Bachelor’s Thesis

Bachelor of Engineering, Information Technology

2019

Jani Salminen

OVERVIEW OF ARTIFICIAL INTELLIGENCE & MACHINE LEARNING

TURKU AMK

TURKU UNIVERSITY OF APPLIED SCIENCES

BACHELOR'S | ABSTRACT

TURKU UNIVERSITY OF APPLIED SCIENCES

Bachelor of Engineering, Information Technology

2019 | Thesis 30 pages

Jani Salminen

OVERVIEW OF ARTIFICIAL INTELLIGENCE & MACHINE LEARNING

Artificial intelligence is a subsection of computer science and machine learning is a subsection of artificial intelligence. They are arguably one of the largest and most influential technologies which have a fascinating and important role in the future of technology and development as a part of self-learning and functioning, devices and software.

Artificial intelligence has quite a fascinating history, since some of the technologies and software has started their development and use as early as 1950s. However, processor performance was a major bottleneck for AI back then. With modern day technology, performance of processors and memory has improved exponentially, which made it possible for the development of AI to keep breaking its limits. If we take the smartest computing machine known to humans, their own brain, and try to imitate that artificially, it would require insane amounts of data to be processed for it to actually function nowhere near the nature of the human brain.

The purpose of this thesis is to study artificial intelligence briefly, but more comprehensively focusing on machine learning and grasp it's history, technologies, terms, safety and ethics plus give a futuristic view of the subject.

What is AI? Where is it being used? What will AI do or become in the future? These are mind-boggling questions that will be focused on in this thesis.

KEYWORDS:

Artificial intelligence, deep learning, machine learning, technology, futuristic, science

OPINNÄYTETYÖ AMK | TIIVISTELMÄ

TURUN AMMATTIKORKEAKOULU

Tieto- ja viestintätekniikan Insinööri

2019 | Opinnäytetyö 30 sivua

Jani Salminen

YLEISKATSAUS TEKOÄLYYN & KONEOPPIMISEEN

Tekoäly on tietotekniikan osa-alue ja koneoppiminen on tekoälyn osa-alue. Tekoäly ja koneoppiminen ovat väittämättä yksiä vaikuttavimmista teknologioista ja toimivat tärkeässä ja kiehtovassa roolissa tulevaisuuden teknologian kehityksessä, osana itsestään oppivia ja toimivia laitteita, sekä ohjelmistoja.

Tekoälyllä on varsin mielenkiintoinen historia, sillä termiin liittyviä teknologioita ja sovellutuksia on alettu kehittämään ja käyttämään jo 1950- luvulta lähtien. Tuolloin prosessoriteknologia toimi pullonkaulana tekoälyn kehityksessä, mutta nykyteknologian avulla prosessorien ja muistien laskentatehot ovat kehittyneet eksponentiaalisesti ja tekoälyn kehittäminen on jatkuvassa kasvussa. Tänäpäivänä mahdollisuudet ovat jo paljon suuremmat. Esimerkkinä historian fiksuin laskentajärjestelmä, ihmisten aivot ja niiden neuroverkot ovat olleet pitkään simuloinnin kohteena neuroverkkoteknologian avulla, mutta täydellinen ihmisten aivojen simulointi vaatisi käsittämättömän määrän laskentatehoa, jotta tuloksena olisi lähellekään samoin toimiva, päätöksiä tekevä laite tai sovellus.

Tämän opinnäytetyön tarkoituksena on tutkia tekoälyä kokonaisvaltaisesti, mutta keskittyä tarkemmin koneoppimiseen ja tuoda esille sen historia, käyttötarkoitukset, teknologiat, termistöt, sekä turvallisuuteen ja eettisyyteen liittyvät kysymykset, sekä tulevaisuuden näkökulma.

Mitä on AI? Missä sitä käytetään? Mitä se tekee tai mitä siitä tulee tulevaisuudessa? Nämä ovat mieltä painavia kysymyksiä, joihin yritetään etsiä vastauksia.

ASIASANAT:

Tekoäly, syväoppiminen, koneoppiminen, teknologia, futuristinen, tiede

CONTENTS

LIST OF ABBREVIATIONS

AI	Artificial intelligence
Algorithm	Definition, functionality, specification
AlphaGo	Software that plays the board game Go
Analytical model building	Mathematical models, functions
(Artificial) Neural network	Enables machines to work independently
Barcode	Data in machine readable form
Big Data	Field of data analysis, huge sets of data
Clustering	Technique used in unsupervised learning
Cognitive	Understanding through thought
Cognitive performance	Cognitive skill
Convergence	The process of converging (Meeting point)
Data	Data, Information
Data analysis	Process of data inspecting and modeling
Decision tree	Technique for supervised learning
Deep learning	Machine learning method for neural networks
Heuristics	Problem solving, technique in computer science
IoT	Internet of Things
Iterative	Repetition of a process
K-Means, K-nearest	Clustering algorithms
Labeled Data	Known, informative data
Layering	Functional components working together
Lisp	Functional programming language
Linear regression	Machine learning algorithm
Linguistic	Relating to language
Logistic regression	Machine learning algorithm
Machine vision	Ability for computers to ‘see’
Moore’s Law	Moore’s interpretation of the amount of transistors

Naïve Bayes	Algorithm using Bayes' theorem
Pattern recognition	Recognizes regularities in data
Pattern matching	Checks similarity in data, used in computer science
PanelScan	Efficient barcode reader
Perceptron	Algorithm for supervised learning
Robotics	Technical science field
Random forest	Classification algorithm
Reinforcement learning	Area of machine learning
Robot	Mechanical machine
Singularity	Superintelligence
Superintelligence	Intelligence surpassing human level intelligence
Supervised learning	Area of machine learning
SVM	Support Vector Machine
Tera, Peta	Petabit = 1000 terabits
Unsupervised learning	Area of machine learning
Volume	Amount of data, term of big data
Velocity	Fast rate data, term of big data
Variety	Different types of data, term of big data
Vitrification	Process of turning or changing matter to glass

1 INTRODUCTION

Artificial intelligence is a huge subject in public these days, which is being wildly debated and the term AI disappears in the context where its being related to various different cases and technologies. People tend to define the idea behind AI wrong for what it really stands for, using logics and algorithms to create complicated functions for machines to calculate and operate in a way which could extremely be comparable to the way of human behavior, thinking and problem solving.

