of human expert in a particular field of endeavor. Rule based Expert Systems try to solve complex problems by implementing series of "if-then-else" rules. One advantage to such systems is that their instructions (what the program should do when it sees "if" or "else") are flexible and can be modified either by the coder, user or program itself. Such expert systems were created and used in the 1970s by Feigenbaum and his colleagues [4.2], and many of them constitute the foundation blocks for AI systems today.

Fuzzy Logic Systems: Fuzzy Logic is a method of reasoning that resembles human reasoning. The approach of Fuzzy Logic system imitates the way of decision making in humans that involves all intermediate possibilities between digital values Yes and No [4.9]. The conventional logic block that a computer can understand takes precise input and produces a definite output as True or False, which is equivalent to human's Yes or No. Fuzzy Logic systems can vary in size and capabilities ranging from small micro-controllers to large, networked, workstation-based control systems.

4.4. MACHINE LEARNING

Arthur Samuel, a pioneer in the field of Artificial Intelligence and computer gaming, coined the term 'Machine Learning'. He defined Machine Learning [4.10] as "Field of study that gives computers the capability to learn without being explicitly programmed". Samuel taught a computer program to play checkers. His goal was to teach it to play checkers better than himself, which is not something he could program explicitly. He succeeded, and in 1962 his program beat the checkers champion of the state of Connecticut.

Machine Learning can be explained as automating and improving the learning process of computers based on their experiences without any human assistance. The process starts with feeding a good quality data and then training the computers by building machine learning models using the data and different algorithms. The choice of algorithms depends on what type of data do we have and what kind of task we are trying to automate.

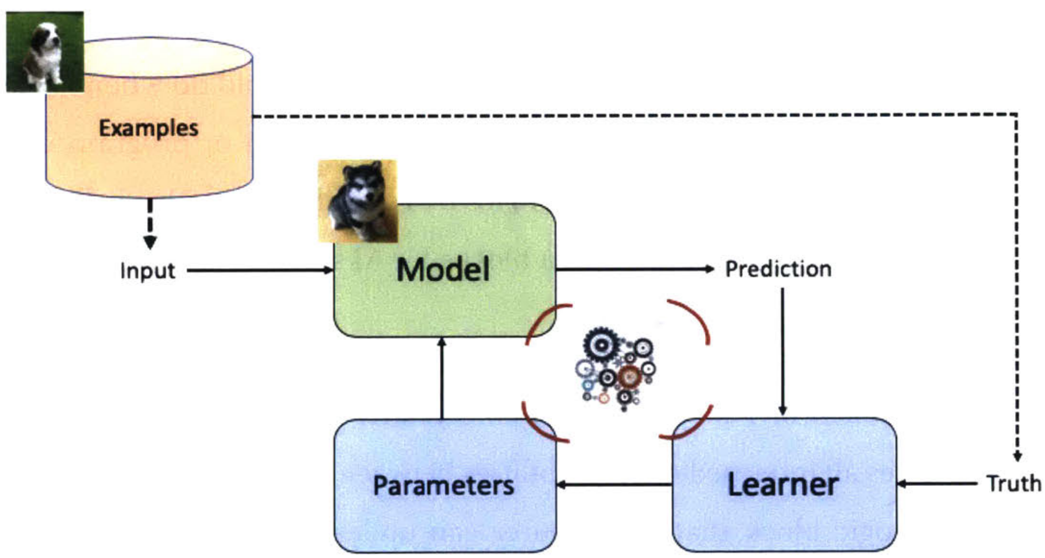

Figure 5: How Machine Learning Works

(Adapted from https://martechtoday.com/how-machine-learning-works-150366)

Machine Learning systems are made up of three major components:

- **Model**: The system that makes predictions or identifications.
- **Parameters**: The signals or factors used by the model to form its decisions.
- **Learner**: The system that adjusts the parameters, and in turn the model, by looking at differences in predictions versus actual outcome.

The "learning" part of machine learning means that ML algorithms attempt to optimize along a certain dimension; they usually try to minimize error or maximize the likelihood of their predictions being true. The Model takes parameters in and form computational logic. The parameters are fed into Learner directly for running tasks. The Model is used internally by the Learner. During training, the Learner takes the Parameters, computes loss and optimization rules based on the model. Learner then updates Model Parameters. During testing, Learner passes the data to model to get prediction.

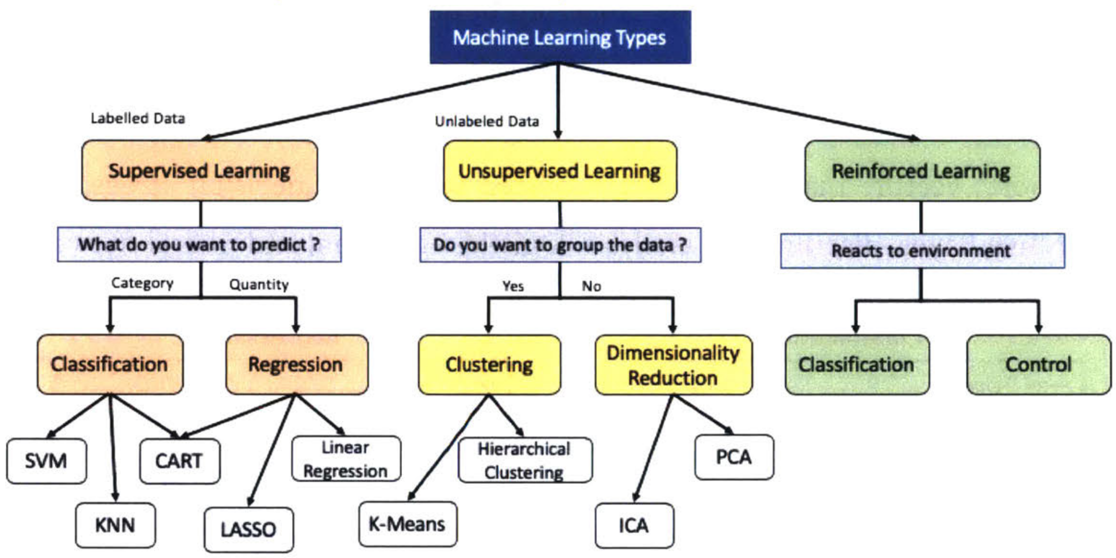

Figure 6: Classification of Machine Learning

4.4.1. SUPERVISED LEARNING

Supervised Machine Learning builds a model that makes predictions based on evidence in the presence of uncertainty [2.14]. A supervised learning algorithm takes a known set of input data and known responses to the data (output) and trains a model to generate reasonable predictions for the response to new data. The supervisor feeds some example data about which the supervisor already knows the answers. The supervisor guides the system by tagging the output. For example, a supervised machine learning system that can learn which emails are 'spam' and which are 'not spam'. The algorithm would be first trained with available input data set (of millions of emails) that is already tagged with this classification to help the machine learning system learn the characteristics or parameters of the 'spam' email and distinguish it from those of 'not spam' emails.

Supervised Learning uses Classification and Regression techniques to develop predictive models.

Classification: These techniques are used to categorize a certain observation into a group. These are used to predict discrete responses [4.11]. For example, whether an email is genuine or spam, or whether a tumor is cancerous or benign. Classification models classify input data into categories. Typical applications include medical imaging, speech recognition, and credit scoring. Common algorithms for performing classification include support vector machine (SVM), boosted and bagged decision trees, k-nearest neighbor, Naïve Bayes, discriminant analysis, logistic regression, and neural networks.

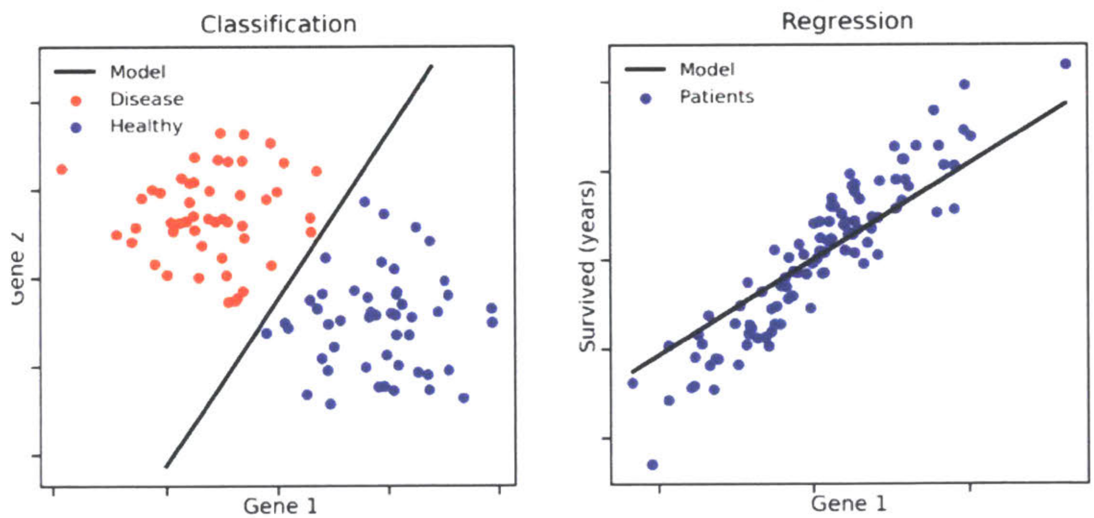

Figure 7: Classification & Regression

(Source: https://aldro61.github.io/microbiome-summer-school-2017/sections/basics/)

Regression: These techniques are used to predict and forecast continuous response values. For example, changes in temperature or fluctuations in power demand [4.11]. Typical applications include electricity load forecasting and algorithmic trading. Common regression algorithms include linear model, nonlinear model, regularization, stepwise regression, boosted and bagged decision trees, neural networks, and adaptive neuro-fuzzy learning.

4.4.2. UNSUPERVISED LEARNING

Unsupervised Learning finds hidden patterns or intrinsic structures in data. It is used to draw inferences from datasets consisting of input data without labeled responses [4.12]. This is useful in searching for a hidden pattern when there is no example data set with known answers. This type of algorithm tends to restructure the data into something else, such as new features that may represent a class or a new series of un-correlated values. They are quite useful in providing insights into the meaning of data and new useful inputs to supervised machine learning algorithms. As a kind of learning, it resembles the methods humans use to figure out that certain objects or events are from the same class, such as by observing the degree of similarity between objects. Some recommendation systems on the web in the form of marketing automation are based on this type of learning.

Clustering and Association (Dimensionality Reduction) are the two most commonly used Unsupervised Learning techniques.

Clustering: It is the most common unsupervised learning technique. It is used for exploratory data analysis to find hidden patterns or groupings in data [4.12]. Applications for cluster analysis include gene sequence analysis, market research, and object recognition. For example, if a cell phone company wants to optimize the locations where they can build cell phone towers, they can use machine learning to estimate the number of clusters of people relying on their towers. A phone can only talk to one tower at a time, so the team uses clustering algorithms to design the best placement of cell towers to optimize signal reception for groups, or clusters, of their customers. Common algorithms for performing clustering include k-means and k-medoids, hierarchical clustering, Gaussian mixture models, hidden Markov models, self-organizing maps, fuzzy c-means clustering, and subtractive clustering.

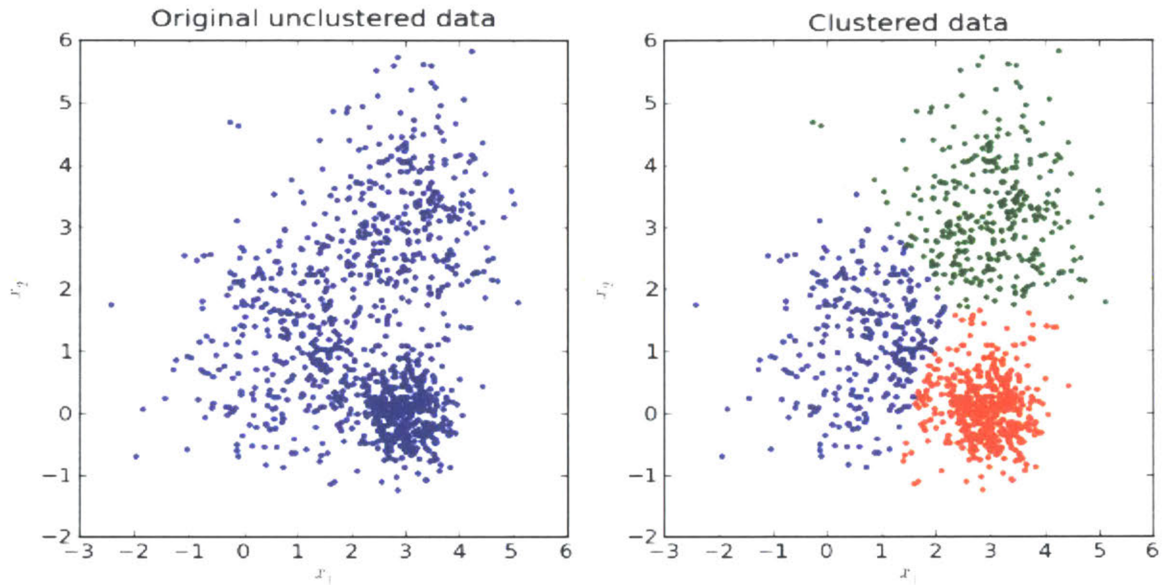

Figure 8: Computer Generated Program Showing K-Means Clustering

(Source: https://brilliant.org/wiki/k-means-clustering/)

Association: It is used to find interesting associations (relationships, dependencies) in large sets of data items [4.12]. The items are stored in the form of transactions that can be generated by an external process, or extracted from relational databases or data warehouses. Association is used to discover the exact rules that will describe the large portions of data. For example, people who buy X are also the ones who tend to buy Y. This is used in building movie or product recommendation algorithms.

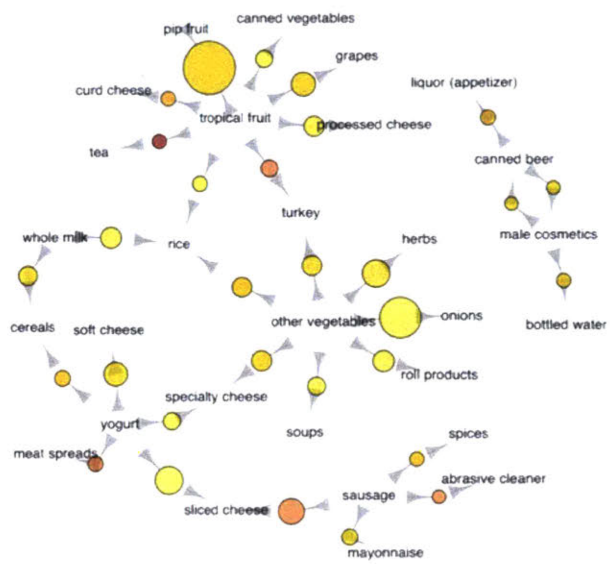

Figure 9: Associations Between Selected Grocery Items

(Source: https://www.kdnuggets.com/2016/04/association-rules-apriori-algorithm-tutorial.html)

4.4.3. REINFORCED LEARNING

Reinforcement learning is the training of machine learning models to make a sequence of decisions. The agent learns to achieve a goal in an uncertain, potentially complex environment [4.13]. In reinforcement learning, an artificial intelligence faces a game-like situation. The computer employs trial and error to come up with a solution to the problem. To get the machine to do what the programmer wants, the AI algorithm gets either rewards or penalties for the actions it performs. Its goal is to maximize the total reward.

Although the designer sets the reward policy, which is the rules of the game, he gives the model no hints or suggestions for how to solve the game. It's up to the model to figure out how to perform the task to maximize the reward, starting from totally random trials and finishing with sophisticated tactics and superhuman skills. By leveraging the power of search and many trials, reinforcement learning is currently the most effective way to hint machine's creativity. In contrast to human beings, artificial intelligence can gather experience from thousands of parallel gameplays if a reinforcement learning algorithm is run on a sufficiently powerful computer infrastructure.

There are two types of Reinforcement.

Positive Reinforcement: Positive Reinforcement is defined as when an event, occurs due to a particular behavior, increases the strength and the frequency of the behavior. In other words, it has a positive effect on the behavior.

Advantages of reinforcement learning are:

- Maximizes performance
- Sustain change for a long period of time

Disadvantages of reinforcement learning:

- Too much reinforcement can lead to overload of states which can diminish the results

Negative Reinforcement: Negative Reinforcement is defined as strengthening of a behavior because a negative condition is stopped or avoided.

Advantages of reinforcement learning:

- Increases behavior
- Provide defiance to minimum standard of performance

Disadvantages of reinforcement learning:

- It only provides enough to meet up the minimum behavior

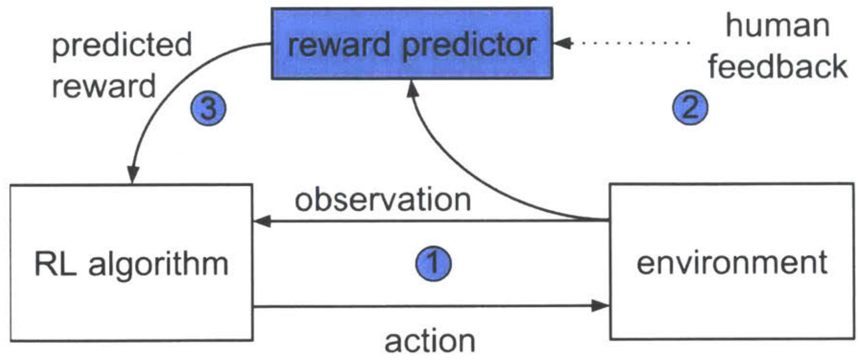

Figure 10: Reinforced Learning

(Source: https://deepmind.com/blog/learning-through-human-feedback/)

4.5. DEEP LEARNING

Deep Learning is a subset of Machine Learning; in fact, deep learning is a technique for realizing Machine Learning [4.14]. Deep Learning was inspired by the structure and function of the brain, namely the interconnecting network of neurons. Our brains are very complex networks with about 10 billion neuron each connected to 10 thousand other neurons. Each of these neurons receives electro-chemical signals and passes these messages to other neurons. Artificial Neural Networks (ANN), modelled after the human brain, are created using layers of artificial neurons to receive input and apply an activation function along with a human set threshold.

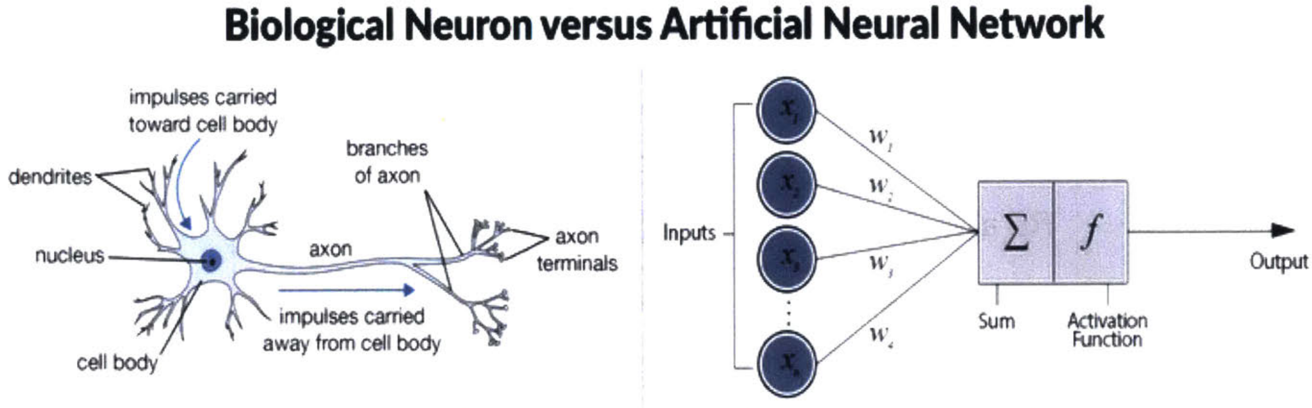

Figure 11: Artificial Neural Network

(Source: https://www.datacamp.com/community/tutorials/deep-learning-python)

4.5.1. ARTIFICIAL NEURAL NETWORK

The simplest neural network is the "perceptron", which, in its simplest form, consists of a single neuron [4.15]. There are five components to artificial neurons:

- **Input nodes**: Each input node is associated with a numerical value, which can be any real number.
- **Connections**: Similarly, each connection that departs from the input node has a weight (w) associated with it and this can be any real number. The ANN runs and propagates millions of times to optimize these "w" values. High computational power is needed to make this in short time.
- **Weighted Sum**: Next, all the values of the input nodes and weights of the connections are brought together. They are used as inputs for a weighted sum.
- **Activation Function**: This result will be the input for a transfer or activation function. Just like a biological neuron only fires when a certain threshold is exceeded, the artificial neuron will also only fire when the sum of the inputs exceeds a threshold. These are parameters set by us.