Artificial intelligence, which contains intelligence, already a very hard term to describe, taking that as a bag of functions and trying to copy or imitate it artificially, surely can make one's imagination fly. So it is no surprise that artificial intelligence can interest and scare people in different ways beyond imagination. Some may believe it is merely used for simple calculations or functioning software in everyday life, others take it as the next uprising of the Terminator series. Whatever the future for AI is, surely is interesting and it will be a huge part of the future of mankind and constantly evolving technologies.

This thesis tries to cover as much grass as possible, but still mainly focusing on machine learning. After introduction in chapter two is the history of artificial intelligence. Section three is about machine learning and its sub-sections. Section four is about some technologies and sub-sections of artificial intelligence. Section five contains superintelligence, whole brain emulation plus where and why artificial intelligence appears and affects. Section six brings up programming languages and after that in sections seven and eight moral and ethics are discussed.

2 HISTORY

One of the founders of artificial intelligence as a concept was an US computer scientist and cognitive scientist Dr. John McCarthy. In 1956, Dartmouth conference was held and McCarthy presented the idea and came up with the term artificial intelligence as “the science and engineering of making intelligent machines”. After the conference, the first wave really took off, mainly focusing on research. In addition to McCarthy, many of these researchers became household names in those early research days and today still are. Among them: Marvin Minsky, Herbert Simon and Allenn Newell, who were promoting AI as a field of computer science that could transform the world. (i-SCOOP, NicoElNino, 2018)

2.1 From 1950 to 2019

After Dartmouth conference in 1956, it took McCarthy 2 years to come up with his own programming language called Lisp, which would become a norm in AI programming. Lisp ended up revolutionizing programming, after all these years it’s still being used in robotic programming and several browser-based services. Lisp was intended to process lists of data. The computer would navigate and process lists of sentences to mimic human reasoning, the ability to answer a question by comparing possible sentences organized to lists. (Kidscodecs 2014)

Artificial intelligence showed less progress from mid 1970s to mid 1990s even with well-funded global effort, since scientists found it incredibly difficult to create intelligence in machines. AI applications require processing of enormous amounts of data. 30 to 40 years back when the algorithms were still primitive and non-efficient, hardware was a huge bottleneck in the progress of AI. These years became known as the *AI Winters*. (Towards Data Science 2018)

Another wave of AI took place in the nineties and up to early 2000s. Computer technology had seen remarkable leaps in hardware performance and made it possible for systems to process bigger amounts of data. There was now more attention from the perspective of applications of AI across diverse domains. On top of this there was the success of the Internet, which led into hype and predictions in the public. It was also those days that the convergence of man and machine became increasingly popular. (i-SCOOP, NicoElNino, 2018)

In the late 1990s, more companies became interested of AI and started to raise funds. Japanese government unveiled plans to develop a fifth generation computer to advance in machine learning. AI enthusiasts believed in computer conversations, translations, interpreting pictures and reasoning like people do. IBM's Deep Blue was the first computer to beat the world chess champion, Garry Kasparov. (Towards Data Science 2018)

Today's artificial intelligence wave is a rapid adoption of AI technologies including cloud, faster processing capabilities, scalability, Big Data, Internet of Things, self-driving cars, robotics, chatbots and more. Common problem across all of these technologies is the overwhelming amount of data. Without AI there isn't that much to do with data to put it simply. (i-SCOOP, NicoElNino, 2018)

AI has lots of potential in the modern world and is already affecting everyday lives by a large margin and due to the nature of the subject and the reason that it looks like its progress wont be stopped anytime soon makes AI a very interesting and up to date topic. This also causes AI to raise lots of ethical, safety, and future related questions.

Market demand for these new technologies is huge and is growing exponentially. During the write of this project, programmers are figuring out new ways to create more and more efficient systems to improve aspects of AI for the future. Machine learning is one of the biggest subsections of AI and that will be introduced in the next chapter.

3 MACHINE LEARNING

Machine learning is a method of data analysis that automates analytical model building. It is a branch of artificial intelligence and is based on the idea that systems could learn from data, identify patterns and make positive decisions with none or minimal human intervention. (SAS 2019) Machines can be considered autonomous or even self-conscious, when they work with complicated algorithms instead of basic, iterative and limited code. The iterative aspect of machine learning is important because as machines or algorithms are exposed to new data, they are able to independently adapt. They learn from previous computations to produce reliable, repeatable decisions and results. Machine learning algorithms have been around for a long time, but the ability to automatically apply complex mathematical calculations to big data over and over, faster and faster is a recent development. (SAS 2019)

Self-driving Google car or a computer software that beats gamers are great examples of AI and machine learning. Recommendations when browsing through the internet such as Amazon discounts or Spotify daily mixes are all applications of machine learning in everyday life. Knowing people's reactions on Twitter or LinkedIn is machine learning combined with linguistic rule creation. Fraud detection is also machine learning applications and algorithms. These are just few of many, many applications that affect the internet and the world today. (SAS 2019)

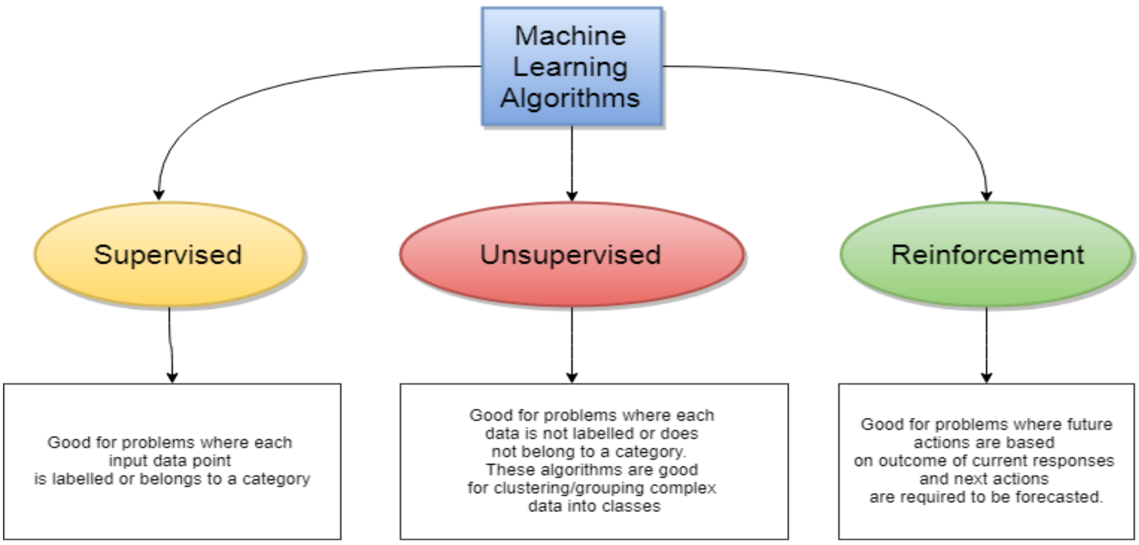

Figure 1. Types of machine learning algorithms. (Miro Medium)

3.1 Unsupervised learning

Unsupervised learning is a branch of machine learning and the opposite of supervised learning. Unsupervised learning learns from data that has not been labeled, classified or categorized. Instead of responding to feedback, unsupervised learning identifies commonalities in the data and reacts based on the presence or absence of such commonalities in each new piece of data. (Towards Data Science 2018)

So with this method the machine can find structures in the data and categorize data by comparing them by their similarities for example. This makes unsupervised learning attractive in applications where data is cheap to obtain, but labels are either expensive or not available. (Science Direct 2014)

3.1.1 Clustering

Clustering is a method of unsupervised learning and is a common technique for statistical data analysis in many fields. Clustering analysis can be used to gain valuable insight from data by seeing what groups the data points fall into when applied to a clustering algorithm. (Towards Data Science 2018)

K-means is one of the most popular algorithms used for clustering. K stands for the number of clusters in your algorithm. K centroids are clusters that the algorithm finds for each data point. Once K-means creates these clusters, finds mean of each cluster and then re-adjusts the points, this step is then repeated until points do not move from clusters, which is called convergence.

This type of algorithms are used to categorize data, find abnormal data items or detect anomalies e.g. finding fraud transactions from a large set of transactions. (Medium 2018)

3.2 Supervised learning

This family of algorithms can be used to find relationships between data. Unlike unsupervised learning, supervised learning usually treats labeled data, or known sets of data and the machine creates predictions based on the algorithms being used. A supervised learning algorithm takes a known set of data and its known responses to the data, which is output data, to learn the regression or classification model. Classification algorithms and regression techniques are used to develop these predictive models. The algorithms include linear regression and logistic regression, apart from decision tree, SVM, random forest, Naïve Bayes, and k-nearest neighbor. Popular or major applications of classification include bank credit scoring, medical imaging, and speech recognition.

3.3 Reinforcement learning

Reinforcement learning is the science of decision making with the goal to select actions to maximise total future reward. The machine learns how to respond to an environment by getting positive responses during it’s work. The machine optimizes its decisions by the feedback it gets, towards a more positive outcome. (Towards Data Science 2018) The following chapter is reinforcement learning at its core. In the next chapter theres a comprehensive example of a machine that learned to beat the world champion in the board game Go.

3.3.1 AlphaGo

AlphaGo is a software developed by DeepMind Technologies that plays the board game Go. Go is known as the most challenging classical game for AI because of its complexity. The goal is to surround and capture opponent’s stones or strategically create spaces of territory. Once all possible moves have been played out, the stones on the board and the empty points are tallied and the highest number wins. Sounds simple, but it surely is more complex than chess, since Go has an astonishing 10^170 possible board configurations, which is more than the number of atoms in the known universe.

AlphaGo Zero is the successor to earlier AlphaGo versions and surpassed its capabilities by a large margin. It was only given the basic input of the game rules and in 40 days, managed to surpass thousands of years of strategy and analyzing by even adding in its very own strategies and moves.

AlphaGo Zero becomes its own teacher because of reinforcement learning. The system starts off with a neural network that knows nothing about the game. Then it repeatedly plays the game against itself, by combining this neural network with a powerful search algorithm. This neural network gets tuned and updated to predict moves. The updated neural network creates new, stronger versions of AlphaGo Zero, and in every new iteration of this process, the performance of the system increases by a small amount, and the quality of the games against itself increases, leading to even more accurate neural networks and stronger version of the software.