- **Output Node**: As a result, you have the output node, which is associated with the function of the weighted sum of the input nodes.

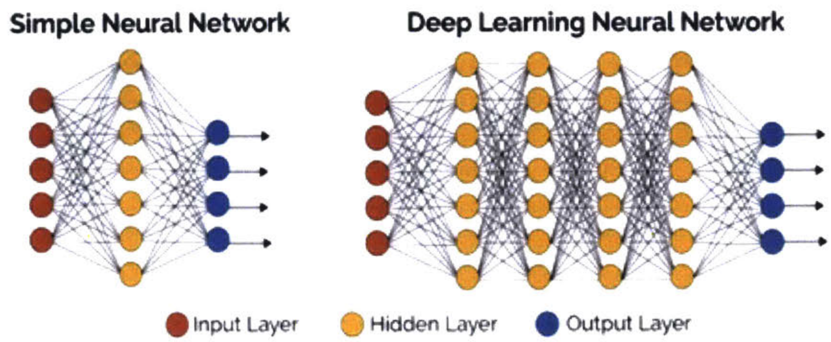

Figure 12: Deep Neural Network

(Source: https://verneglobal.com/blog/deep-learning-at-scale)

The term "Deep" usually refers to the number of hidden layers in the Neural Network. Traditional neural networks only contain 2-3 hidden layers, while deep networks can have as many as 150. Among the layers, there are distinct input layer, hidden layers and an output layer.

4.5.2. CONVOLUTIONAL NEURAL NETWORKS

Convolutional neural networks are deep artificial neural networks that are used primarily to classify images (naming of images), cluster them by similarity (photo search), and perform object recognition within scenes [4.16]. They are algorithms that can identify faces, individuals, street signs, tumors, features and many other aspects of visual data.

Convolutional networks perform optical character recognition (OCR) to digitize text and make natural-language processing possible on analog and hand-written documents, where the images are symbols to be transcribed. CNNs can also be applied to sound when it is represented visually as a spectrogram. More recently, convolutional networks have been applied directly to text analytics as well as graph data with graph convolutional networks.

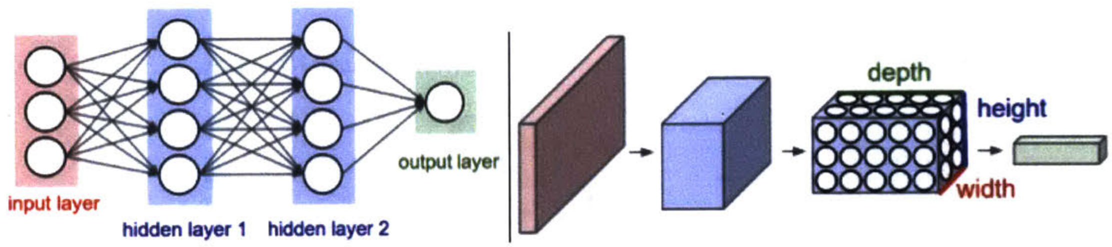

Left: A regular 3-layer Neural Network. Right: A ConvNet arranges its neurons in three dimensions (width, height, depth), as visualized in one of the layers. Every layer of a ConvNet transforms the 3D input volume to a 3D output volume of neuron activations. In this example, the red input layer holds the image, so its width and height would be the dimensions of the image, and the depth would be 3 (Red, Green, Blue channels).

Figure 13: Convolution Network Architecture

(Source: http://cs231n.github.io/convolutional-networks/)

Convolutional Neural Networks are a bit different. First of all, the layers are organized in 3 dimensions: width, height and depth. Further, the neurons in one layer do not connect to all the neurons in the next layer but only to a small region of it. Lastly, the final output will be reduced to a single vector of probability scores, organized along the depth dimension.

The special technique used in CNNs is convolution, where a filter slides over the input and merges the input value and the filter value on the feature map. In the end, the goal is to feed new images to CNN so it can give a probability for the object it thinks it sees or describe an image with text.

4.5.3. RECURRENT NEURAL NETWORK

Recurrent Neural Network (RNN) are the type of Neural Network where the output from previous step are fed as input to the current step [4.16]. In traditional neural networks, all the inputs and outputs are independent of each other, but in cases like when it is required to predict the next word of a sentence, the previous words are required and hence there is a need to remember the previous words. Thus, RNN came into existence, which solved this issue with the help of a Hidden Layer. The main and most important feature of RNN is Hidden state, which remembers some information about a sequence.

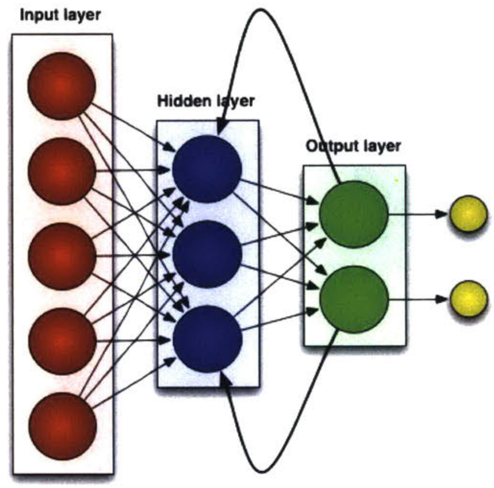

Figure 14: Recursive Neural Network

(Source: https://towardsdatascience.com/notes-on-artificial-intelligence-ai-machine-learning-ml-and-deep-learning-dl-for-56e51a2071c2)

RNN has a "memory" which remembers all information about what has been calculated. It uses the same parameters for each input as it performs the same task on all the inputs or hidden layers to produce the output. This reduces the complexity of parameters, unlike other neural networks. The ability to remember their former inputs gives RNN a big edge over other Artificial Neural Networks when it comes to sequential, context-sensitive tasks such as speech recognition.

4.5.4. GENERATIVE ADVERSARIAL NETWORK

Generative Adversarial Network (GAN) was introduced in a paper by Ian Goodfellow [4.17] and other researchers in 2014. Generative adversarial networks (GANs) are deep neural net architectures comprised of two nets, pitting one against the other, thus the "adversarial".

A generative adversarial network is composed of two neural networks: a Generative Network and a Discriminative Network. In the starting phase, a Generator model takes random noise signals as input and generates a random noisy (fake) image as the output. Gradually with the help of the Discriminator, it starts generating images of a particular class that look real.

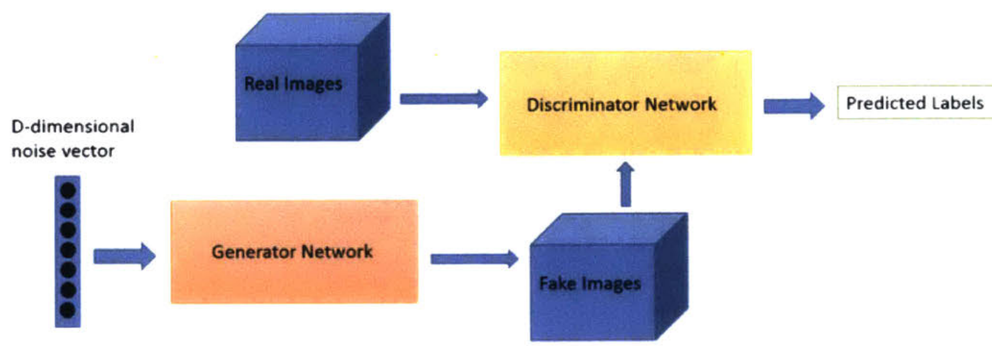

Figure 15: Generative Adversarial Network Architecture
(Source: https://www.oreilly.com/learning/generative-adversarial-networks-for-beginners)

The Discriminator which is the advisory of Generator is fed with both the generated images as well as a certain class of images at the same time, allowing it to tell the generator how the real image looks like. After reaching a certain point, the Discriminator will be unable to tell if the generate image is a real or a fake image, and that is when we can see images of a certain class (class that the discriminator is trained with) being generated by the Generator that never actually existed before.

GANs could be used for increasing the resolution of an image, recreating popular images or paintings or generating an image from text, producing photo realistic depictions of product prototypes, generate realistic speech audio of real people as well as producing fashion/merchandise shots.

CHAPTER 5: ARTIFICIAL INTELLIGENCE – COMMERCIAL PERSPECTIVE

5.1. ARTIFICIAL INTELLIGENCE IN HEALTHCARE

Artificial Intelligence in healthcare represents a collection of multiple technologies enabling machines to manage and analyze data, make decisions and conduct conversations so they can perform administrative and clinical healthcare functions. There are endless opportunities ranging from chronic diseases management to radiology and risk assessment to leverage the transformative power of Artificial Intelligence. According to an analysis done by Accenture [5.1], here is a list of AI application in healthcare in terms of estimated potential annual benefits by 2026.

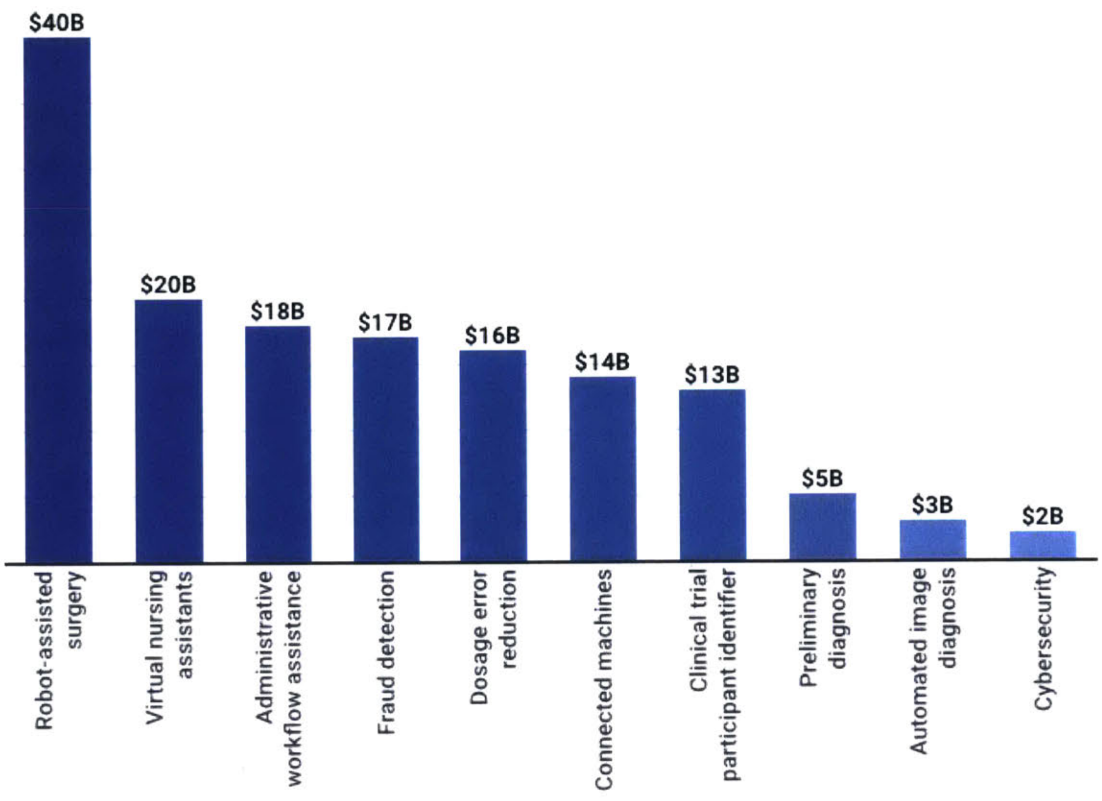

Figure 16: Most promising AI Healthcare Applications

(Source: https://www.accenture.com/_acnmedia/PDF-49/Accenture-Health-Artificial-Intelligence.pdf#zoom=50)

5.1.1. NEXT GENERATION RADIOLOGY POWERED BY AI

Radiological images obtained by MRI machines, CT Scanners and X-ray machines offer non-invasive visibility into the inner-workings of the human body. Rapid technological advances in electronics provided improvements from greyscale and still images to real-time moving images. Historically, in radiology practice, trained doctors visually assessed medical images for the detection, characterization and monitoring of diseases. Doctors and pathologists have to examine hundreds of images and tissue slides for traces of abnormalities to diagnose patients.

Machine Learning have demonstrated great potential in supplementing and verifying the work of the clinicians. Artificial Intelligence methods excel at automatically recognizing complex patterns in imaging data and providing quantitative assessments of radiographic characteristics. According to Radiology Today [5.2], "Artificial Intelligence technology is already being used in radiology in a number of ways, such as computer-aided detection of cancer, auto-segmentation of organs in 3D postprocessing, natural language processing to facilitate critical results reporting, consultation of best guidelines for recommendations, and quantification and kinetics in postprocessing. However, deep learning algorithms have not made their way into widespread clinical use".

Some Neural Networks are approaching performance levels comparable to trained doctors in detecting malignant and benign growths in images such as mammograms. Deep Neural Networks can now identify pathologies in radiological images such as bone fractures and analyze the white matter lesions on MRI scans of the brain [5.3]. Chest radiographs are one of the most utilized radiological modalities in the world and large datasets are currently available for Machine Learning researchers. A team of researchers from Stanford [5.4] found that Deep Learning models detected clinically important abnormalities on chest radiography at a performance level comparable to practicing radiologists. The Deep Learning algorithm called CheXNext, can evaluate X-rays for a multitude of possible abnormalities and return results that are consistent with that of a radiologist.

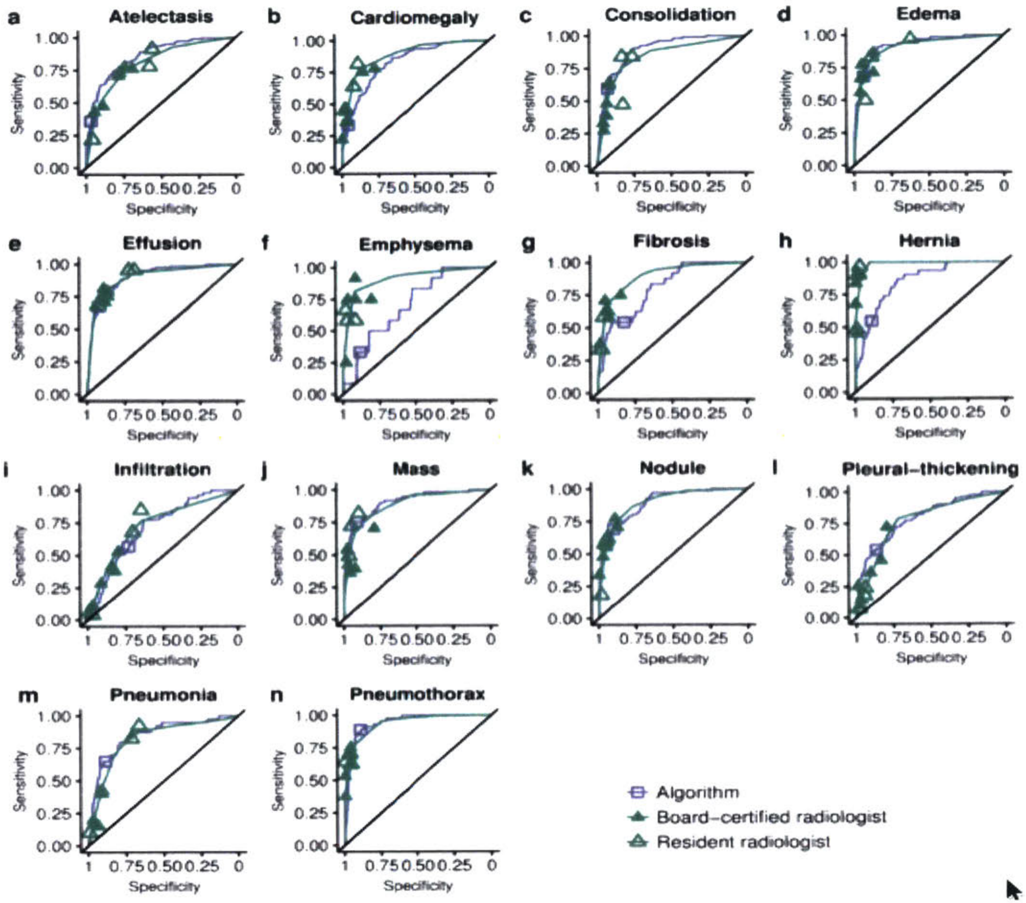

Figure 17: Comparison of CheXNeXt's Performance to that of 9 Radiologists (Source: https://stanfordmlgroup.github.io/projects/chexnext/)

5.1.2. OPHTHALMOLOGY AND AI

Early detection and appropriate treatment of eye diseases are of great significance to prevent vision loss. Traditionally, diagnosis methods tremendously depend on the professional experience and knowledge of doctors (ophthalmologist). The four leading causes of adult blindness are diabetic retinopathy (DR), glaucoma, age-related macular degeneration (AMD), and cataract [5.5].

Diabetic Retinopathy is a leading cause for blindness. Diabetic patients are at very high risk for Diabetic Retinopathy and don't notice changes in their vision in the early stages of the disease. As it progresses, DR usually causes vision loss that in many cases cannot be reversed. Hence it is important that people with diabetes have yearly screenings. One of the most common ways to detect diabetic eye disease is to have an eye specialist examine the pictures of the back of the eye. The eye specialist then rates these pictures for the presence of the disease and severity, determined by the type of lesions present.

A team of researcher from Google AI developed a Deep Learning algorithm with excellent diagnostic performance. A specific type of Neural Network optimized for image classification called a Deep Convolutional Neural Network was trained using a dataset of 128175 retinal images. The dataset was used to train the algorithm to detect referable diabetic retinopathy. The performance of the algorithm is on par with trained ophthalmologists [5.6].

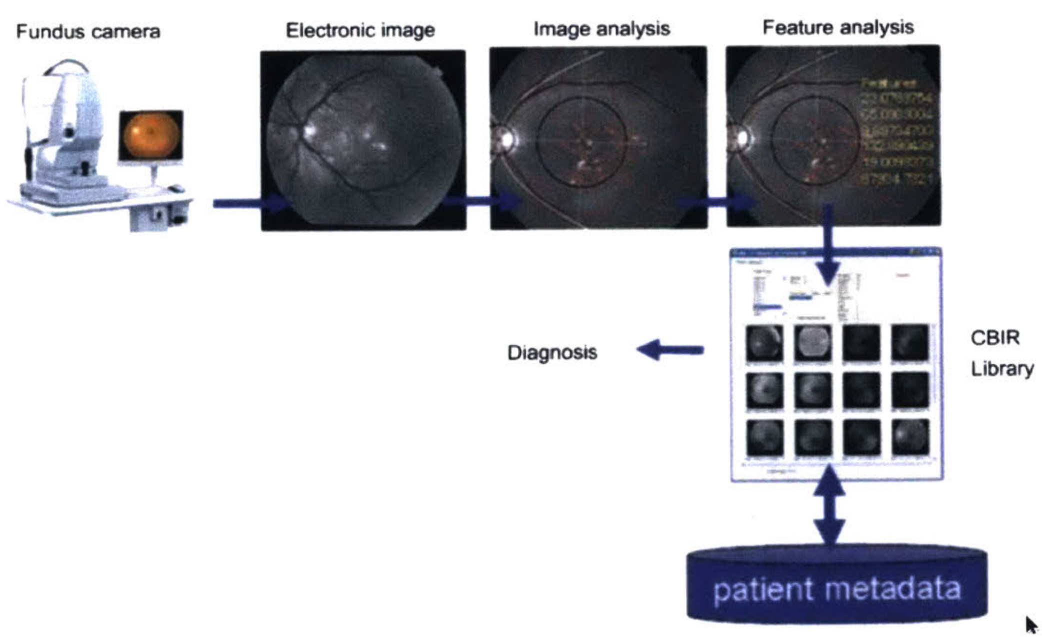

Figure 18: Deep Learning Algorithm for Detecting

(Source: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6133903/)

Refractive error such as myopia and hyperopia are some of the most common visual impairments in the world. If left untreated, refractive errors could become pathological. Microsoft in partnership with LV Prasad Eye Institute in India, has developed a Machine Learning algorithm for prediction and progression of refractive error in children and young adults [5.7]. The algorithm developed by Microsoft is currently being validated in Electronic Medical Record systems at 174 centres of LV Prasad Eye Institute in India.