This makes AlphaGo Zero much more powerful than previous versions of AlphaGo, because it's no longer constrained by the limits of human knowledge and instead, it learns from the strongest player in the world: AlphaGo itself. (DeepMind 2019)

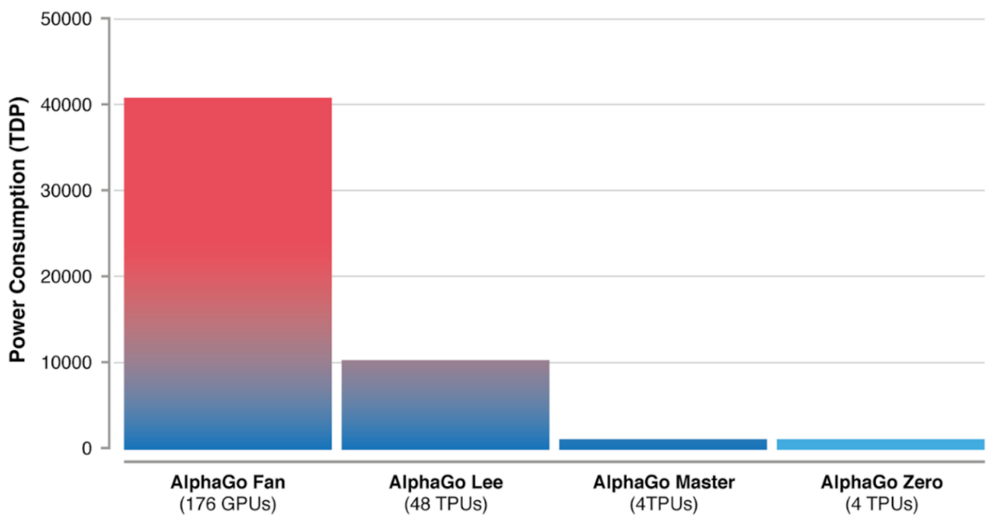

Figure 2. Interesting table showing the progress in efficiency from 2015 to 2019 (DeepMind 2019)

4 TECHNOLOGIES & TERMS

4.1 Neural networks & deep learning

Deep learning makes use of artificial neural networks that behave similar to the neural networks in human brain. A neural network functions when some input data is fed to it. This data is then processed via layers of perceptrons to produce a desired output.

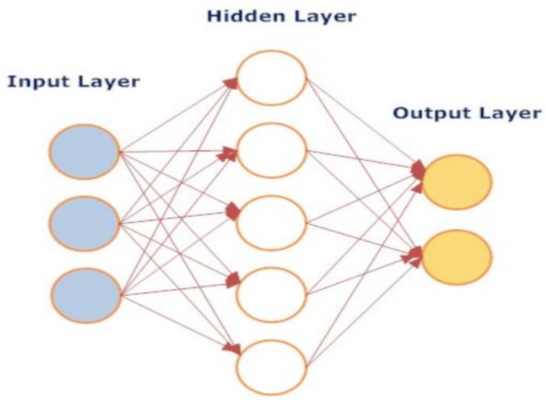

Figure 3. Example of layering in a simple neural network. (Marketing Insider Group)

With the help of deep learning, Google can instantly translate between more than 100 different languages, convert text to speech, speech to text and even produce software for self driving cars.

Layering and depth in the neural networks allows the network to learn complicated structures and rules based on input data. Each layer decides and delivers attributes to the next layers. Unlike traditional architectures it’s possible to modify and save data inside the neural networks. Long term memory inside the neural networks is based on a weight factor between the data cycles.

4.1.1 Pattern recognition

Pattern recognition is a branch of machine learning that is very present in modern day as it is used to recognize patterns and shapes from data. In IT, pattern recognition can be supervised, where previously known patterns can be found in given data, or unsupervised, where entirely new patterns will be discovered. Pattern recognition should not be confused with pattern matching, since pattern matching looks exactly for the same pattern and is not even considered part of machine learning, even though in some cases it can lead to similar results as pattern recognition. (DZone 2018)

4.1.2 Machine vision

Machine vision is widely used especially in industrial applications. Sensors and cameras are used to imitate human eyes to capture images, which are processed by software to analyze and measure various characteristics for decision making. (Cognex 2019)

There are many applications for machine vision in manufacturing, a simple example is barcodes. Reading, identifying and processing countless of barcodes per day is not something for the human eye to do at a large scale. The ever-growing tech market requires PCB manufacturers to step up their game, so a machine vision based solution PanelScan was developed for reading many barcodes at once. Before, a human would do this task by using a handheld barcode scanner to scan one barcode at a time. (Dev Team Space 2019)

4.2 Big Data

Big data in IT literally means huge amounts of unlabeled datamasses. There is however, something to be found from these masses of data, if it’s used wisely.

The three versions of Big data: Volume, Velocity and Variety.

Volume is the amount of data. You will have to process high volumes of low-density, unlabeled data. This can be from social media feeds, clickstreams from a webpage or some sensor-enabled equipment. For some use cases, this may be tens of terabytes of data, for others, hundreds of petabytes of data.

Velocity is fast rate data, where it’s received and maybe acted on. Normally the high velocity data streams directly into memory, rather than written on a disk. Some internet smart products operate almost in real time and will require immediate evaluation and action.