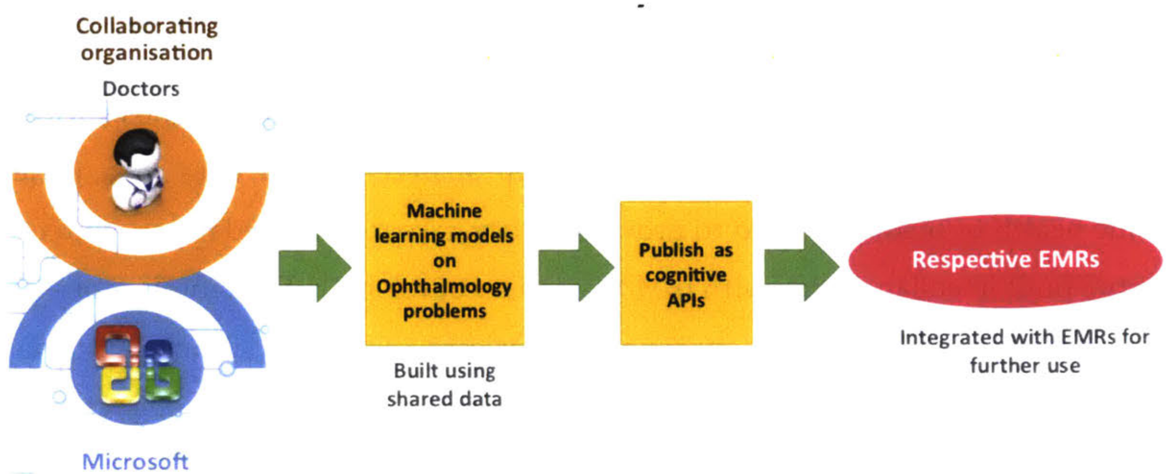

Figure 19: Microsoft Intelligent Network for Eyecare (MINE)

(Source:https://www.visionmonday.com/CMSDocuments/2017/04/VM%20Summit_march%202017_final%20version_2.pdf)

The algorithm has been trained on a dataset of 335,799 instances from 176,037 patients aged between 0 and 25 years. The model has proved to be statistically accurate and provided optimal results when used to predict cases of myopia that could become severe and cause a change in the value of spherical power in the eye of children. The ultimate aim of this program is to help eyecare providers with a tool to screen patients who are at high risk of eventual blindness and implement anti-blindness strategies at scale to prevent loss of vision.

5.1.3. TELEHEALTH MEETS AI

According to the Health Resources and Services Administration (HRSA) of the US Department of Health and Human Services, Telehealth is defined as the "use of electronic information and telecommunications technologies to support and promote long-distance clinical health care, patient and professional health-related education, public health and health administration Technologies include videoconferencing, the internet, store-and-forward imaging, streaming media, and terrestrial and wireless communications [5.8]."

Infusing Machine Learning into Telehealth platform will mean better diagnosis with less human effort. Machine Learning algorithms help analyze clinical data in a patient's electronic health or medical record to provide patient care recommendations. iDirectDoc is a startup built in collaboration with IBM Watson, to develop Artificial Intelligence based mobile telemedicine solution that manages patient care so that physicians can focus on delivering care [5.9]. The Artificial Intelligence algorithm deployed by iDirectDoc can derive insights from patient-generated data and medical research. The platform is helping to connect physicians with patients for everything from routine checkups and home health services to stroke assessments and surgical consultations.

5.2. ARTIFICIAL INTELLIGENCE IN EDUCATION

While the role of teachers in the education system is irreplaceable, Artificial Intelligence can be an important tool to make education more accessible and personalized. Artificial Intelligence can automate and expedite a lot of administrative work, thus freeing up teachers for more valuable work. Natural Language Processing algorithms can be used in education as tutors, Machine Learning algorithms can be used to grade student's work and Artificial Intelligence can power complicated virtual reality simulations to help students learn.

5.2.1. PERSONALIZED LEARNING

Application of Artificial Intelligence allows the creation of individualized learning that is impossible for teachers who have to manage a cohort in each class. Century Tech uses AI technology to provide a personalized learning journey to students, and offers real-time insights and analytics to educators [5.10]. The platform utilizes cognitive neuroscience and data analytics to create personalized learning plans and reduce workloads for instructors. The platform tracks student progress, identifies knowledge gaps and offers personal study recommendations and feedback. By having a better understanding of the student, educators are better positioned to utilize their training skillset to address the knowledge requirement of students.

5.2.2. AUGMENTED REALITY FOR THE CLASSROOM

Educators are constantly challenged to find ways to keep students' attention in the classroom. Gamifying the learning experience helps in keeping a student engaged. Augmented Reality and Virtual Reality can be used to increase engagement and shared experience among students. The company Blippar has created products that combine computer vision intelligence technology and augmented reality to enhance the way students learn in the classroom [5.11]. The platform lets students look at real-world objects enhanced with text, music and digital graphics on their smartphone and tablets. The platform also allows teachers to add multiple choice and open-ended questions, characters with speech bubbles, while creating virtual experiences.

5.2.3. ASSITIVE LEARNING

Many companies are harnessing the power of Artificial Intelligence to make life easier for the people with disabilities. As the technology improves, assistive technology will continue producing new platforms to help create a better standard of living for those individual.

Reading and writing has been difficult for individuals with visual impairment, including blindness. Artificial Intelligence has enhanced the quality of life of such people through technology such as Speech-To-Text.

Microsoft has created a ground-breaking Seeing AI app that provides a wide range of capabilities for a person with low vision or blindness [5.12].

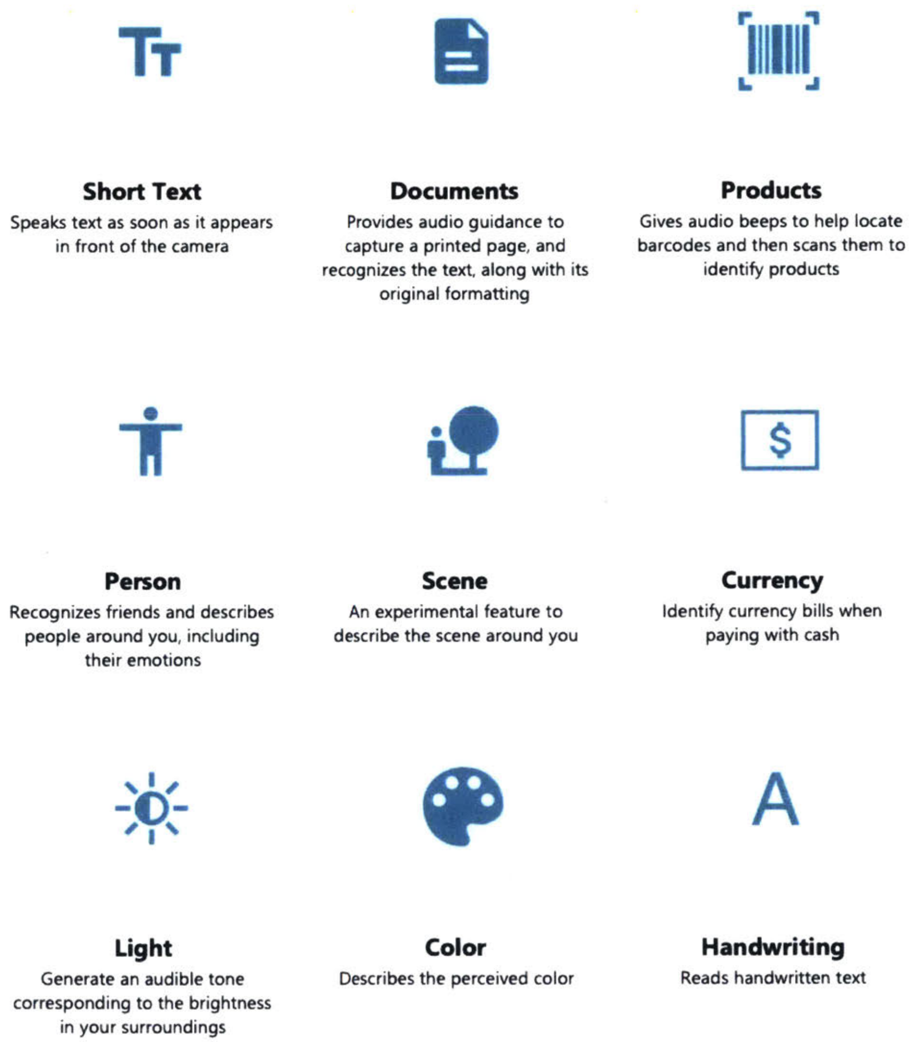

Figure 20: Functions of the Seeing AI App

(Source: https://www.microsoft.com/en-us/seeing-ai)

A feature called the Short Text mode turns on the camera and searches for printed material continuously. When it finds printed text, it translates it into machine text, and then speaks it aloud. The Document mode works like many text recognition apps; it automatically clicks a photo of the document and sends it to Microsoft servers for recognition. The app scans barcodes to identify products and product details such as nutrition labeling, calorie counts and even cooking instructions. The app helps the user to create a photo library of friends and family, and the app will tell the user when someone it recognizes is near. The app will inform the user when a face is in focus and offer up the person's name if they are in the app image database. The scenes mode will help the user to snap a photo of his or her environment and the app will attempt to describe it. The app also acts as a currency identifier.

5.3. ARTIFICIAL INTELLIGENCE IN FINANCE

Artificial Intelligence is taking the Financial services industry by storm. Financial services seeking competitive advantages are utilizing Artificial Intelligence to improve customer service operations, risk assessment and optimize marketing and sales processes. Artificial system can examine billions of data points and find patterns and trends people may miss. Based on historical data, these systems can also predict future patterns. Natural Language Processing can be used to build conversational applications that let customers converse and perform specific actions by chat and voice applications. Artificial Intelligence can be used in processes which involve auditing of financial transactions.

Analysts estimate that Artificial Intelligence will save the banking industry more than $1 trillion by 2030 [5.13]. Analysts predict a savings of $490 billion in Front Office from reduction in the scale of retail branch networks, security, tellers, cashier and distribution staff. Analysts predict a savings of $350 billion in Middle Office by applying Artificial Intelligence to compliance, KYC, authentication and other forms of data processing.

Analysts also predict a savings of $200 billion in Back Office by deploying Artificial Intelligence methods to underwriting and collections systems.

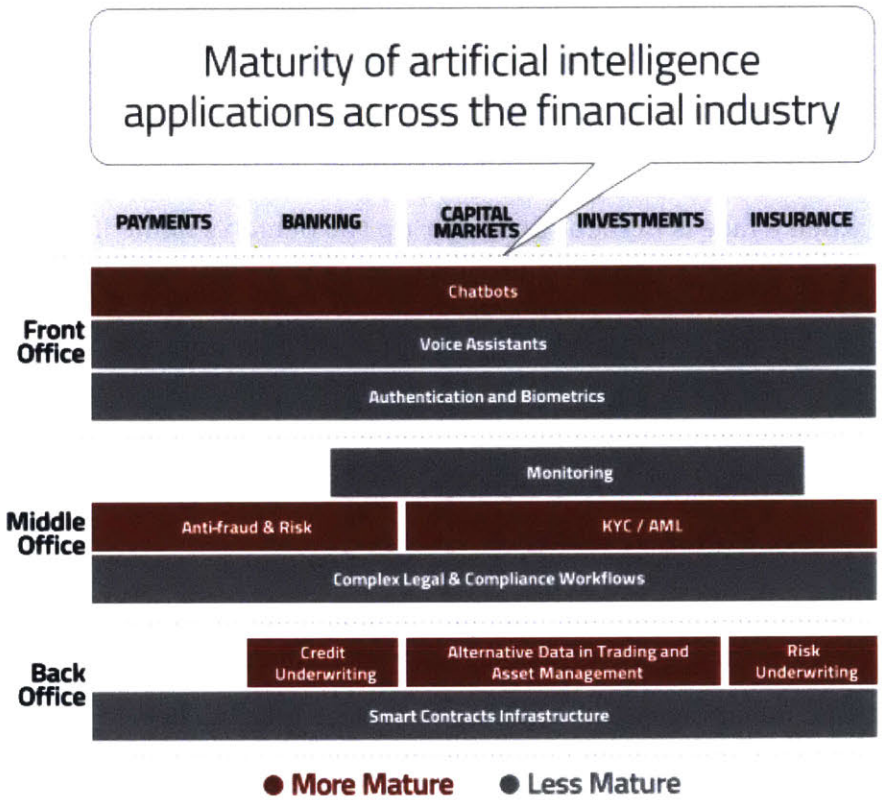

Figure 21: Maturity of AI Application Across Financial Industry

(Source: https://thefinancialbrand.com/72653/artificial-intelligence-trends-banking-industry/)

5.3.1. ALGORITHMIC TRADING

Artificial Intelligence can help manage and automate rules and trading decisions based on algorithms created for managing trading rules. Intelligent Trading systems monitor both structured data (spreadsheets, databases) and unstructured data (social media, news) in a fraction of time. Such trading algorithms can conduct trades autonomously.

The algorithm puts together recommendations for the strongest portfolios depending on a specific investor's short-term and long-term goals. Multiple financial institutions also trust AI to manage their entire portfolios. The predictions for stock performance are more accurate, due to the fact that algorithms can test trading systems based on past data and bring the validation process to a whole new level before pushing it live. The company Kavout [5.14] uses Machine Learning and quantitative analysis to process huge sets of unstructured data and identify real-time patterns in financial markets. The intelligence platform generates K Score, which is an AI powered stock ranker. The K score analyzes vast amounts of data including company fundamentals such as SEC filings, technical factors such as price patters and alternate data such as sentiment for its analysis. Deep Learning methods are applied to capture nonlinear insights and correlations. Higher the K score, better is the performance of the stock.

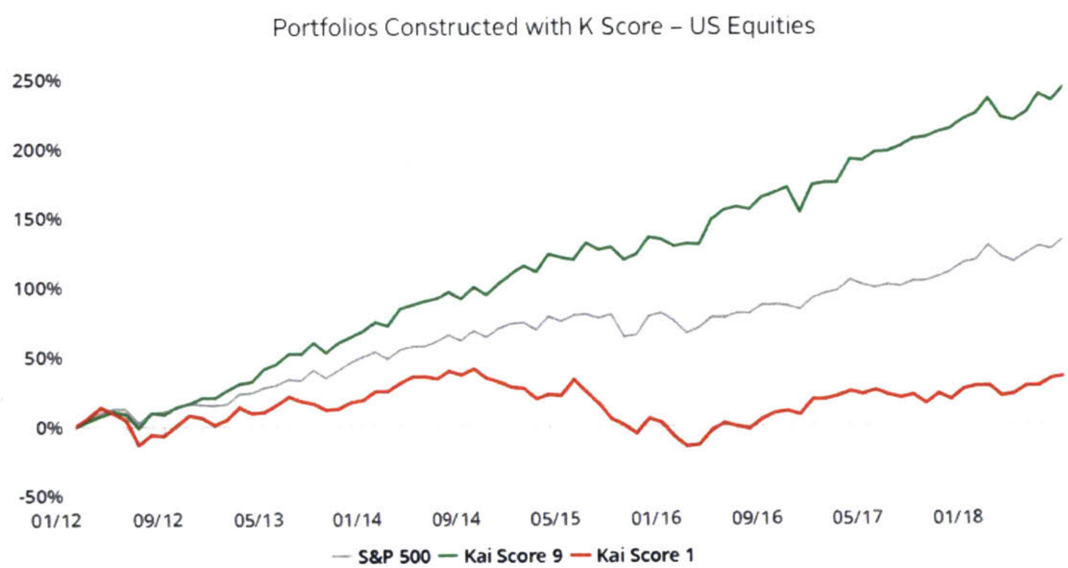

Figure 22: Comparison of Portfolios Constructed Using K Score Against S&P500

(Source: https://www.kavout.com/k-score/)

5.3.2. CREDIT RATING

The access of credit to the unbanked population is very limited in emerging markets. A sizeable chunk of the population conducts cash-based transactions and do not have assets to serve as collateral. Financial service providers in such economies deemed it too risky and costly to provide loans to the unbanked population. But the wealth of data that is being generated from the explosive growth in smartphone penetration is changing that and alternative data points are providing a new means for credit access.

The easy access to credit has the potential to jump-start economic progress in an emerging economy. Artificial Intelligence based credit scoring systems can identify hidden relationships between variables that are not always considered by traditional credit scoring systems which look at one variable at a time. The fintech company ZestFinance [5.15] has created an Artificial Intelligence powered underwriting solution called Zest Automated Machine Learning (ZAML) platform that helps financial agencies determine the credit worthiness of potential borrowers with no credit history. The Machine Learning algorithm was trained on a huge dataset comprising of social media posts, geolocations, browsing activities and other relevant data points to identify which data points correlate to the types of borrowers that more most likely to pay back their loans.

Some of the common Machine Learning methods that are currently used in credit scoring are LOCO, LIME and PI. "LOCO (Leave One Column Out) substitutes the 'missing' data point for a variable and recomputes the model's prediction. If the score changes a lot, it implies that the variable that was left out must be really important. LIME (Local Interpretable Model-agnostic Explanations) fits a new linear model around a given applicant's real data values to explain how the more complex model behaves. PI (Permutation Impact) replaces a variable with a randomly selected value and recalculates the model's prediction. If the score changes a lot, the variable that was replaced must be really important. The proprietary ZAML model derived from game theory and multivariate calculus and determines the relative importance of each variable to the final score by looking at how it interacts with other variables" [5.16].

Does Your Explainability Method...?	Why it Matters	LOCO	LIME	PI	ZAML
Run just as fast as the number of model variables rises.	Speed	✗	✓	✗	✓
Analyze the inner workings of the model, not just inputs and outputs'	Accuracy	✗	✗	✗	✓
Analyze the final model, not a refitted or proxy model.	Accuracy	✗	✗	✓	✓
Analyze in a multivariate manner, not just univariate.	Accuracy	✗	✓	✗	✓
Reflect global not local feature importance.	Accuracy	✓	✗	✓	✓
Make itself user-friendly and not require expert tuning	Consistency	✓	✗	✓	✓
Use actual data, not synthetic or randomly shuffled data?	Consistency	✓	✗	✗	✓
Accuracy in explaining: *					
XGBoost		90%	82%	84%	**99%**
Neural Network		96%	92%	93%	**100%**
Variance (consistency at explaining):					
XGBoost		0.02	0.07	0.03	**0.002**
Neural Network		0.004	0.01	0.009	**0.0002**
Explanation time (sec) *					
XGBoost		0.003	35	0.02	**0.006**
Neural Network		0.005	49	0.04	**0.26**

Figure 23: Evaluation of the ML Model Used by ZestFinance

(Source: https://www.zestfinance.com/snake-oil-explainability)

5.3.3. ROBO ADVISOR

Robo Advisors are Artificial Intelligence powered digital platform that provide automated, algorithm-based financial planning services with little to no human interaction and supervision. These platforms offer personalized financial advice and help individuals achieve their financial goals. These Robo Advisors track income, essential recurring expenses and spending habits to offer an optimized financial plan.

The robo advisor Betterment [5.17] invests and manages individual, IRA and 401K accounts for its users. The Machine Learning algorithm uses information about the user's tax filing status, income, debt, annual household income and investing strategies to provide the user an overview of their financial standing and recommends relevant investment opportunities to the user. The company also uses Artificial Intelligence to select lots to sell based in part on the tax consequences of the transaction, which is aimed at improving tax impact for transactions [5.18].

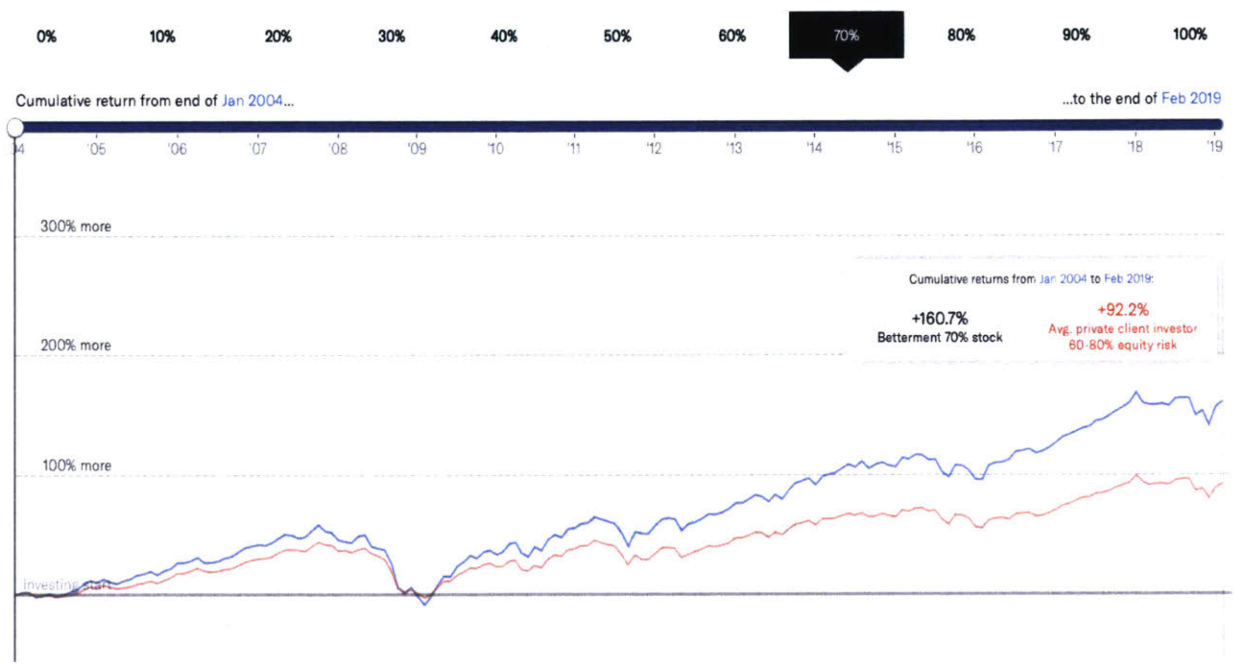

Figure 24: Performance of Betterment Stock Allocation

(Source: https://www.betterment.com/portfolio/)

According to Betterment, "Investing in a Betterment 70% stock portfolio since the end of Jan 2004 through the end of Feb 2019 would have produced a cumulative return of +160.7%, which is the same as an average annual return of +6.6%. Assuming similar risk, the cumulative return would have differed from the average private client investor by +68.6%".

5.3.4. FRAUD DETECTION

The growth in new business models from instant delivery of goods to digital cash to online downloads has also led to a corresponding growth in online fraud and fraud losses. Payment fraud prevention is one of the ideal use cases of Artificial Intelligence. Machine Learning helps data scientists efficiently determine which transactions are most likely to be fraudulent, while significantly reducing false positives. The techniques are extremely effective in fraud prevention and detection, as they allow for the automated discovery of patterns across large volumes of streaming transactions. Similarly, Neural Networks can be trained to use a set of comparative algorithms to distinguish good checks from anomalous ones to prevent check frauds.

The company Vectra [5.19] has created an AI powered platform that automates threat detection, reveals hidden attackers specifically targeting financial institutions, accelerates investigations after incidents, and even identifies compromised information. According to the company "Sensors are deployed across enterprise, cloud and data center environments, where they extract relevant metadata from traffic and ingest external threat intelligence and Active Directory and DHCP logs. The characteristics of every flow are recorded, including the ebb and flow, timing, traffic direction, and size of packets. Security researchers build and continually tune scores of self-learning behavioral models that enrich the metadata with machine learning-derived security information. Scores of custom-built attacker behavior models detect threats automatically and in real time, and provides the most relevant information by augmenting detections with actionable context" [5.20].

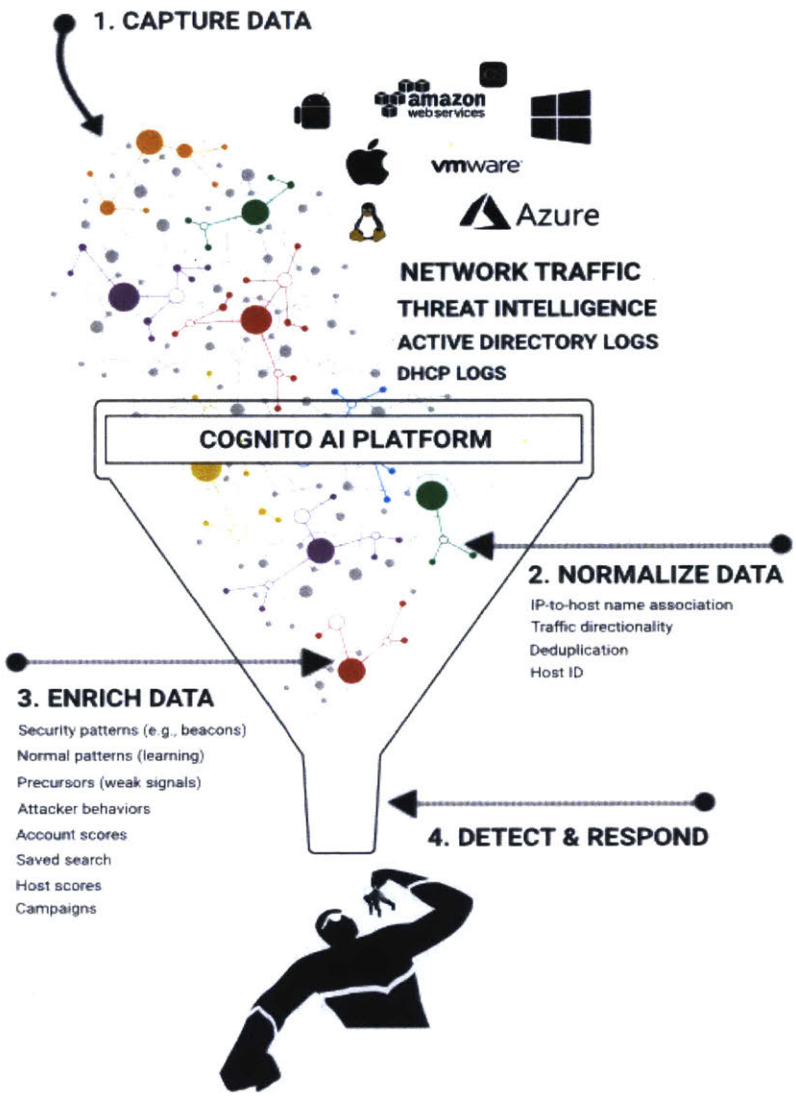

Figure 25: The COGNITO AI Platform for Fraud Detection

(Source: https://www.vectra.ai/product/how-it-works)

5.4. ARTIFICIAL INTELLIGENCE IN MEDIA AND ECOMMERCE

A lot of data being generated today in the Media and Ecommerce industry today is dedicated to understand the audience or customer. Marketing companies are training Machine Learning algorithms to help develop film trailers and design advertisements. Entertainment providers are using Artificial Intelligence to recommends personalized content based on user data. Ecommerce companies are using Artificial Intelligence to improve speed and efficiency of the search process for the user and to organize visual assets.

5.4.1. VISUAL SEARCH

Consumers are often disappointed with e-commerce experience because the product shown are often irrelevant. Artificial Intelligence methods such as Natural Language Processing is utilized to narrow, contextualized and improve search results for online shoppers. Artificial Intelligence also enables shoppers to discover complimentary products.

The company Zorroa [5.21] offers a platform that allows companies to manage a very large collections of visual assets such as images, videos, etc. and integrates Machine Learning algorithms to perform content searches within large databases.

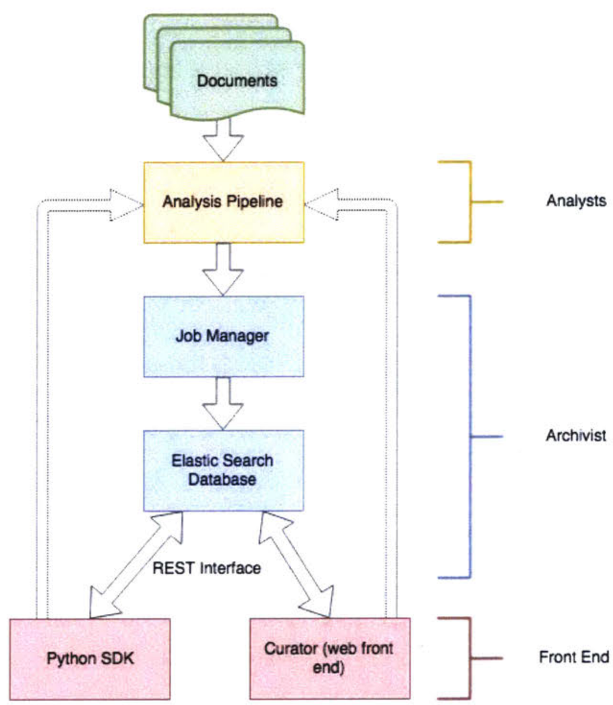

Figure 26: Overview of the Zorroa Platform

(Source: https://zorroa.com/an-introduction-to-the-zorroa-platform/)

Algorithms are trained to recognize specific components of visual content which can then be organized and catalogued to deliver robust search results. Image Classification uses a neural network to classify the image into a set of predefined categories. Face Recognition is used to find faces and decide if they are known, if so tag the asset with a name. Zorroa's architecture allows the deployment of multiple algorithms simultaneously to manage visual assets such as documents, images and video and takes seconds to run intelligent visual searches.

5.4.2. RECOMMENDATION ENGINE

The company Netflix [5.22] uses Machine Learning algorithms to adapt the entire user experience to each individual subscriber, including the rows selected for the homepage, the titles selected for those rows, the visuals for each movie as well as the recommendations of other movies etc. This movie recommendation engine is powered by algorithms that are based on the assumption that similar viewing patterns represent similar user tastes.

For most traditional internet companies, recommendations are precomputed in a batch processing fashion, but such a model cannot react quickly based on member interactions, title interests, popularity etc. Batch processing was not quick enough for time sensitive scenarios such as new title launch campaigns or strong trending popularity cases. The company moved towards a Near Real-Time Recommendation (NRT) recommendation process in order to accelerate the learning process and roll out test results. The company "uses a real-time Spark Streaming ecosystem and leverages real-time data for model training to provide the right personalized videos in a member's account and the right personalized images soon after the launch of a new show."

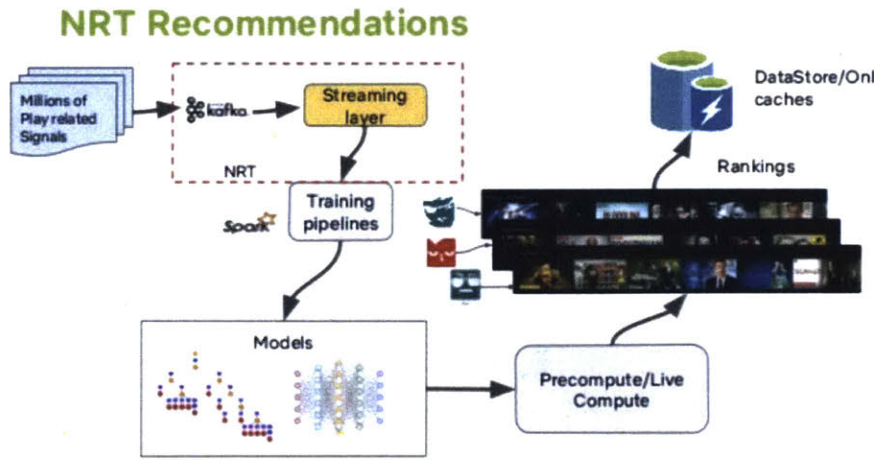

Figure 27: ML Infrastructure for near Real-Time Recommendations (Source: https://medium.com/netflix-techblog/netflix-at-spark-ai-summit-2018-5304749ed7fa)

5.4.3. EFFICIENT SALES PROCESS

One of the biggest problems that an ecommerce retailer faces is that of the abandoned cart. It is estimated that over 70% of online shopping carts are abandoned before customers complete their purchases. Marketing leads are not followed up by the sales team. This means that pre-qualified potential buyers interested in a product are just left behind. Artificial intelligence can help enhance the sales cycle, by creating a strong sales message that reaches consumers at the right time on the right platform. Integrating AI systems that enable Natural Language Processing and voice input, such as Siri, Alexa or Google Assistant allows a CRM system to answer customer queries, solve their problems and even identify new opportunities for the sales team.

The e-commerce retailer The North Face [5.23] employed IBM Watson's text analyzing capabilities as part of its digital strategy. The company uses the Expert Personal Shopper (XPS) software to provide shoppers with product recommendations based on a Q&A with the customers [5.24]. "XPS allows shoppers to have a conversation like they would naturally in a store with an associate. The platform, through a dialogue, helps shoppers find the right products specific to their needs and preferences". The platform provides the customer with a customized list of items, improving efficiency and conversion rate for The North Face.

5.4.4. VIDEO GAMES

The application of Artificial Intelligence in the gaming industry is leading to games with more realistic animations that mimic the natural ways humans move about. Artificial Intelligence in video games is also largely used to determine the behaviour of non-player characters in games. Neural Networks are being developed to create avatars that can move through complex landscapes in new and unexpected ways, while games themselves learn from and adapt to the actions individual players take.

An advanced Artificial Intelligence method used to enhance the personalized gaming experience is called the Monte Carlo Search Tree [5.25].

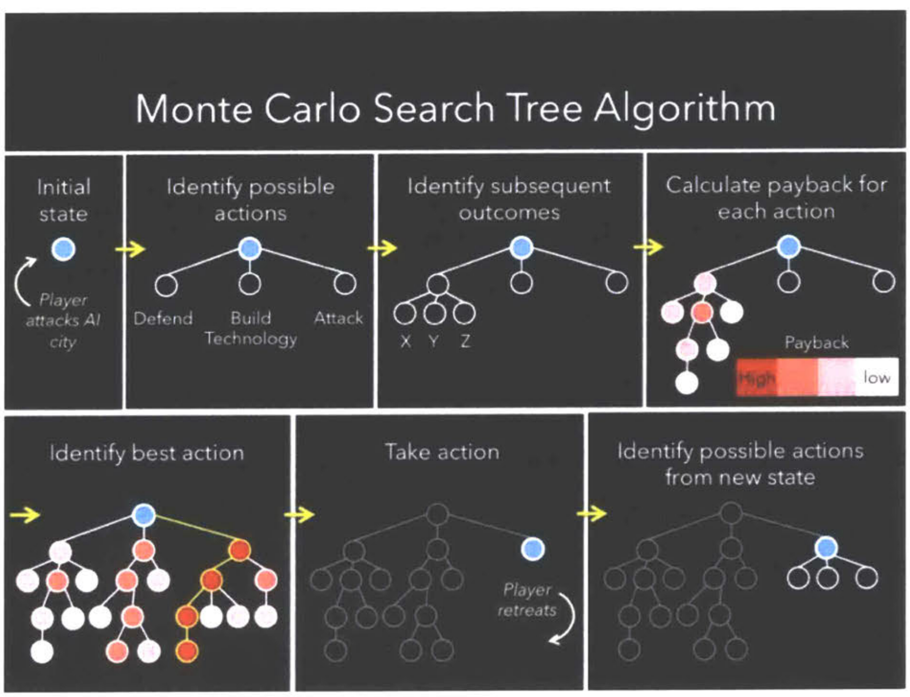

Figure 28: Monte Carlo Search Tree Algorithm

(Source: http://sitn.hms.harvard.edu/flash/2017/ai-video-games-toward-intelligent-game/)

According to Harbing Lou [5.25], "In a game like Civilization [5.26] in which players compete to develop a city in competition with an AI engine which is doing the same thing, it is impossible to pre-program every move for the AI. Instead of acting only based on current status, a Monte Carlo Search Tree (MSCT) AI algorithm evaluates some of the possible next moves, such as developing 'technology', attacking a human player, defending a fortress, and so on. The algorithm then performs the MCST to calculate the overall payback of each of these moves and chooses whichever is the most valuable. Complicated open-world games like Civilization employ MCST algorithm to provide different AI behaviors in each round. In these games, the evolution of a situation is never predetermined, providing a fresh gaming experience for human players every time."

5.5. ARTIFICIAL INTELLIGENCE IN ROBOTICS

Advances in Robotics and Artificial Intelligence has enabled machines to efficiently and quickly perform not only dangerous jobs but also more mundane and time-consuming tasks, thus allowing professionals to take up more innovative job roles. Building on the advances made in mechatronics, electrical engineering and computing, robotics is developing increasingly sophisticated motor functions that give machines the ability to adapt to their changing environment. Artificial Intelligence is used in Robotics to try to optimize its level of autonomy through learning. Artificial Intelligence is also used in Robotics to improve robotic perception of humans such as understanding activities, emotions and non-verbal communications.

5.5.1. MACHINE VISION

Machine Vision integrates image capturing hardware systems with computer vision software algorithms to provide automatic inspection and robot guidance in the execution of their functions.

Kawasaki Robotics [5.26] offers Robotic Assembly systems that uses 2D and 3D vision systems for inspection and motion guidance. According to the company, "these vision systems can identify the style of part to be assembled, or uniquely locate a part feature in order to accurately assemble the product. Advances in control hardware and software allows the robot to gently comply with tooling or the part to correct for any positional error. For large assembly applications, multiple robots can be synchronized to lift or assemble a large component using cooperative motion control. Servo end of arm tooling and multi-axis positioners can all be controlled by the Kawasaki robot controller to complete multiple assembly operations in a single work cell."

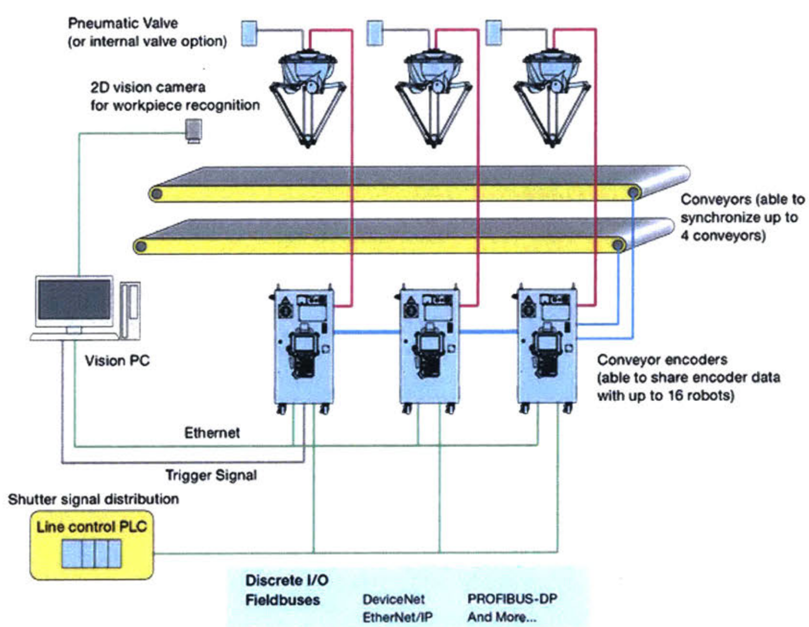

Figure 29: Kawasaki 2D Vision System

(Source: https://robotics.kawasaki.com/en1/products/other/vision-systems/?language_id=4)

5.5.2. HUMAN ROBOT INTERACTION

Human-Robot Interaction (HRI) is a field dedicated to understanding, designing, and evaluating robotic systems for use by or with humans. This multidisciplinary research area draws from Robotics, Artificial Intelligence, Human-Computer Interaction, and Cognitive Psychology. Robots are increasingly being developed for real world application areas, such as robots in rehabilitation, eldercare, or robots used in robot-assisted therapy and other assistive or educational applications.