Variety refers to the types of data that is available. With the rise of big data, data comes in new unlabeled data types. Unlabeled and semilabeled data types such as txt, audio and image, require additional processing to obtain meaningful and supportive metadata. (Oracle 2019)

4.3 Robotics

Robotics is the science behind robots, which contains of electronic engineering, mechanical engineering, information engineering and computer science. However, robotics and AI are two almost completely different things unless they are combined, which is in the least of cases. While robots are programmable and usually semi-autonomous, they are working in environments with sensors and actuators and are able to carry on with iterative tasks only. Tasks which do not require AI at all. Artificially intelligent robots on the other hand have been programmed with algorithms to perform more complex tasks and to learn from their mistakes. Say a vacuum robot which is programmed to go through the rooms and using its sensory data to avoid obstacles and this way vacuum a house. Well this is not very advanced nor efficient and is definitely not an artificially intelligent robot at this point. When software is added to make the robot scan for the room’s size, identifying obstacles and remembering these by mapping everything, it may also find the most efficient routes for vacuuming. This self-deploying robot can then determine how much vacuuming there is based on these values and instead of roaming mindlessly all around it can perform even without human assistance. (BuiltIn 2018)

4.4 Linear Regression

Regression techniques are used to create models between variables. One of these algorithms is linear regression. Goal here is to set the best line between the dots so that it minimizes the distance between each one (Figure 4). (Commonly used Machine Learning Algorithms 2017)

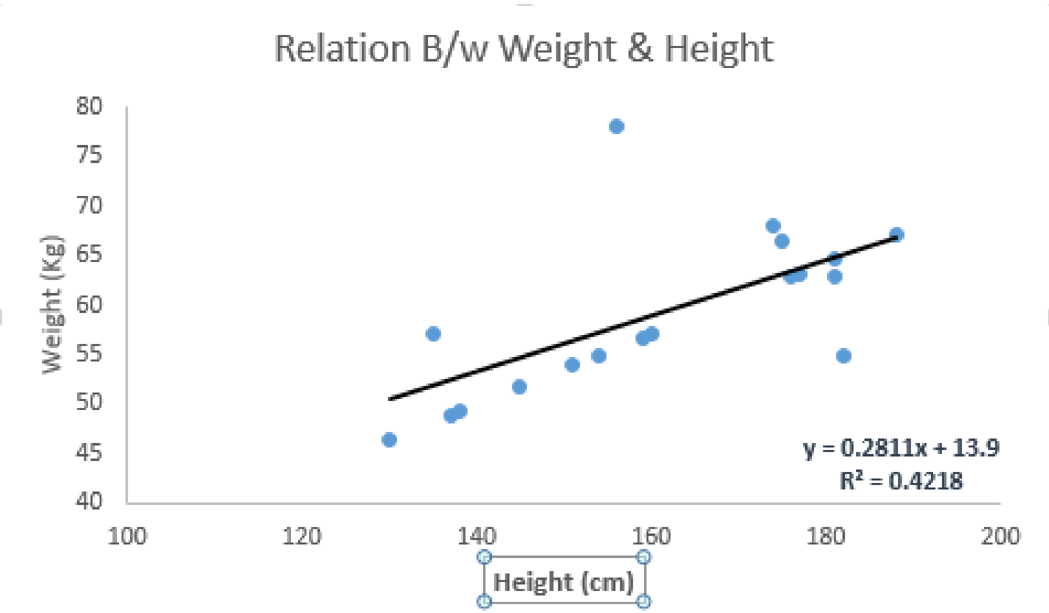

Figure 4. Finding weight, knowing the height of a person. (Commonly used Machine Learning Algorithms 2017)

4.5 Logistic Regression

Logistic regression is a classification algorithm. It is being used to estimate values like 0/1, yes/no or true/false. These estimates are based on given set of independent variables. It predicts the probability of occurrence of an event by fitting data to a logit function. (Commonly used Machine Learning Algorithms 2017)

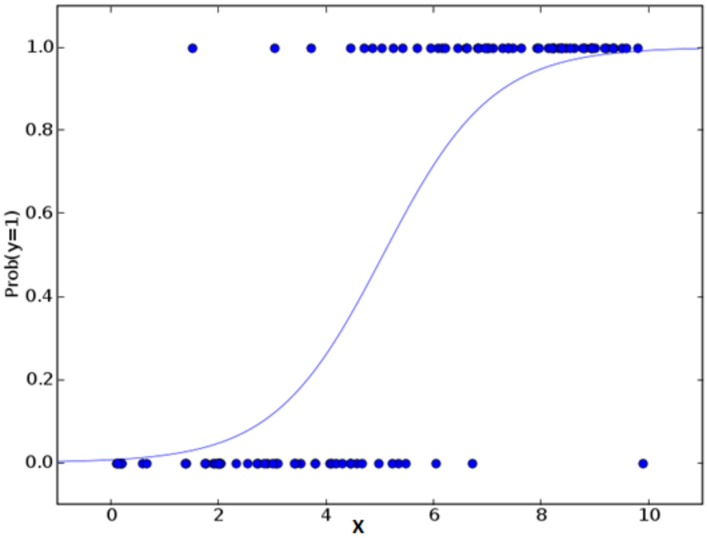

Figure 5. Logistic regression algorithm. (Commonly used Machine Learning Algorithms 2017)

4.6 Decision Tree

Decision trees are one of the most often used algorithms in machine learning. It's decisions are based on the data available, by showing decisions in a tree-like form. The population is divided into four different groups based on multiple attributes to identify if they will play or not (Figure 6). (Commonly used Machine Learning Algorithms 2017)