In the field of assistive technology and rehabilitation robotics, Artificial Intelligence is being used to develop Socially Assistive Rehabilitation Robots that are able to interact with humans. The company Loro [5.28], that was part of the MIT Fuse and Harvard ILAB Incubation program, has developed an AI powered smart companion robot for a wheelchair user to navigate safely and to communicate efficiently.

Figure 30: Loro's Companion Robot for Wheelchair Users

(Source: https://www.loro.xyz/our-technology)

The assistive robot mounts to a wheelchair and offers its occupant the ability to see and interact with the people and things around them in powerful ways. The smart robot uses Facial Recognition technology to recognize predefined faces and helps to keep the screen centered on the person to enhance the quality of communication. The robot also uses

Natural Language Processing, Speech-To-Text and Text-To-Speech to aid people with speech or hearing impairments.

5.5.3. SELF DRIVING CAR

Artificial Intelligence, Cameras, Radar and Lidar sensor technologies, Laser-light and network infrastructure are critical components behind safe and efficient autonomous vehicle operation. Training a self-driving car to behave like a human driver, or to drive better than a human, is the current focus of artificial intelligence research in autonomous driving. Deep Learning algorithms are developed to interpret, predict, and respond to data accrued from millions of miles driven on public roads and in simulations. These algorithms model not only how care recognize objects in the road, but also how human behaviour affects how cars should behave.

Waymo [5.29] uses Artificial Intelligence to train its cars for situations like maneuvering through construction zones, moving over for emergency vehicles, and making room for cars that are parallel parking. This is done by feeding the system examples based on observed real-world situations. Researchers at Waymo are using different scenarios to teach a Recurrent Neural Network (RNN) dubbed the ChauffeurNet, how to drive and then ran the model on a real car.

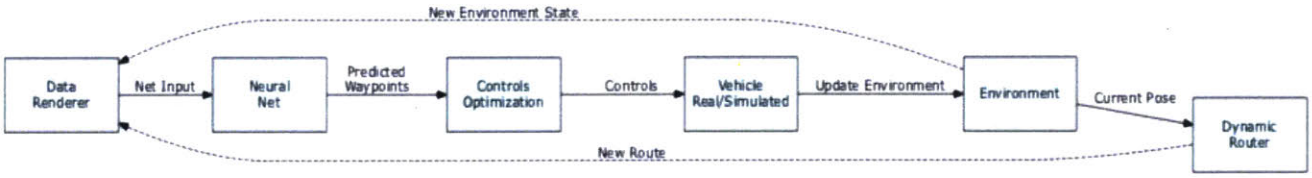

Figure 31: Software Architecture for the End-to-End Driving Pipeline

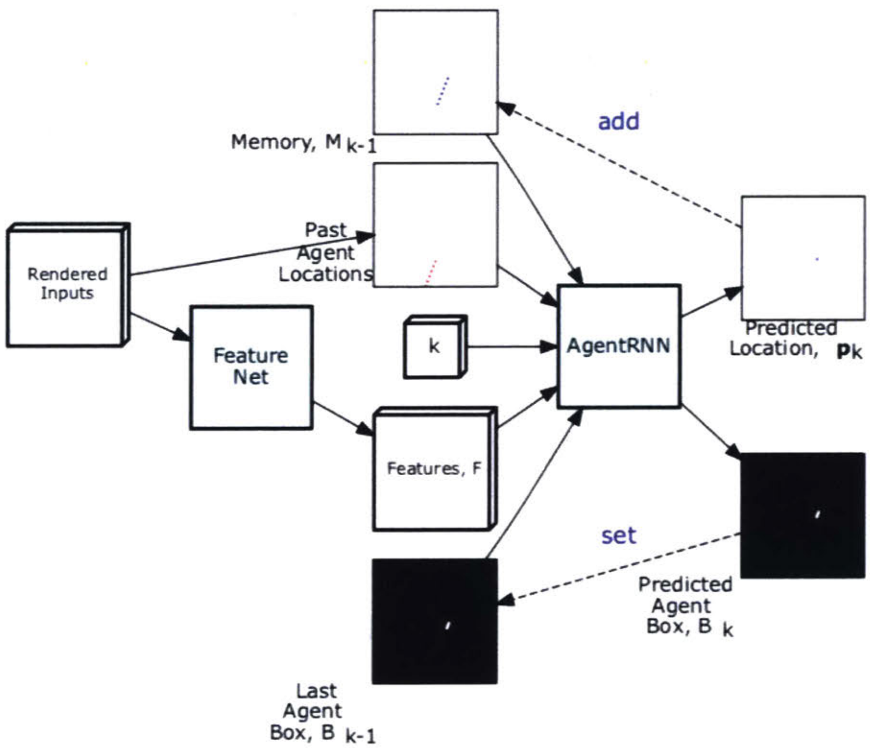

Figure 32: Schematic of ChauffeurNet
(Source: https://arxiv.org/pdf/1812.03079.pdf)

According to the researchers at Waymo [5.30], "ChaeufferNet learned to generate a driving trajectory by observing a combination of real and simulated data, including a map, surrounding objects, traffic lights states, and the past motions of cars. A low-level controller converted the ten-point trajectory to steering and acceleration commands, allowing the AI model to drive both real and digital cars. Waymo vehicles today uses a combination of machine learning and explicit reasoning to continuously evaluate a large number of possibilities and make the best driving decisions in a variety of different scenarios". Waymo is now building Machine Learning algorithms to train its Neural Networks to drive on streets with unclear markings, so that self-driving cars can making guidance decisions in the most challenging driving environments without white lines or clear demarcations at the edge of the road.

5.6. ARTIFICIAL INTELLIGENCE IN AGRICULTURE

Climate change, population growth and food security concerns have propelled the agriculture industry to seek more innovative solutions to protect and improve crop yield. Machine learning, image recognition, and predictive modeling are being applied in the agriculture industry as ways to boost productivity and efficiency. These approaches could be important steps in the effort to produce more food for a growing global population by helping farmers reduce chemical inputs, detect diseases sooner, buffer against labor shortages, and respond to weather conditions as the climate changes.

5.6.1. AGRICULTURAL ROBOT

In many farms in the United States, Japan and Australia, robotic harvesting, irrigation, weeding and seeding devices are being field-tested. Machine Learning systems are helping farmers overlay more data to understand their farms better and make their precision methods more accurate.

The company BlueRiver [5.31] is using automation and robotics to help farmers find more efficient ways to protect their crops from weeds. The company has built smart farms machines using Computer Vision and Artificial Intelligence to monitor and precisely spray herbicides to control the growth of weeds. The company claims that its "precision technology eliminates 80 percent of the volume of chemicals normally sprayed on crops and can reduce herbicide expenditures by 90 percent".

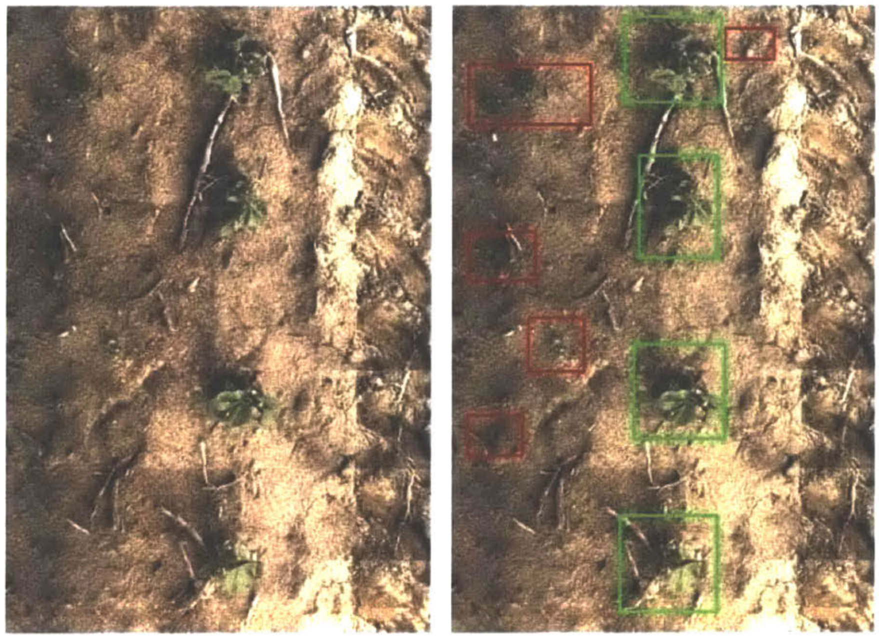

Figure 33: Blue Rivers Farming Robots uses AI & Computer Vision

(Source: http://www.bluerivertechnology.com)

5.6.2. CROP HEALTH DIAGNOSIS

Many plants and crops in Africa are susceptible to not only droughts but also to many diseases and pests. The symptoms of a diseased plant develop slowly, so it can be difficult for farmers to diagnose these problems in time. The company PlantVillage [5.32] has developed a solution using Machine Learning that could help farmers better identify and manage various diseases that afflict crops in Africa.

According to the researchers [5.33], "they annotated thousands of cassava plant images, identifying and classifying diseases to train a machine learning model using TensorFlow. Once the model was trained to identify diseases, it was deployed in the app. Farmers can wave their phone in front of a cassava leaf and if a plant had a disease, the app could identify it and give options on the best ways to manage it".

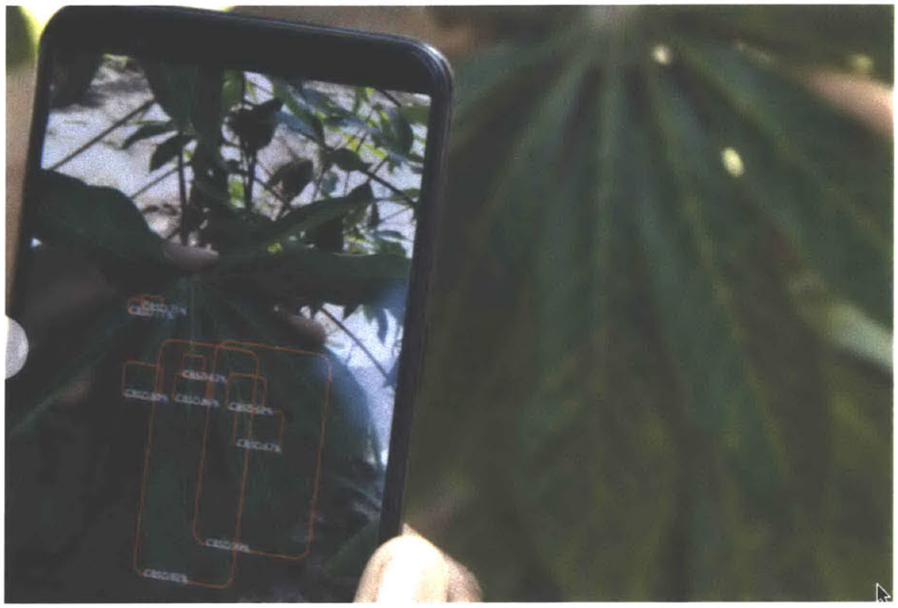

Figure 34: Nuru App Used in Identifying Specific Diseases
(Source: https://www.blog.google/technology/ai/ai-takes-root-helping-farmers-identity-diseased-plants/)

5.6.3. WEATHER FORECASTING

The agriculture industry has started to use Artificial Intelligence to improve weather forecasting, which ultimately determines the best time to plant, fertilize, spray, irrigate and harvest crops.

The company aWhere [5.34] uses machine learning algorithms in connection with satellites to predict weather, analyze crop sustainability and evaluate farms for the presence of diseases and pests. According to the company, "the platform delivers localized information and recommendations to the farmers who use its services. It also forecasts the output of crops around the world, which helps farmers foresee price volatility and respond accordingly". The company also claims that it provides its users with access to over a billion points of agronomic data on a daily basis. Data sources include temperature, precipitation, wind speed, and solar radiation, along with comparisons to historic values.

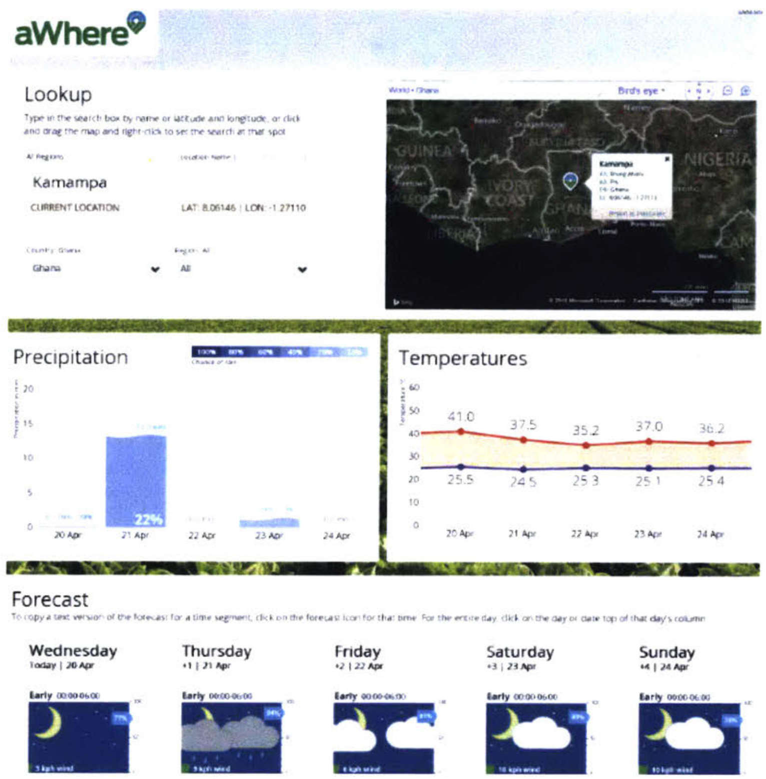

Figure 35: Localized Weather Prediction of aWhere
(Source: https://sdtimes.com/agribusiness/helping-modern-farms-increase-production-power-apis/)

CHAPTER 6: ETHICAL, LEGAL AND SOCIETAL IMPLICATIONS OF AI

Although the benefits and industrialization of Artificial Intelligence have been researched extensively, there is an urgent need to examine the ethical, legal and socio-political implications of AI. There are significant concerns about how the rapidly growing development and utilization of various Artificial Intelligence technologies relates to societal good, laws and ethics. There is a strong and complex relationship between ethics and law. Codes of ethics are nested within the appropriate legal jurisdictions. Although the law of the land generally embodies some ethical principles, moral judgements often concern issues which go beyond the concerns of the law.

6.1. ETHICS AND AI

One of the most interesting thought experiments in ethics is called the "Trolley Problem" [6.1]. The problem describes a situation where "a runaway trolley is going down the railway lines and there are five people tied to the track ahead. You are standing next to a lever. If you pull the lever, the trolley will switch to a different set of tracks. However, there is another person tied to that set of tracks. Do you pull the lever or not"?

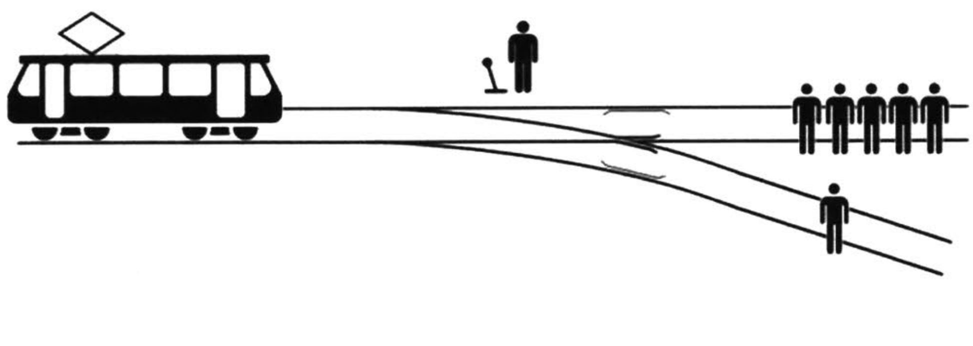

Figure 36: Trolley Problem

(Source: https://en.wikipedia.org/wiki/Trolley_problem)

There is no definitive answer to this question. Some people might pull the switch, thus extinguishing one life instead of five. Whereas others may say that pulling the switch is wrong because there is a difference between killing someone intentionally versus letting them die through circumstances beyond your control.

A clear example of ethical consideration when it comes to the utilization of Artificial Intelligence is the driverless car. In the absence of mechanical fault with the car, the driver is typically liable for loss it causes. The introduction of AI, however, has the potential to shift that liability up the supply chain to the manufacturer or to another supply chain participant. If a self-driving vehicle is travelling along a road in a situation where an accident is unavoidable, the question then arises as to whose lives should take priority: the lives of the passengers or the lives of the pedestrians. Massachusetts Institute of Technology created a special website [6.2] to gather human perspective on moral decisions when it comes to self-driving cars. According to a survey done by the MIT Media Lab called Machine Ethics [6.3], many of the moral principles that guide a driver's decisions vary by country. Most people from various countries believed that the moral thing to do was to spare humans over pets, the young over the old and groups of people over individuals. The study called Moral Machine, reveals cultural nuances that governments and makers of self-driving cars might want to take into account if they want the vehicles to gain public acceptance.

German automaker Mercedes Benz has stated that its self-driving cars will save the car's drivers and passengers, even if that means sacrificing the lives of pedestrians, in a situation where those are the only two options [6.4]. German rival automaker Audi says it will assume full legal responsibility for any crashes or fatalities from its first Level 3 self-driving car. Germany's Ethics Commission on Automated Driving created initial guidelines for automated vehicles [6.5]. One of their key dictates is a prohibition against such decision-making by a car's operating system.

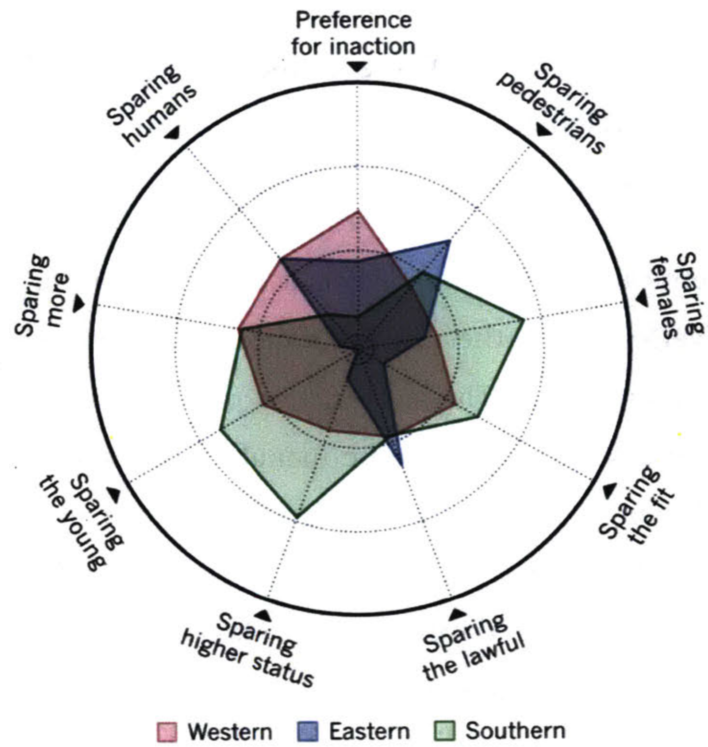

Figure 37: Variation in Moral Principles that Guide Driver's Decisions

(Source: https://www.nature.com/articles/s41586-018-0637-6)

Another ethical consideration is regarding the Weaponization of Artificial Intelligence. AI is being utilized to create lethal autonomous weapons and there are concerns of the impact if such technologies fall into the hands of a rogue state or terrorist organizations. Artificial Intelligence has enabled the creation of unmanned combat aerial vehicles that are capable of replacing human-piloted fighter planes. Robotic tanks and submarines that can operate autonomously are being developed. Current Artificial Intelligence may contain significant bias that affects its decision making. This has led to many associations around the world to carry out awareness-raising actions to contain the use of AI in creating weapons.

6.1.1. ETHICAL CONSIDERATIONS

Here are some of the major ethical considerations for Artificial Intelligence:

- AI systems should use training data and models that are free of bias.
- AI systems should be safe and secure and not vulnerable to tampering.
- AI systems should include details of their development, deployment and maintenance so that they can be audited throughout their lifecycle.
- AI algorithms should incorporate societal norms, policies and regulations that correspond to well establish human norms.