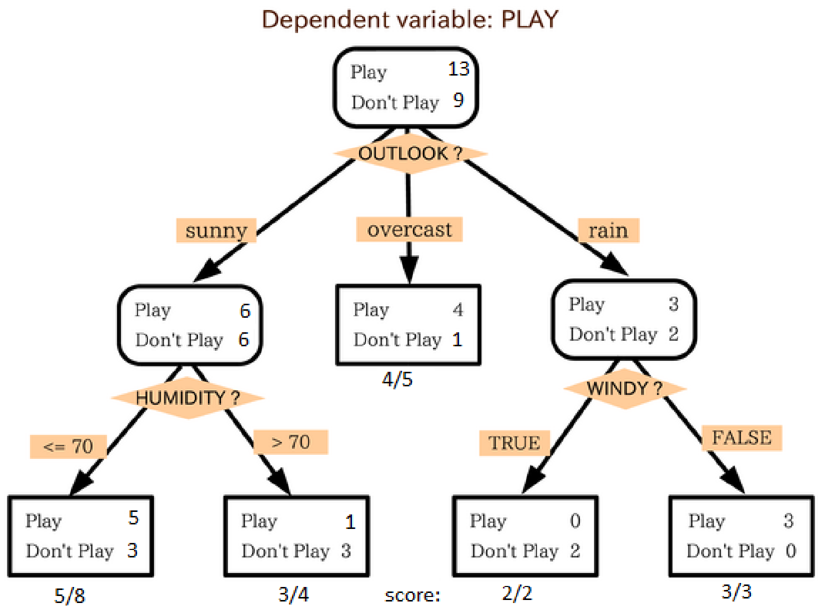

Figure 6. Decision tree algorithm. (Commonly used Machine Learning Algorithms 2017)

5 ARTIFICIAL INTELLIGENCE WHERE AND WHY

At this point of the study it is safe to say that AI is indeed very present in the modern world. It's just quite good at hiding behind the scenes. In principle, it's machine's and software's capability to react in a way that is atleast somewhat comparable to the way of the human thinking.

Scientific explanation is usually a computer which has achieved self-consciousness and learns by itself. This varies from the traditional way of thinking, where computer can only do things as it has been taught, running an algorithm or iterative code given to it by its operator.

However AI can be an independent machine or software which can do more than just simple given operations. Analytics softwares are a great example. This type of software is dependent on massive data storages like big data, human taught data identifications, algorithms and heuristics. They may use parameters such as the amount of processes, transferred data, latency, system's history, date, time of day or used software. Analytics software may use all this data to come to the conclusion if there's a potential problem or just a regular one. The software may then draw conclusions about the information they have available or given in these data sets. This way large amounts of formerly useless data can be analyzed and used for something productive, like improving industrial business, sales or healthcare for example.

5.1 Superintelligence

Superintelligence or Singularity sounds exciting and has mostly been heard and seen in movies and fairytales. It's however, something that has been and will be a topic ever since AI has come around, afterall AI stands for "the science and engineering of making intelligent machines" as McCarthy originally stated.

What we can define as superintelligence is any intellect that exceeds the cognitive performance of humans in virtually all areas. "Observations have been made by philosophers and scientists, including David Chalmers and Hans Moravec, to argue that human-level AI is not only theoretically possible but feasible within this century" (Nick Bostrom 2014, 23).

“Evolutionary aspect for the feasibility of AI is the idea that we could, by running genetic algorithms on fast enough computers, to achieve results comparable to those of biological evolution. The downside is that the computer performance required to simply replicate the required evolutionary processes on Earth that produced human-level intelligence are severely out of reach and will remain so even if Moore’s law were to continue for a century” (Nick Bostrom 2014, 24).

The fact that there are many paths that could eventually lead to superintelligence increases the confidence to get there.

5.2 Whole brain emulation

Achieving whole brain emulation requires some inspiration from nature: barefaced plagiarism. These steps could be needed to achieve this.

First a sufficiently detailed scan of a particular human brain is created. This might involve stabilizing the brain post-mortem through vitrification. A machine could then dissect the tissue into thin slices, which could be fed into another machine for scanning, perhaps by an array of electron microscopes. Various stains might be applied at this stage to bring out different structural and chemical properties. (Nick Bostrom 2014)

Next the raw data from the scanners is fed to a computer for automated image processing to reconstruct the three-dimensional neural network that implemented cognition in the original brain. The resulting neural network is then combined with a library of neurocomputational models of different types of neurons or of different neuronal elements such as somekind of synaptic connectors. (Nick Bostrom 2014)

In the third stage, the neurocomputational structure is implemented on a sufficiently powerful computer. If completely successful, the result would be a digital reproduction of the original intellect, with memory and personality. The emulated human mind now exists as software on a computer. The mind can either inhabit a virtual reality or interface with the external world by means of robotic appendages. The whole brain emulation path does not require that we figure out how human cognition works or how to program an artificial intelligence. It requires only that we understand the low-level functional characteristics of the basic computational elements of the brain. No fundamental conceptual or theoretical breakthrough is needed for whole brain emulation to succeed. There is a good reason to think that the requisite enabling technologies are attainable, though not in the near future. Pattern and image recognition software has been developed that can trace axons and dendrites through a stack of two-dimensional images though these are not reliable yet. (Nick Bostrom 2014)

Capabilities needed for whole brain emulation consist of scanning, translation and simulation. Capability of scanning at sufficient resolution and to scan entire brain volumes in reasonable time and expense, identifying cell types, synapses and their connectivity, storing information to databases and so on, would however require insane amount of performance and accuracy. (Nick Bostrom 2014)

A great expansion of neurocomputational libraries and major improvements in automated image processing and scan interpretation would also be needed. In general, whole brain emulation relies less on theoretical insight and more on technological capability than artificial intelligence. Just how much technology is required for whole brain emulation depends on the level of abstraction at which the brain is emulated. (Nick Bostrom 2014)

The availability of the brain as a template provides strong support for the claim that superintelligence is ultimately feasible. This does not enable to predict when it will be achieved because it is hard to predict the future discoveries in science. The further into the future we look, the greater the likelihood may be that the secrets of the brain’s functionality will be decoded sufficiently enough to enable the creation of superintelligence in this manner. (Nick Bostrom 2014)

5.3 Applications of AI

AI algorithms and machine learning work every domain of technology today. Speech recognition, machine vision, recommendations, advertisements, search engines, industrial machines, electronics have already strong influences. (Figure.4)