6.2. LAW AND AI

Artificial Intelligence technologies can be regulated as items subject to copyright or as property. But difficulties around legal frameworks arise since AI technologies can act autonomously. Many machines that are powered by Artificial Intelligence are becoming more removed from any explicit programming and are in turn based on Machine Learning methods. So, it has become harder to attribute the question of fault. The existing legal frameworks deal comfortably with traceable defects, which are machine decisions that are traceable back to defective programming or incorrect operation. Legal frameworks begin to fail however where defects are inexplicable or cannot be traced back to human error [6.6].

The company Cambridge Analytica [6.7] harvested the data from 50 million Facebook profiles. The company then developed Artificial Intelligence algorithms that could both predict and influence voting behaviour by creating campaigns that were designed to appeal to individuals' beliefs and preferences. This is an example of Artificial Intelligence and analytics deployed without data subject's consent. It is very important that legal protections are put in place to prevent any organization from using AI to abuse public trust in technology platforms.

Academics and regulators from around the world are working towards building regulatory measures and technical standards on the governance of Artificial Intelligence. More than 20 countries have put forward new AI Strategies. These initiatives aim to address legal and ethical framework, privacy and protection of personal data as well as principles of accountability and transparency [6.8].

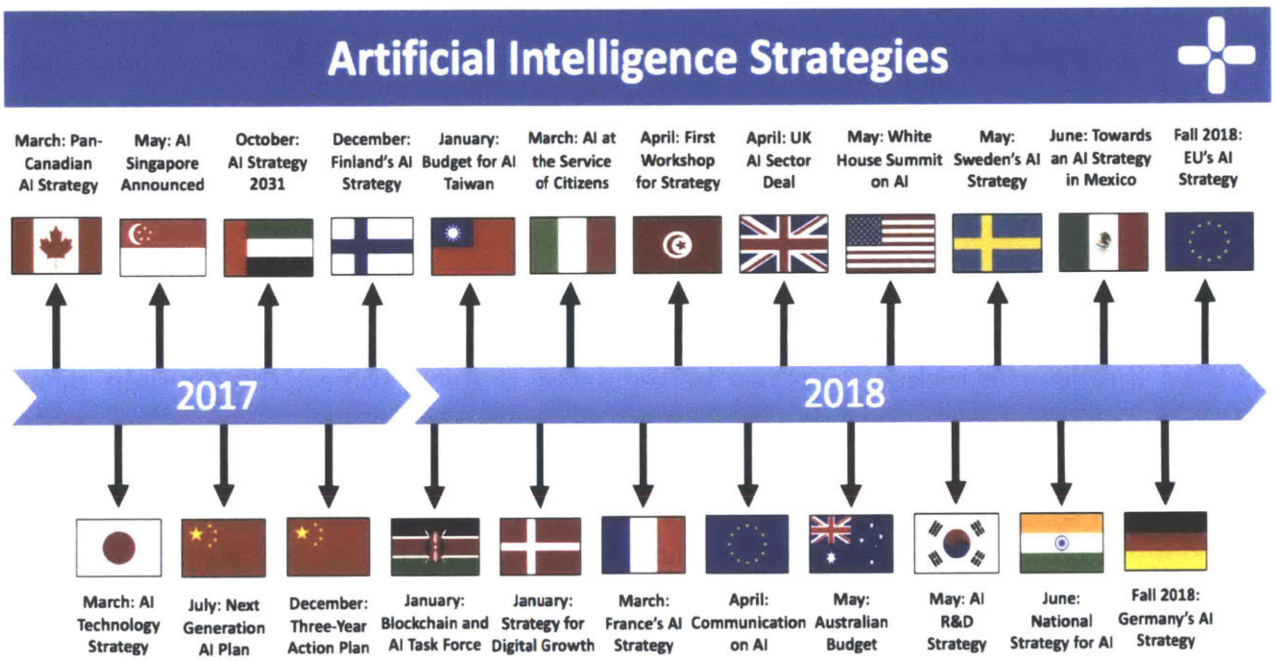

Figure 38: Artificial Intelligence Strategies of various Nations
(Source: https://medium.com/politics-ai/an-overview-of-national-ai-strategies-2a70ec6edfd)

6.2.1. LEGAL CONSIDERATIONS

Here are some of the major ethical considerations for Artificial Intelligence:

- Accountability of AI systems should lie on someone on whom the responsibility for loss or damage can be attributed.

- It is important to address liability issues for AI systems: making autonomous systems legally liable for their actions or making programmers and users of autonomous systems liable for the actions of those systems.
- Legal frameworks should regulate the potential impact of the AI systems on people.
- It is also important to regulate the impact of AI systems on goods and services, potential antitrust implications and intellectual property rights.

6.3. SOCIETY AND AI

While Artificial Intelligence has the potential to solve many of the problems in our social life and the business world, there is a concern about the impact on jobs. The pace at and extent to which automation will be adopted and impact actual jobs will depend on several factors such as technical feasibility, cost of adoption, labor supply quantity and associated wages. According to a study carried out by McKinsey [6.9], "in advanced economies with relatively high wage levels, such as France, Japan, and the United States, jobs affected by automation could be more than double that in India, as a percentage of the total". However, it is important to note that AI will also create millions of jobs in almost all the industry sectors. AI will eliminate low-skilled jobs and create high-skilled job opportunities. In the face of the potential social consequences, business leaders and regulators should reimagine work to meet the needs of all workforce segments in all job types.

An important societal consequence of AI adoption is the algorithmic bias of models trained on data that reflects historical gender and racial discrimination. Algorithmic bias can have damaging effect in areas such as law enforcement, recruitment, security and loan approval. If not addressed, algorithmic bias can magnify human biases. Measures should be taken to ensure that the data that is used to train algorithms is diversified. Statistical methods should be used to spot hidden biases in algorithms. Individuals from marginalized and under-represented sections of the society should find fair representations in AI research groups. Benefits of the AI revolution should reach all sections of the society.

CHAPTER 7: AI COMMERCIALIZATION FRAMEWORK

There are various business and technical scenarios that should be evaluated before commercializing Artificial Intelligence technology in a specific industry. It is useful to establish a framework to model the environment and the internal and external forces which can accelerate or delay adoption of AI. This proposed framework is intended to be used by key decision makers in conjunction with the technical and business perspective highlighted in this thesis.

7.1. OUTLINE OF THE FRAMEWORK

The proposed Framework is organized into four parts:

- Reviewing the Enterprise Model
- Identifying AI use cases and Opportunities
- Assessing the Technology Readiness
- Productizing AI

The **Review of the Enterprise Model** guides the decision maker in evaluating the AI maturity and strategic intent of an organization, resulting in four distinct level of AI organizations – AI Pioneers, AI Investigators, AI Experimenters and AI Laggards.

The **Identification of AI use-cases and Opportunities** is an assessment of the practical application and the economic potential of advanced AI, ML and DL technologies across industries and business function to explore potential opportunities.

The **Assessment of Technology Readiness** is a set of toolkits to evaluate the organization's ability and readiness to use AI to generate business value as well as understand the current state of different AI technologies to identify areas that need to be addressed.

The **Productization of AI** guides the decision maker in determining product requirements, architecting an end-to-end AI system and selecting a go-to-market strategy to create a successful AI product.

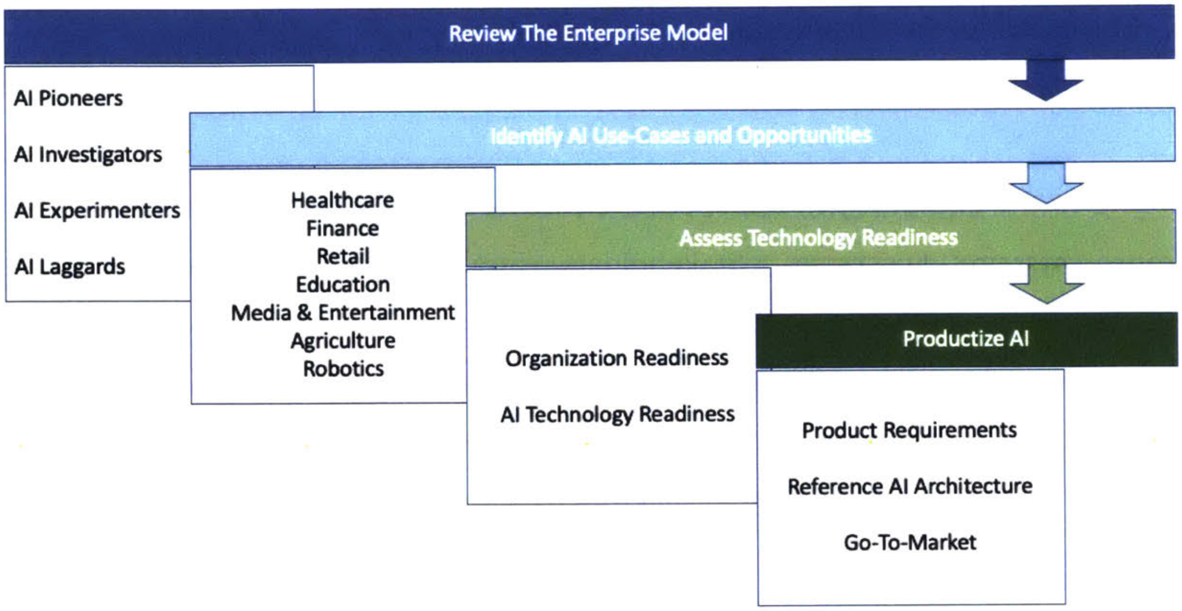

Figure 39: Artificial Intelligence Commercialization Framework

7.2. REVIEW THE ENTERPRISE MODEL

The first part of the Commercialization Framework is the Reviewing of the Enterprise Model. Industries and organizations will need to undergo major transformations in response to the disruptive nature of Artificial Intelligence. A framework for Enterprise Architecting provides a holistic approach to analyzing the 'current state' enterprises and designing the 'future state' enterprise system. Enterprise Architecting is a systems approach to business and organizational design. Deborah Nightingale and Donna Rhodes [7.1] developed a framework for enterprise transformation. According to the authors, there are seven architecting imperatives that should be followed:

1. "Make architecting the initial activity in transformation.
2. Develop a comprehensive understanding of the enterprise landscape.
3. Understand what stakeholders' value and how that may change in the future.
4. Use multiple perspectives to see the whole enterprise.
5. Create an architecting team suited to the transformation challenges.

6. Engage all levels of leadership in transformation.
7. Architect for the enterprise's changing world".

According to the authors [7.1], ten specific elements are used to describe and design an enterprise system: Ecosystem, Stakeholders, Strategy, Organization, Information, Knowledge, Processes, Services, Products, and Infrastructure. A few of the ten principles are used to evaluate the organization maturity: Stakeholder (degree of support & sponsorship of various stakeholders), Strategy (the state and stage of AI as part of the digital transformation journey, AI budget consideration), Organization (structural support of AI within the organization), Knowledge (understanding and identification of operational use cases for AI), Processes, Services, Products (degree to which AI is being harnessed in the organization). Based on the understanding and adoption of AI, four distinct clusters of organization maturity can be seen, as summarized in Table 3.

Table 3: Organization Maturity for AI adoption

AI Pioneers	Organizations that understand the implications of AI and have adopted AI.
AI Investigators	Organization that understand the implications of AI but have only deployed AI in pilot stage.
AI Experimenters	Organizations that are adopting AI without a deep understanding of the implications of AI.
AI Laggards	Organizations with very little understanding of AI and no plans of AI adoption.

(Adapted from https://sloanreview.mit.edu/projects/reshaping-business-with-artificial-intelligence/)

Organizational flexibility is core to the commercialization and adoption of AI. Investments in Artificial Intelligence should begin with securing the buy-in from the impacted business process stakeholders. AI Pioneers and AI Investigators are the best placed organizations to commercialize and adopt Artificial Intelligence.

7.3. IDENTIFY AI USE CASES AND OPPORTUNITIES

The next step is creating a baseline adoption for specific Artificial Intelligence categories in specific industries. This will help in creating the understanding of the current state of adoption and the timeframe over which the adoption will progress from the pioneers to the laggards.

McKinsey Global Institute carried out a survey across industry sectors to understand the sector-by-sector adoption of Artificial Intelligence [7.2].

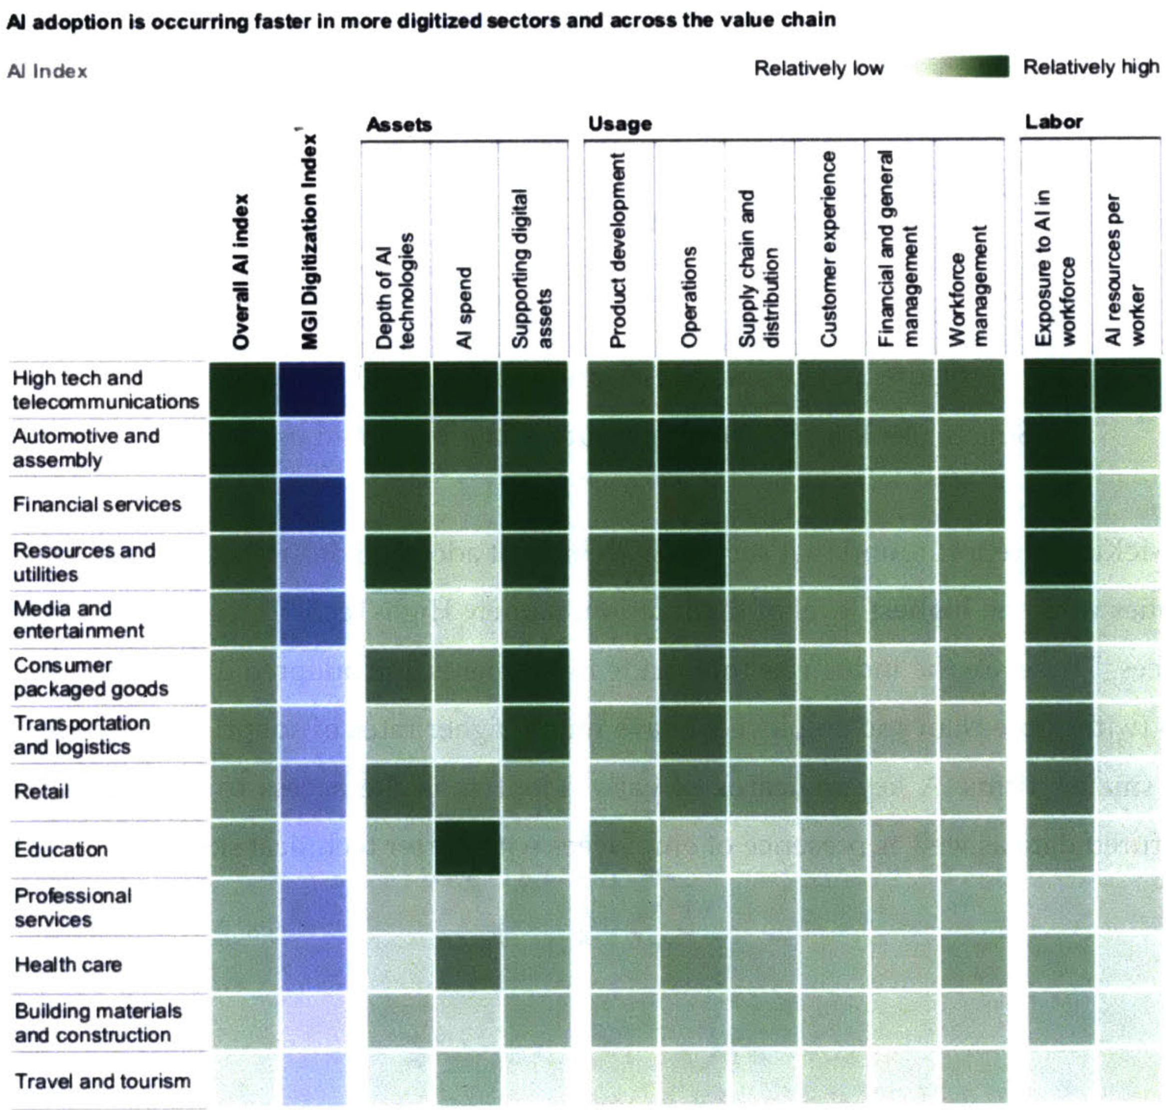

Figure 40: Adoption of AI Across Industries [7.2]

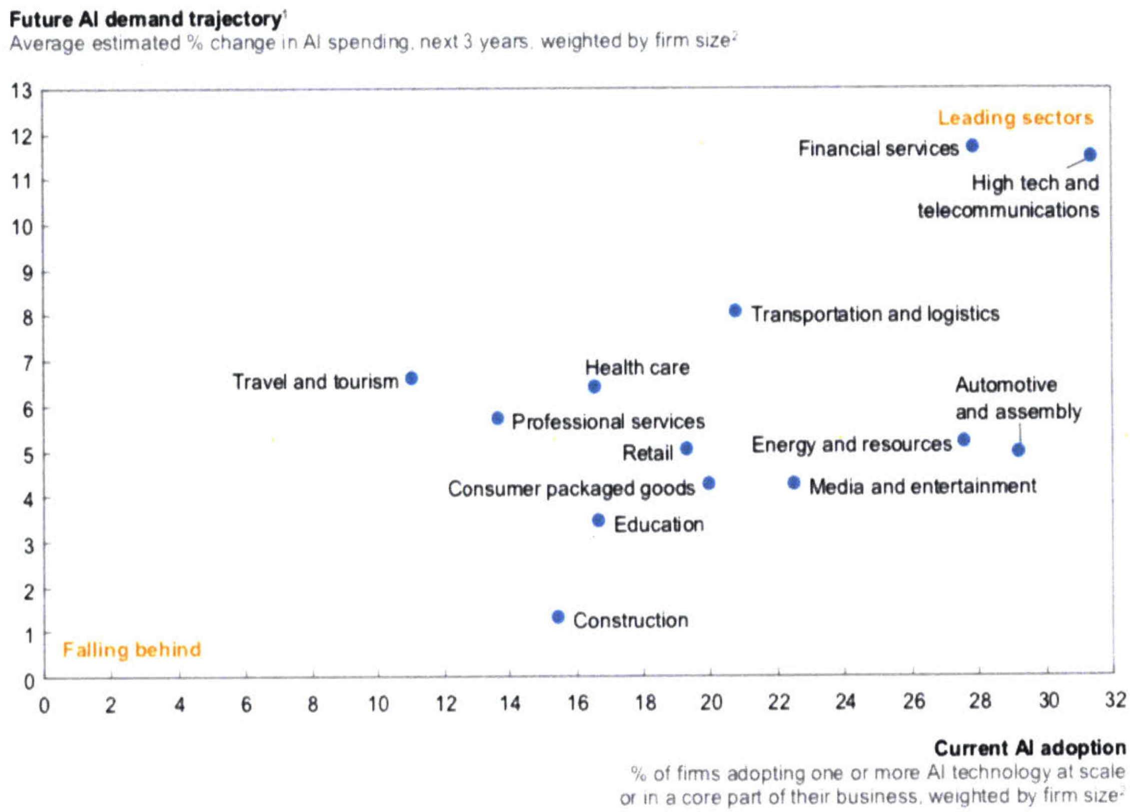

Figure 41: Current AI Adoption vs Future AI Demand

(Source: McKinsey Global Institute AI adoption and use survey)

The McKinsey report found that sectors with highest adoption of Artificial Intelligence are the ones with the highest level of digitization, namely High-Tech, Telecom and Financial Services. These are the industries that made investments and adopted digitization. Larger firms (with more than 500 employees) have much higher rates of adoption and awareness than smaller firms. A key indicator of early adopters is the access to more and better structured data as well as presence of employees with better technical skills.

7.3.1. ARTIFICIAL INTELLIGENCE USE CASES

Based on the extensive research done on the adoption of AI across industry sectors, primary research with Subject Matters Experts in the leading AI service providers and exploratory research conducted, here are some of the prominent use cases of AI:

- Table 4 summarizes AI use-cases in the Healthcare sector.
- Table 5 summarizes AI use-cases in the Finance sector.
- Table 6 summarizes AI use-cases in the Retail sector.
- Table 7 summarizes AI use-cases in the Education sector.
- Table 8 summarizes AI use-cases in the Media and Entertainment sector.
- Table 9 summarizes AI use-cases in the Agriculture sector.
- Table 10 summarizes AI use-cases in the Robotics sector.

Table 4: AI use-cases in Healthcare

SECTOR	USE CASE	APPLICATION	AI METHODS
Healthcare	1. Early Diagnosis	Leverage medical images & data to enable early diagnosis of diseases	Computer Vision, Neural Networks, Medical Imaging
	2. Patient Data Analytics	Analyze patient data to discover insights and suggest actions	Machine Learning, Advanced Analytics
	3. Drug Discovery	Discover new drugs based on scientific data and medical intelligence	Machine Learning, Advanced Analytics
	4. AI assisted Surgery	Identify distinct body structures from medical images (MRI, XRAY) to aid surgery	Machine Learning, Computer Vision

	5. Medical Transcription	Transcribe doctor's voice into medical documents such as EHR	Natural Language Processing
	6. Tele Medicine	Provide remote access to medical consultation and pathology services	Natural Language Processing, Computer Vision

Table 5: AI use-cases in Finance

SECTOR	USE CASE	APPLICATION	AI METHODS
Finance	1. Credit Scoring	Create credit scores for applicants that lack a credit history but indicate a high likelihood of being good borrowers	Machine Vision, Machine Learning, Natural Language Processing
	2. Robo Advisors	Use chatbots and voice assistants to monitor and manage personal finances	Natural Language Processing
	3. Fraud Detection	Detect fraudulent and abnormal financial behavior	Machine Learning Prescriptive Analytics, Anomaly Detection
	4. Insurance Claims	Process insurance claims, optimize claims process and reduce overpayments	Natural Language Processing, Computer Vision
	5. Algorithmic Trading	Automate micro level trading decisions for equities, futures contracts	Reinforced Learning, Deep Learning, Quantitative Model

	6. Regulatory Compliance	Scan large legal and regulatory documents to extract information	Natural Language Processing

Table 6: AI use-cases in Retail

SECTOR	USE CASE	APPLICATION	AI METHODS
Retail	1. Sales & Customer Relationship Management	Generate account insights, prioritize leads and convert social media interactions	Machine Learning, Natural Language Processing
	2. Product Recommendation	Use customer's data to recommend products they would be interested in	Machine Learning, Regression Analysis
	3. Warehouse Automation	Automate decision making and operations in the warehouse	Machine Learning, Computer Vision, Robotics
	4. Social Analytics	Analyze and act upon data generated by potential customers in social media	Machine Vision, Natural Language Processing
	5. Demand Forecasting	Predict localized demand by understanding customer choice, product assortment and pricing	Machine Learning, Advanced Analytics
	6. Pricing and Promotions	Allow companies to optimize pricing to minimize cannibalization & maximize revenues	Machine Learning, Advanced Analytics

Table 7: AI use-cases in Education

SECTOR	USE CASE	APPLICATION	AI METHODS
Education	1. Smart Content	Create smart content from digitized textbooks	Machine Learning
	2. Tutoring System	Provide personalized tutoring and real-time feedback for students	Machine Learning
	3. Foreign Language Translation	Personalize each foreign language lesson based on a user's progress	Natural Language Processing, Machine Learning
	4. Administrative Task Automation	Automate time consuming activities like grading for multiple choice questions	Optical Character Recognition, Machine Learning
	5. Teaching Assistants	Design chatbots to provide answers to student queries to increase student engagement	Natural Language Processing
	6. Augmented Learning	Create interactive and 3D virtual environment for students to learn	Computer Vision, Artificial Intelligence

Table 8: AI use-cases in Media and Entertainment

SECTOR	USE CASE	APPLICATION	AI METHODS
Media and Entertainment	1. Advertisement	Create movie trailers by analyzing scenes based on visual, audio composition	Machine Learning
	2. Sports Highlights	Generate a highlight reel from sports events by	Machine Learning

		analyzing video footage and fan reaction	
	3. Movie and Show Recommendation	Match video content to users based on their preferences	Neural Network
	4. Search Optimization	Perform searches and run analytics on visual assets within large databases	Machine Learning
	5. Video Games	Generate responsive, adaptive and intelligent behaviors and generate content in games	Artificial Intelligence
	6. Marketing Analytics	Utilize data assets to manage campaigns and improve awareness	Machine Learning, Predictive Analytics

Table 9: AI use-cases in Agriculture

SECTOR	USE CASE	APPLICATION	AI METHODS
Agriculture	1. Robotic Farming	Autonomous machines to handle agricultural tasks such as seeding, weeding and harvesting	Computer Vision, Artificial Intelligence
	2. Crop Health Diagnosis	Process data captured by drones to monitor crop	Computer Vision, Deep Learning
	3. Soil Monitoring	Identify potential defects and nutrient deficiencies in soil	Computer Vision, Deep Learning, IOT
	4. Weather Forecasting	Provide accurate localized weather forecasts	Machine Learning, Sensor

	5. Predictive Analytics	Track and predict various environmental impacts on crop yield	Machine Learning, Advanced Analytics

Table 10: AI use-cases in Robotics

SECTOR	USE CASE	APPLICATION	AI METHODS
Robotics	1. Robotic Surgery	Carry out micro-surgical procedures to help reduce surgeon variation	Computer Vision, Machine Learning
	2. Assistive Technology	Develop devices that can sense, process sensory information & perform actions for people with disabilities	Computer Vision, Natural Language Processing, Artificial Intelligence
	3. Self-Driving Cars	Create fully autonomous self-driving vehicles	Deep Learning, Computer Vision
	4. Shared Mobility Services	Allocate rides, adjust ride sharing prices and optimize traffic routes	Machine Learning
	5. Driver Assistance	Provide driving assistance by tracking the road, detecting eye movements and detecting speech	Computer Vision, Natural Language Processing
	6. Automated Manufacturing	Completely automate the manufacturing and quality control processes	Computer Vision, IoT, Predictive Maintenance

7.3.2. MACHINE LEARNING USE CASES

Table 11 summarizes some of the use-cases of utilizing Machine Learning methods.

Table 11: Machine Learning Use Cases

ML METHOD	SUB TYPE	APPLICATION
Supervised Learning	Classification	Diagnostics
		Image Classification
		Customer Retention
		Identity Fraud Detection
	Regression	Weather Forecasting
		Advertising Popularity Prediction
		Market Forecasting
		Estimating Life Expectancy
		Population Growth Prediction
Unsupervised Learning	Dimensionality Reduction	Feature Elicitation
		Meaningful Compression
		Big Data Visualization
		Structure Discovery
	Clustering	Recommender Systems
		Customer Segmentation
		Targeted Marketing
Reinforced Learning		Real-time Decisions
		Learning Tasks
		Robot Navigation
		Skill Acquisition
		Game AI

(Adapted from a Brief Taxonomy of AI)

7.4. ASSESS TECHNOLOGY READINESS

The next step in the commercialization of Artificial Intelligence Framework is to understand where to prioritize efforts. This part of the thesis aims to provide guidance on how to judge an organization's ability and readiness to use AI to generate business value as well as understand the current state of different AI technologies.

7.4.1. ORGANIZATION READINESS TOOLKIT

Organizational readiness can be clustered in three broad categories:

- Foundational Readiness: An important pre-requisite for AI is the presence of appropriate infrastructure and interfaces.
- Operational Readiness: Management buy-in and governance mechanisms are essential for sustainable AI solutions.
- Cultural Readiness: The transformational ability of an organization to maximize the value it derives from AI.

Table 12: Organizational AI Readiness Matrix

Foundational Readiness	1. Server Infrastructure
	2. Data Storage
	3. Data Sources
	4. Cloud Resources
	5. Software Packages
Operational Readiness	6. Operational Management
	7. Agile Delivery
	8. Skills and Expertise
	9. Cybersecurity
	10. Governance, Compliance and Risk
	11. Strategic Leadership

Transformational Readiness	12. Business Opportunity
	13. Clarity of Business Case
	14. Business Acceptance

(Adapted from Intel's AI Readiness Model)

7.4.2. TECHNOLOGY READINESS OF ARTIFICIAL INTELLIGENCE

The maturity level of different Artificial Intelligence technology varies. By gathering together information from several studies carried out in both academia and industry, Callaghan Innovation [7.3] has created a technology readiness map of Artificial Intelligence based on both the time span of development and being operational.

Table 13: Artificial Intelligence Technology Readiness

TIMELINE	ARTIFICIAL INTELLIGENCE TECHNOLOGY
Available Now	1. Machine Learning
	2. Deep Learning
	3. Neural Networks
	4. Pattern Recognition
	5. Natural Language Processing
	6. Autonomous Systems
	7. Chatbots
1 – 2 Years	8. Real-time Emotion Analytics
	9. Real-time Universal Translation
	10. Virtual Companions
	11. Thought-controlled Gaming
2 – 5 Years	12. Next-gen Cloud Robotics
	13. Autonomous Surgical Robotics

	14. Robotic Personal Assistants
	15. Cognitive Cyber Security
	16. Optimized Farming Platforms
More than 5 Years	17. Neuromorphic Computing
	18. Fully Autonomous Cars
	19. Fully Autonomous Shipping
	20. Predictive Personalized Health

(Adapted from Callaghan Innovation's AI Demystified Blog)

7.5. PRODUCTIZE AI

The last step of the Framework is to gather product requirements, architect an end-to-end AI system and successfully launch the product or services.

7.5.1. PRODUCT REQUIREMENTS

Here are some of the important considerations when developing Artificial Intelligence, Machine Learning or Deep Learning solutions to solve real-world problems:

1. **Performance:** Commercial application of Artificial Intelligence requires the highest degree of functionality and reliability. It is very important to set the correct scope for production and well as to shape public perception.

- In the product development of an AI solution, it is very crucial to limit performance to a specific use-case or a certain environment.
- If the AI system can perform at a level higher than humans, it will be commercially viable as a stand-alone product.
- If the performance of the AI would not be sufficient to completely replace humans, the AI system can be productized as a decision support system.

- It is important to associate a level of uncertainty to a prediction of a machine learning or deep learning model.
- The model has to work not only on training datasets but also in the real world. The model has to be aware of all the possible elements, conditions and their outcomes regarding the target environment, including all the changes over a period of time.
- New and more accurate test datasets should be continuously built to incorporate changes in data and target environment.

2. Model: When designing the Artificial Intelligence model, it is important to optimize and align the performance metric to the business goal.

- Negative consequences of misclassification or wrong prediction can be catastrophic, especially in the context of medical diagnosis or autonomous driving.
- The algorithms should be explainable in terms of the inputs, outputs and internal workings.
- It is important to note that the complexity of a model is inversely proportional to its explainability. Trade-offs between the two should be made.
- It is crucial to consider computational constraints and connectivity constraints while designing models for industry applications.
- Models and applications need not be over-engineered for real world application. Models should work in multiple scenarios in order to provide robust systems.

3. Data: The cost-benefit tradeoff for data acquisition is one of the most crucial aspect of commercializing Artificial Intelligence models.

- Large datasets are not only the prerequisite for AI solutions to be developed, but the quality and amount of data determines the accuracy and commercial benefit of the respective AI model.
- If there is an exclusive access to preferred datasets, then building the AI model in-house might be necessary to maintain competitive advantage.

- In the absence of exclusive access to relevant data, buying external AI solutions should be the default option.
- If there is an exclusive access to preferred datasets, but this data is not fully used in the core business, there might be an opportunity to co-develop AI solutions with a strategic partner.

7.5.2. REFERENCE ARCHITECTURE

This thesis proposes a reference architecture that is intended to be used as a reference by business decision makers, data scientists and IT professionals in defining, deploying and commercializing Artificial Intelligence solutions for any industry.

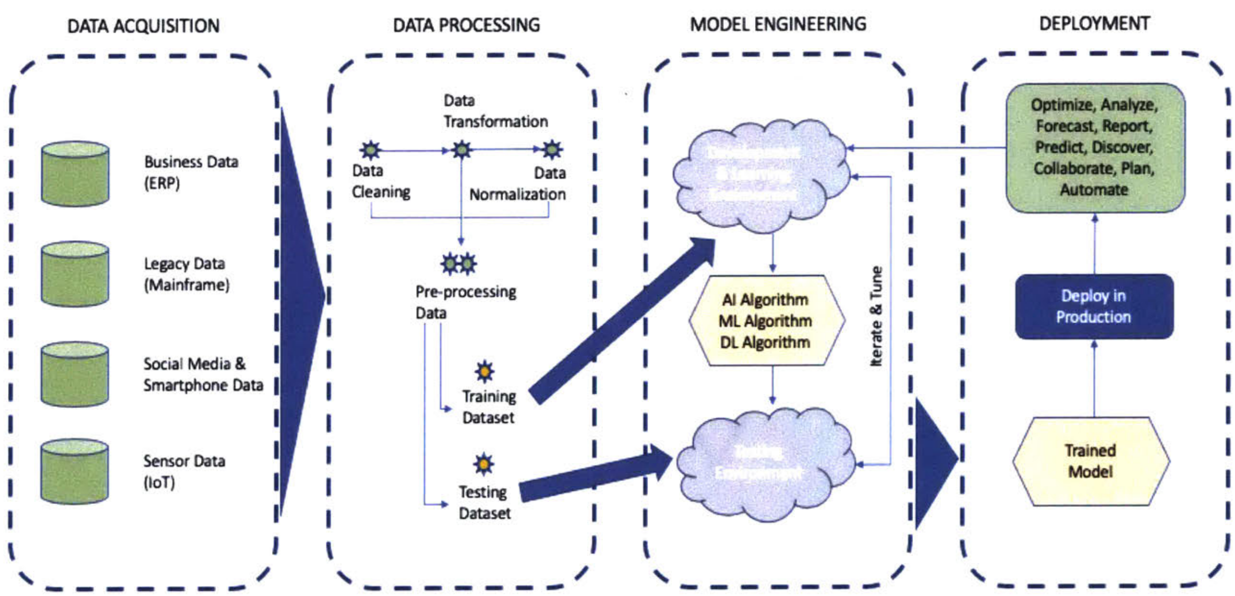

Figure 42: Reference AI Architecture

Data Acquisition: The first component of the AI architecture begins with the collection of high volume of data from a variety of sources such as business data from ERP systems, legacy data sourced from mainframes, user generated data from smartphone apps and social media as well as sensor data from IoT system. Discrete data should usually be handled through a data warehouse whereas continuously streaming data should be handled through a stream processing platform.

Data Processing: The second component of the AI architecture deals with processing of data and includes modules to perform cleaning, transformation and normalization of data. A portion of the dataset is used for training the algorithm whereas the rest of the dataset is used for testing the algorithm. Based on whether data is processed at rest or continuously, throughput computing maybe needed. It is also important to secure the data processing part to ensure privacy and prevent adversarial attacks.

Model Engineering: The third component of the AI architecture is where Artificial Intelligence algorithms, Machine Learning algorithms or Deep Learning algorithms or a combination of three are developed to address a technical or business problem. These algorithms can be built in-house from scratch or algorithm toolkits can be sourced from the marketplace and deployed in the cloud.

Deployment: The fourth component of the AI architecture is setting the stage for the execution of the trained algorithm in the production environment. The algorithm will be tested and tuned repeatedly to optimize performance. A key consideration is the availability of processing power to execute the AI, ML or DL models. The processing of a simple ML algorithm could be handled by a regular desktop server. However, a Deep Neural Network will need high-throughput computing in the form of Graphical Processing Units (GPU), High Performance Computing (HPC) cluster or Cloud Computing.

7.5.3. GO TO MARKET STRATEGY

Solutions are productized if they can be leveraged across multiple installments with little customization, enabling the company to grow at scale. Based on the level of productization and the number of use cases that the solution targets, there are three strategies to go-to-market:

- AI Vertical Product: High productization focused on few use cases
- AI Platform: High productization focused on many use cases
- AI Service: Low productization

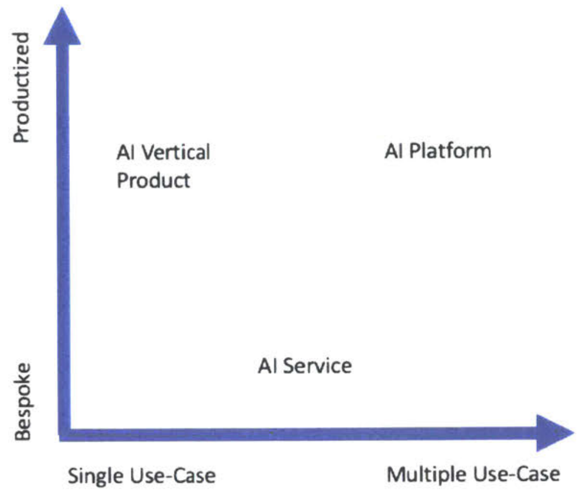

Figure 43: Productization of AI

AI Vertical Product solves a particular business problem or set of problems with a productized solution. Enterprise customers can directly leverage AI without needing to build or maintain models in-house. Such a solution will power not only core functions but also supporting operations, including sales, marketing, finance, customer support and internal communication. The addressable market for such a solution is the second highest.

AI Platform is the underlying infrastructure on top of which specific AI solutions exist. The platform allows end users to import data, perform data transformation, train models and perform model validation. Platforms are highly productized as they are only dependent on the underlying computational framework and hardware that supports them. The addressable market for such a solution is the highest.

AI Service solutions are defined by their low level of productization. AI services provide bespoke solutions to business requirements, as AI models require customization and are highly dependent on customer data. Such solutions cannot be scaled quickly. The addressable market is the smallest of the three.

7.6. CHALLENGES IN COMMERCIALIZING AI

The business world is just beginning to harness AI technologies and their benefits. Here are some of the challenges that enterprises need to overcome to successfully commercialize and leverage AI solutions:

1. Obtaining datasets that are sufficiently large and comprehensive to be used for training models.
2. Complexities of models which makes it difficult to explain the factors that led to a decision or prediction and how it was made.
3. Functional silos in organizations and industries that constrain end-to-end deployment.
4. Dealing with law and regulation.
5. Concerns about data privacy and security.
6. Biases in algorithms that make vital decisions could lead to unethical and unfair consequences.
7. Optimizing algorithms to perform on lower quality hardware without losing significant accuracy.
8. Optimizing product market fit.
9. Lack of talent with relevant technical skillset.
10. Transforming the culture of an organization, people and processes to derive value from insights provided by AI.

CHAPTER 8: CONCLUSION

Artificial Intelligence offers the potential to not only solve real-world and business problems, but fundamentally reshape society and business models. As businesses seek to leverage the advancement in the field of Artificial Intelligence, Machine Learning and Deep Learning to drive commercial growth and efficiencies, decision makers are confused with the complexities, benefits and challenges of these technologies. There is also a lack of frameworks to guide the commercialization of Artificial Intelligence based solutions and products for various industry sectors.

As such, this thesis has provided an overview of Artificial Intelligence, Machine Learning and Deep Learning from both technical and commercial perspectives. This thesis takes a deep-dive into the application and AI use-case for various industries such as healthcare, financial services, education, technology, agriculture and robotics. The goal of this thesis is to present a realistic baseline that allows companies to compare their AI ambitions and efforts. The thesis then proposes an AI Commercialization Framework which guides the decision maker to model the environment and the internal and external forces to commercialize and deploy AI solutions.

8.1. RESEARCH FINDINGS

The thesis provides an overview of the definitions and concepts, landscapes and techniques of Artificial Intelligence and its sub-field: Machine Learning and Deep Learning. Domain and industry specific case studies of organizations that have successfully commercialized AI technologies as well as deployed AI technologies and approaches have been documented in this thesis.

The adoption of Artificial Intelligence is on the rise inside the healthcare industry, solving a variety of problems and improving patient care through radiology diagnosis, blindness prevention programs, telemedicine and drug discovery, paving the way to a broader understanding of health sciences. AI is also finding significant traction in the Financial Services sector, from managing investments to anticipating changes in the stock market to assessing credit worthiness of applicants. AI is being deployed in Education to assist students with learning difficulties, created personalized learning experiences and provide insights into student performance. Transportation is the most hyped industry where a lot of AI investment is being made, to create self-driving fully autonomous cars and intelligent routing and distribution networks. AI in manufacturing is seeing adoption in preventive maintenance to automation of human tasks. Retail industry has also seen heavy adoption of AI in creating immersive catalog visualization and improving the sales and marketing function. AI is increasingly serving as an engine that drives richer, more realistic experiences in media and entertainment, from video games to movies. There is a tremendous potential of using AI in agriculture to reduce the use of chemicals and to empower farmers to manage land more sustainably.