Amazon	One of the leading companies in AI and machine learning. Amazon has managed to create an environment, where machines are working and learning all the time, efficiently serving customers with fully automated warehouses as their goal. Amazon uses AI throughout all of their customer service.
Emotech	Emotech is a company that created Olly, a voice-controlled AI robot similar to Amazon's Alexa or Google's Home, but it has an evolving personality. Olly's personality is a mix of machine learning algorithms so it becomes more and more like its owner. It is the most advanced assistant and can even suggest some of your favorite music if it notices you resting your head after a long day at work.
Facebook	Analyzes data and photos to improve in machine learning.
Conversational/Chatting bots	Virtual assistants or smart assistants in customer service, healthcare or games and entertainment. Few examples are Amazon Alexa, Apple Siri, Google Now, Google Assistant.
Google	Google uses AI in various of softwares. Google Smart Maps are a great example of their work. When you want to get from A to B in a huge city for example, google's algorithms can determine the optimal route to take, let it be on foot, car, bike, bus or train. It can count it all in a matter of seconds.
Tinder	Simple dating app that everyone knows. However, behind the scenes they run different algorithms and machine learning to learn from their users and how to connect people with each other. Machine learning works behind their smart photos, super likeable and top picks.

Figure 5. Example uses of artificial intelligence

5.3.1 Healthcare

Healthcare might be one of the most revolutionary aspect of AI in the future. While artificially intelligent robots are not yet ready to perform full surgical operations, that's still one possible path. AI is not threatening human surgeons, but it's becoming part of the operations by assisting and improving even the best of surgeons performance. AI controlled robots can provide three-dimensional magnification for articulation to perform with more precision. There are still numerous complex tasks during surgical operations, which require the skills of a skilled surgeon, so we will not be seeing a robot perform large surgical operations by it's own anytime soon.

AI has improved healthcare in detecting diseases, for example cancer where the use of AI in analyzing patients and reviewing them with radiology images has sped up the processes up to 30 times with 99% accuracy. Google's DeepMind technology has trained a neural network to detect more than 50 types of eye disease by analyzing 3D retinal scans. (Kolabtree Blog 2014)

IBM Watson and Google DeepMind are the leaders in mining medical records. Their objective is to create a "cognitive assistant" equipped with all the possible clinical knowledge for analytical and reasoning capabilities. Without these applications enormous amounts of data would be wasted, because no human could manually browse through all of the patients clinical data and make decisions based on those, with right algorithms AI can improve itself and make right choices for each patients path.

6 PROGRAMMING LANGUAGES

There are quite a few programming languages used in AI, which of most used ones are Lisp, Python, Prolog, C++ and Java. I'm going to introduce Lisp, since it's one of the oldest high-level programming languages and was invented by John McCarthy in 1958.

6.1 Lisp (List Processing)

Lisp was invented by John McCarthy while he was at the Massachusetts Institute of Technology (MIT) and it's still actively being used, mostly in AI.

Common Lisp is a dialect of Lisp, which appeared between 1980s and 1990s, it's goal was to standardize, unify and extend the features of MacLisp dialects.

Here we have a function called *averagenum* that will print the average of our given four numbers. Functions are shown as lists, where the function's name comes first, and the rest are parameters.

```
(defun averagenum (n1 n2 n3 n4)
   (/ ( + n1 n2 n3 n4) 4)
)
(write(averagenum 10 20 30 40))
```

Figure 6. Example of average number calculator in Lisp language. (Tutorialspoint/Lisp)

"Lisp is unbelievably simple and straight forward as it looks at everything you do as a programmer as just a function call. So if you want to load a file, you call a function that loads the file. If you want to do some math, whether it's addition or finding the square root, you call a function" (Quora, Max Thompson).

7 ETHICS & MORAL

All the flash and hype about AI is accompanied by the negative side. Data privacy issues and ethical usage of AI. Some of the ethical concerns surrounding AI include questions like who would be held accountable for machine errors which lead to mismanagement of healthcare? Would patients in hospitals be informed beforehand how AI is playing in the background with their treatment? How come people find anymore jobs if intelligent machines take all of theirs? Could they eventually lose authority and autonomy over intelligent machines?

Clearly as an upcoming and influential technology, AI concerns people and it requires careful thought and responsible use. If used responsibly the associated ethical and data privacy criteria, AI can potentially lead to an efficient transformation for example how healthcare works. While this transition is in the works, it is important to train and inform professionals of all crafts about the use of AI.

Artificial intelligence and digitalization of the future is described as making things easier or getting rid of tedious repetitive working stages. However if it goes too far, it will put some people out of their jobs simply because machines can handle some tasks more efficiently and cost productively, especially in the future they will. Optimistic view is that these machines need some kind of supervisors and servicians to be taken care of resulting in more jobs.

Huge amounts of data moves around the internet, where algorithms search through these datamasses for valuable information. These datamasses of information contain private data, search information, face recognitions, speech recognitions and valuable private information. Safety related questions are important, as artificial intelligence evolves. How do we protect ourselves from data leaks and hackers?

8 SAFETY

Safety related questions bug people because people value privacy. AI might take away some of this privacy due to its nature of being and functionality. In the near future, a very important problem may need to be solved, since programmers and developers need to create safe software with improved AI in mind. It might not be the end of the world when a laptop crashes or a windows software crashes, but when it is the software that controls a car, a plane, or any other autonomous software with people's life at stake.