Based on the research work conducted for this thesis, a Theoretical Framework for Commercializing and Adopting Artificial Intelligence has also been proposed. The four parts of the Framework are: i) Review the Enterprise Model, ii) Identify AI Uses-Cases and Opportunity, iii) Assess Technology Readiness and iv) Productize AI.

The first step of the Commercialization Framework is to identify whether the organization is an AI pioneer or AI Investigator. The second part of the Framework is to identify opportunities from a business perspective. The greatest value capture will initially be in the enterprise micro verticals – specific use cases within select industries. The third step in the Framework is to assess the technology stack of the organization to identify areas that need to be addressed as well as to evaluate the current state of different AI technologies. The final step of the Framework guides the development of AI products and services and architecting of the end-to-end AI systems.

8.2. DISCUSSION

The strategic objective driving investments in AI is profitable business growth and competitive advantage. At a strategic level, companies will need to develop an enterprise-wide view of compelling AI opportunities and transform parts of their current business processes. Organizations will need robust data acquisition, governance processes, talent with appropriate skillsets and overcome functional silos to deploy and leverage AI. The role of data strategy, data architecture, data management, data security and data privacy are central to the success of AI adoption. As new data sources emerge over time, a strong governance model becomes essential. As AI systems are moving towards real-time decision-making capabilities, it would necessitate an agile based enterprise architecture.

Alongside the economic benefits, AI will impact society in a positive way as it helps tackle societal challenges ranging from healthcare to nutrition. However, it is essential to address some of the negative consequence of AI such as job loss due to automation, use in surveillance applications, social media manipulations and misuse in criminal justice system where the impact has social consequences. Data privacy and use of personal information are also critical issues to address if AI has to find greater acceptance in business and society.

8.3. SCOPE FOR FUTURE WORK

This thesis examined use-cases of AI adoption across industries ranging from healthcare to agriculture. New patters will like evolve as additional industries as well as public and private sector firms with traditionally weaker technology focus are included in the analysis. It would also be interesting to drill down into specific use cases of failures within the industry sectors that have been covered in this thesis. This will enable a decision maker to fine-tune the strategies to deal with the most dominant barrier of AI adoption.

The thesis also looks at the ethical, legal and societal implication of large-scale adoption of Artificial Intelligence. These complex issues are of vital importance to our future as the advances in AI technologies are inevitable. Future research can focus on more granular ethical and related legal risks that need to be managed by a business developing or using AI in whatever industry sector it occupies. For AI to gain acceptance in a given sector and society, a business will need to take into account the ethical considerations and the legal factors that flow from them.

The Artificial Intelligence Commercialization Framework needs to be tested in a real-world scenario and optimized based on the industry sectors. Qualitative survey should be conducted to observe the development and incubation of nascent AI technologies from around the world.

REFERENCES

[1.1] Crawley, E., Cameron B., Selva D. (2015). System Architecture: Strategy and Product Development for Complex Systems, Pearson

[2.1] Roberts, Jacob. (2016). Thinking Machines: The Search for Artificial Intelligence. [Retrieved from: https://www.sciencehistory.org/distillations/magazine/thinking-machines-the-search-for-artificial-intelligence] [Accessed: 22-Sep-2018]

[2.2] Sharples, M., Hogg, D., Torrance, S., Young, D., Hutchinson, C. (1963). Computers and Thought: A Practical Introduction to Artificial Intelligence. MIT Press.

[2.3] Edwards, B. (2008). The Computer Mouse Turns 40. Macworld. [Retrieved from: https://www.macworld.com/article/1137400/mouse40.html] [Accessed: 22-Sep-2018]

[2.4] IJCAI. (1969). Proceeding of the First International Joint Conference on Artificial Intelligence. [Retrieved from: https://www.ijcai.org/proceedings/1969] [Accessed: 22-Sep-2018]

[2.5] Amber, A., Barrow, H., Brown, C., Burstall, R., Popplestone, R. (1973). A Versatile Computer-Controlled Assembly System, Proc. Third International Joint Conference on AI, Stanford, California. [Retrieved from: https://pdfs.semanticscholar.org/f461/65aee5006d1cd19e44a2481fbfc889f845c0.pdf] [Accessed: 07-Oct-2018]

[2.6] Miller, R., Pople, H., Myers, J. (1982). INTERNIST-1, An Experimental Computer-Based Diagnostic Consultant for General Internal Medicine. [Retrieved from: http://people.dbmi.columbia.edu/~ehs7001/Clancey-Shortliffe-1984/Ch8.pdf] [Accessed: 07-Oct-2018].

[2.7] Kurenkov, A. (2015). A Brief History of Neural Nets and Deep Learning. [Retrieved from: http://www.andreykurenkov.com/writing/ai/a-brief-history-of-neural-nets-and-deep-learning/] [Accessed: 07-Oct-2018]

[2.8] Minsky, M. (1988). The Society of Mind, Simon & Schuster

[2.9] Hsu, F., (2002). Behind Deep Blue: Building the Computer That Defeated the World Chess Champion, Princeton University Press

[2.10] NASA, (1997). Pathfinder: Exploring Mars with the Sojourner Rover. [Retrieved from: https://www.nasa.gov/redplanet/sojourner.html] [Accessed: 23-Dec-2018]

[2.11] Luger, G., (1989). Artificial Intelligence: Structures and Strategies for Complex Problem Solving, Addison Wesley. [Retrieved from: http://iips.icci.edu.iq/images/exam/artificial-intelligence-structures-and-strategies-for--complex-problem-solving.pdf] [Accessed: 23-Dec-2018]

[2.12] Fetzer, J., (1995). Minds and Machines: Behaviorism, Dualism and Beyond. [Retrieved from: https://web.stanford.edu/group/SHR/4-2/text/fetzer.html] [Accessed: 23-Dec-2018]

[2.13] Kulstad, M., Carlin, L. (2013). Leibniz's Philosophy of Mind. Stanford Encyclopedia of Philosophy. [Retrieved from: https://plato.stanford.edu/entries/leibniz-mind/] [Accessed: 23-Dec-2018]

[2.14] Norvig, P., Russell, S., (1994). Artificial Intelligence: A Modern Approach, Prentice Hall

[2.15] Padua, S., (2015). The Thrilling Adventure of Lovelace and Babbage: The True Story of the First Computer, Pantheon

[2.16] Boole, G., (1952). Studies in Logic and Probability, Dover Publications, New York

[2.17] Raatikainen, P., (2013). Godel's Incompleteness Theorems. [Retrieved from: https://plato.stanford.edu/entries/goedel-incompleteness/] [Accessed: 12-Nov-2018]

[2.18] Barrett, L., Connell, M., (2005). Jevons and the Logic Piano. [Retrieved from: http://www.rutherfordjournal.org/article010103.html] [Accessed: 14-Nov-2018]

[2.19] Turing, A., (1950). Computing Machinery and Intelligence. Mind: A quarterly review of Psychology and Philosophy. [Retrieved from: http://phil415.pbworks.com/f/TuringComputing.pdf] [Accessed: 15-Nov-2015]

[2.20] Nilsson, N., (2009). The Quest for Artificial Intelligence: A History of Ideas and Achievements, Cambridge University Press [Retrieved from: https://ai.stanford.edu/~nilsson/QAI/qai.pdf] [Accessed: 20-Nov-2019]

[4.1] Cheng, L., Yu, Tao., (2019). A review and perspective on Machine Learning Technologies Applied to Smart Energy and Electric Power Systems. International Journal of Energy Research. [Retrieved from: https://onlinelibrary.wiley.com/doi/full/10.1002/er.4333#er4333-fig-000] [Accessed: 5-Dec-2018]

[4.2] Feigenbaum, E., Engelmore, R., (1993). Expert Systems and Artificial Intelligence. [Retrieved from: http://www.wtec.org/loyola/kb/c1_s1.htm] [Accessed: 10-Dec-2018]

[4.3] Samuel, A., (1959). Some Studies in Machine Learning Using the Game of Checkers. Retrieved from IBM Journal of Research and Development

[4.4] Source: https://deepmind.com/research/alphago/ [Accessed: 14-Dec-2018]

[4.5] Tweedie, M. (2017). Three Types of AI: Narrow, General and Super AI. Codebots. [Retrieved from: https://codebots.com/ai-powered-bots/the-3-types-of-ai-is-the-third-even-possible] [Accessed: 15-Dec-2018]

[4.6] Bostrom, N. (1998). How Long Before Superintelligence. International Journey of Future Studies. [Retrieved from: https://nickbostrom.com/superintelligence.html] [Accessed: 15-Dec-2018]

[4.7] Poole, D., Mackworth, A. (2017). Artificial Intelligence: Foundations of Computational Agents. Cambridge University Press. [Retrieved from: https://artint.info/html/ArtInt_8.html] [Accessed: 15-Dec-2018]

[4.8] Fagin, R., Moses, Y., Halpern, J., Vardi, M. (2003). Reasoning About Knowledge. MIT Press

[4.9] Cintula, P., Fermiller, C., Noguera, C. (2016). Fuzzy Logic. Stanford Encyclopedia of Philosophy. [Retrieved from: https://plato.stanford.edu/entries/logic-fuzzy/] [Accessed: 15-Dec-2018]

[4.10] Arthur, S. (1959). Some Studies in Machine Learning Using the Game of Checkers. IBM Journal of Research and Development. [Retrieved from: http://www.cs.virginia.edu/~evans/greatworks/samuel1959.pdf] [Accessed: 20-Dec-2015]]

[4.11] Maini, V. (2017). Machine Learning for Humans, Part 2: Supervised Learning. [Retrieved from: https://medium.com/machine-learning-for-humans/supervised-learning-740383a2feab] [Accessed: 20-Dec-2015]

[4.12] Maini, V. (2017). Machine Learning for Humans, Part 3: Unsupervised Learning. [Retrieved from: https://medium.com/machine-learning-for-humans/unsupervised-learning-f45587588294] [Accessed: 20-Dec-2015]

[4.13] Maini, V. (2017). Machine Learning for Humans, Part 5: Reinforcement Learning. [Retrieved from: https://medium.com/machine-learning-for-humans/reinforcement-learning-6eacf258b265] [Accessed: 22-Dec-2015]

[4.14] Hof, R. (2013). Deep Learning. MIT Technology Review. [Retrieved from: https://www.technologyreview.com/s/513696/deep-learning/] [Accessed: 23-Dec-2018]

[4.15] Veloso, M. (2001). Perceptrons and Neural Networks. CMU. [Retrieved from: https://www.cs.cmu.edu/afs/cs.cmu.edu/academic/class/15381-f01/www/handouts/110601.pdf] [Accessed: 23-Dec-2018]

[4.16] Josh, A. (2015). Everything You Need to Know About Artificial Neural Networks. [Retrieved from: https://medium.com/technology-invention-and-more/everything-you-need-to-know-about-artificial-neural-networks-57fac18245a1] [Accessed: 23-Dec-2018]

[4.17] Goodfellow, I., Abadie, J., Mirza, M., Xu, B., Farley, D., Ozair, S., Courville, A., Bengio, Y., (2014). Generative Adversarial Networks. [Retrieved from: https://arxiv.org/abs/1406.2661] [Accessed: 25-Dec-2018]

[5.1] Collier, M., Fu, R., Yin, L., Christiansen, P., (2017). Artificial Intelligence: Healthcare's New Nervous System. [Retrieved from https://www.accenture.com/_acnmedia/PDF-49/Accenture-Health-Artificial-Intelligence.pdf#zoom=50] [Accessed: 14-Jan-2019]

[5.2] Loria, K., (2019). Putting the AI in Radiology. Radiology Today.

[5.3] Schier, R., (2018). Artificial Intelligence and the Practice of Radiology: An Alternative View. JACR; 15(7): 1004–1007.

[5.4] Rajpurkar, P., Irvin, J., Ball, RL., Zhu, K., Yang, B., Mehta, H., et al. Deep Learning for Chest Radiograph Diagnosis: A Retrospective Comparison of CheXNeXt to Practicing Radiologists. PLoS Med. 2018;15(11): e1002686. [Retrieved from https://journals.plos.org/plosmedicine/article?id=10.1371/journal.pmed.1002686] [Accessed: 15-Jan-2019]

[5.5] Du, X., Li, W., and Hu, B., (2018). Application of Artificial Intelligence in Ophthalmology. International Journal of Ophthalmology. [Retrieved from https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6133903/] [Accessed: 20-Jan-2019]

[5.6] Gulshan V., Peng, L., Coram, M., Stumpe, M., et.al, (2016). Development and Validation of a Deep Learning Algorithm for Detection of Diabetic Retinopathy in Retinal Fundus Photographs. [Retrieved from https://static.googleusercontent.com/media/research.google.com/en//pubs/archive/45732.pdf] [Accessed: 22-Jan-2019]

[5.7] Hillman, L. (2017). Microsoft and International Group of Eye Institutes Team Up to Develop Real-Time Machine Learning. [Retrieved from: https://www.eyeworld.org/microsoft-and-international-group-eye-institutes-team-develop-real-time-machine-learning] [Accessed: 25-Jan-2019]

[5.8] Health IT. (2017). Telemedicine and Telehealth. [Retrieved from: https://www.healthit.gov/topic/health-it-initiatives/telemedicine-and-telehealth] [Accessed: 30-Jan-2019]

[5.9] Cooper, T. (2018). Embracing AI to Help Doctors Develop Communities of Care. [Retrieved from: https://www.ibm.com/blogs/client-voices/embracing-ai-to-help-doctors-develop-communities-of-care/] [Accessed: 1-Feb-2019]

[5.10] Source: https://www.century.tech [Accessed: 4-Feb-2019]

[5.11] Source: https://www.blippar.com/build-ar/augmented-reality-education [Accessed: 5-Feb-2019]

[5.12] Source: https://www.microsoft.com/en-us/seeing-ai [Accessed: 7-Feb-2019]

[5.13] Source: https://thefinancialbrand.com/72653/artificial-intelligence-trends-banking-industry/ [Accessed: 8-Feb-2019]

[5.14] Source: https://www.kavout.com [Accessed: 9-Feb-2019]

[5.15] Source: https://www.zestfinance.com [Accessed: 10-Feb-2019]

[5.16] Source: https://www.zestfinance.com/snake-oil-explainability [Accessed: 10-Feb-2019]

[5.17] Source: https://www.betterment.com [Accessed: 11-Feb-2019]

[5.18] Khentov, B. (2014). Lower Taxes with our Improved Cost Basis Accounting Method. [Retrieved from: https://www.betterment.com/resources/lowering-your-tax-bill-by-improving-our-cost-basis-accounting-methods/] [Accessed: 11-Feb-2019]

[5.19] Source: https://www.vectra.ai [Accessed: 14-Feb-2019]

[5.20] Source: https://www.vectra.ai/product/how-it-works [Accessed: 14-Feb-2019]

[5.21] Source: https://zorroa.com/homepage [Accessed: 17-Feb-2019]

[5.22] Source: https://www.netflix.com [Accessed: 20-Feb-2019]

[5.23] Source: https://www.thenorthface.com [Accessed: 22-Feb-2019]

[5.24] Harris, R. (2016). The North Face Brings AI to Ecommerce. [Retrieved from: http://marketingmag.ca/brands/the-north-face-brings-ai-to-ecommerce-165328/] [Accessed: 22-Feb-2019]

[5.25] Lou, H. (2017). AI in Video Games: Towards a More Intelligent Game. [Retrieved from: http://sitn.hms.harvard.edu/flash/2017/ai-video-games-toward-intelligent-game/] [Accessed: 26-Feb-2019]

[5.26] Source: https://civilization.com [Accessed: 26-Feb-2019]

[5.27] Source: https://robotics.kawasaki.com/en1/applications/robotic-assembly/index.html [Accessed: 28-Feb-2019]

[5.28] Source: https://www.loro.xyz [Accessed: 28-Feb-2019]

[5.29] Source: https://waymo.com [Accessed: 28-Feb-2019]

[5.30] Bansal, M., Krizhevsky, A., Ogale, A. (2018). ChauffeurNet: Learning to Drive by Imitating the Best and Synthesizing the Worst. [Retrieved from: https://arxiv.org/pdf/1812.03079.pdf] [Accessed: 28-Feb-2019]

[5.31] Source: http://www.bluerivertechnology.com [Accessed: 7-April-2019]

[5.32] Source: https://plantvillage.psu.edu [Accessed: 8-April-2019]

[5.33] Alcober, F. (2018). AI takes root, helping farmers identify diseased plants. [Retrieved from: https://www.blog.google/technology/ai/ai-takes-root-helping-farmers-identity-diseased-plants/] {Accessed: 8-April-2019]

[5.34] Source: https://www.awhere.com [Accessed: 10-April-2019]

[6.1] Edmonds, D. (2013). Would You Kill the Fat Man: The Trolley Problem and What Your Answer Tells Us about Right and Wrong. Princeton University Press.

[6.2] Source: http://moralmachine.mit.edu [Accessed: 14-April-2019]

[6.3] Awad, E., Dsouza, S., Kim, R., Schulz, J., Henrich, J., Shariff, A., Bonnefon, J., Rahwan, I. (2018). The Moral Machine Experiment. Nature. [Retrieved from: https://www.nature.com/articles/s41586-018-0637-6] [Accessed: 14-April-2019]

[6.4] Taylor, M. (2016). Self-Driving Mercedes Benz Will Prioritize Occupant Safety over Pedestrian. Car and Driver. [Retrieved from: https://www.caranddriver.com/news/a15344706/self-driving-mercedes-will-prioritize-occupant-safety-over-pedestrians/] [Accessed: 15-April-2019]

[6.5] BMVI. (2017). Ethics Commission. Automated and Connected Driving. [Retrieved from: https://www.bmvi.de/SharedDocs/EN/publications/report-ethics-commission.pdf?__blob=publicationFile] [Accessed: 16-April-2019]

[6.6] Asaro, P. (2016). The Liability Problem for Autonomous Artificial Agents. AAAI Symposium on Ethical and Moral Considerations in Non-Human Agents. Stanford University. [Retrieved from: http://peterasaro.org/writing/Asaro,%20Ethics%20Auto%20Agents,%20AAAI.pdf] {Accessed: 17-April-2019]

[6.7] Middleton, C. (2018). Cambridge Analytica vs Facebook: Why AI laws are Inadequate. [Retrieved from: https://internetofbusiness.com/cambridge-analytica-vs-facebook-why-ai-laws-are-inadequate/] [Accessed: 17-April-2019]

[6.8] Dutton, T. (2018). An Overview of National AI Strategies. [Retrieved from: https://medium.com/politics-ai/an-overview-of-national-ai-strategies-2a70ec6edfd] [Accessed: 18-April-2019]

[6.9] Manyika, J., Bughin, J. (2018). The Promise and Challenge of the Age of Artificial Intelligence. McKinsey Global Institute. [Retrieved from: https://www.mckinsey.com/featured-insights/artificial-intelligence/the-promise-and-challenge-of-the-age-of-artificial-intelligence] [Accessed: 18-April-2019]

[7.1] Rhodes, D., Nightingale, D. (2015). Architecting the Future Enterprise. MIT Press

[7.2] Bughin, J., Hazan, E., Ramaswamy, S., Chui, M., Allas, T., Dahlstrom, P., Henke, N., Trench, M. (2017). Artificial Intelligence: The Next Digital Frontier. McKinsey Global Institute. [Retrieved from: https://www.mckinsey.com/~/media/McKinsey/Industries/Advanced%20Electronics/Our%20Insights/How%20artificial%20intelligence%20can%20deliver%20real%20value%20to%20companies/MGI-Artificial-Intelligence-Discussion-paper.ashx] [Accessed: 19-April-2019]

[7.3] Miller, J. (2017). AI Demystified. Disruptive Influence. Callaghan Innovation. [Retrieved from: https://www.callaghaninnovation.govt.nz/blog/ai-demystified] [Accessed: 20-April-2019]