Another way that questions safety is when AI is applied in societies without consideration and analysis of the long-term ethical and moral implications it might cause. There's also danger to build a society where certain decisions are made purely by AI algorithms. For example, who will receive jail time for crimes they have conducted? In US, "Criminal risk assessment algorithms" are already being used to analyze whether a person is likely to reoffend in the future. There lies a problem where these tools may make wrong kind of suggestions and send innocent people to prison, or vice versa. (Lasse Rouhiainen 2019)

Hackers may also take advantage of AI by using machine learning and big data as their advantage. Same artificial intelligence technologies that are used to power speech recognition, self-driving cars and deep fake videos have tha ability to turn upside down. For example creating viruses which can morph faster than defenders and antivirus companies can keep up, automatized phishing emails that are indistinguishable from real messages written by actual people and intelligently lurking inside huge companies systems, they only need to get in once. (Why It's Time to Prepare for AI Wielding Hackers 2019)

9 CONCLUSION

AI is part of computer science and machine learning is part of AI. Due to their nature AI or machine learning are not simply defined or classified. Technologies and real-life examples that demonstrate what machine learning is and how it functions have been introduced in this thesis. Some futuristic views have also been shown and questioned.

A friend of mine tends to talk often about psychology and after some of his theories about categorizing people by their way of acting to some life event I convinced him to think more deeply how shallow it’s to define people into few categories just like that. Artificial intelligence seems to be so similar in this aspect, we know something and make conclusions about it, but its roots go way deeper and spread extensively, plus it’s affected by so many external things it’s hard to define or classify when under constant change and evolve.

As AI is making remarkable progress in the daily lives of humans and plays a big role in the future of mankind, realistically thinking it simply cannot have only positive effects surrounding it. People question whether it’s smart at all to apply AI technologies as part of our future. Getting familiar with AI is important, then it’s possible to get an idea how it can be utilized, to make a more realistic approach and not completely alienate it as a dangerous subject that should never have made its appearance on earth or totally accept it with bright blue eyes.

Favourite part of this thesis was the idea of superintelligence and the theories behind it. Superintelligence is mostly speculation, but studying robotics, technologies and features of AI were interesting and fascinating, because it’s futuristic content.

Cutting this thesis down to machine learning felt like a smart idea at first, but realizing how big of a rabbit hole that topic itself was, this thesis had to be limited somehow. Overall view and understanding of the topic felt like a good idea. Ideally it would probably have been a better choice to pick a smaller subsection of machine learning and focus on one software or technology at a time. Now that a general idea of the subject is achieved, work can progress to deeper dive into smaller subsections of the topic.

REFERENCES

Analytics Vidhya, Sunil, R. (2017) Commonly used Machine Learning Algorithms. Accessed 25.10.2019 https://www.analyticsvidhya.com/blog/2017/09/common-machine-learning-algorithms/

BGR, Chris, S (2019) Facebook uses AI to analyze and catalog every single photo you upload. Accessed 7.10.2019 https://bgr.com/2019/07/05/facebook-outage-reveals-ai-catalogs-every-photo-you-upload/

BuiltIn, Sam, D. (2019) Artificial Intelligence Shaking Up Business As Usual. Accessed 3.10.2019 https://builtin.com/artificial-intelligence/examples-ai-in-industry

Cognex, What is Machine Vision? Accessed 30.9.2019 https://www.cognex.com/what-is/machine-vision/what-is-machine-vision

DataCenter Knowledge, (2019) Why It's Time to Prepare for AI Wielding Hackers. Accessed 25.10.2019 https://www.datacenterknowledge.com/machine-learning/why-its-time-prepare-ai-wielding-hackers

DeepMind, (2019) AlphaGo, AlphaGo Zero. Accessed 9.10.2019 https://deepmind.com/research/case-studies/alphago-the-story-so-far

DevTeam.Space, Using Machine Vision in Manufacturing. Accessed 30.9.2019 https://www.devteam.space/blog/10-examples-of-using-machine-vision-in-manufacturing/

DZone, Chandu, S. (2018) Machine Learning and Pattern Recognition. Accessed 30.9.2019 https://dzone.com/articles/machine-learning-and-pattern-recognition

Figure 1, Miro.Medium, Accessed 9.9.2019

https://miro.medium.com/max/951/1*r0Nx8dRjXYeYbCzPxhCUrw.png

Figure 2, Ecommerce-Digest, Accessed 9.9.2019 https://marketinginsidergroup.com/content-marketing/artificial-neural-networks-every-marketer-know/

Figure 3, DeepMind, AlphaGo Zero. Accessed 9.9.2019 https://deepmind.com/blog/article/alphago-zero-starting-scratch

Figure 4, Analytics Vidhya, Commonly used Machine Learning Algorithms. Accessed 25.10.2019 https://www.analyticsvidhya.com/wp-content/uploads/2015/08/Linear_Regression.png

Figure 5, Analytics Vidhya, Commonly used Machine Learning Algorithms. Accessed 25.10.2019 https://www.analyticsvidhya.com/wp-content/uploads/2015/08/Logistic_Regression.png

Figure 6, Analytics Vidhya, Commonly used Machine Learning Algorithms. Accessed 25.10.2019 https://www.analyticsvidhya.com/wp-content/uploads/2015/08/IkBzK.png

IBM, (2018) What is Artificial Intelligence in medicine? Accessed 9.10.2019 https://www.ibm.com/watson-health/learn/artificial-intelligence-medicine

i-SCOOP, agsandrew. (2018) AI: Understanding and debating fears of artificial intelligence. Accessed 5.4.2019 https://www.i-scoop.eu/artificial-intelligence-cognitive-computing/understanding-debating-fears-artificial-intelligence/

i-SCOOP, Roland, S. (2018) AI: Solving the information and big data challenge of artificial intelligence. Accessed 5.4.2019 https://www.i-scoop.eu/artificial-intelligence-cognitive-computing/solving-information-big-data-challenge-artificial-intelligence/

i-SCOOP, NicoElNino. (2018) AI and cognitive computing: what, why and where. Accessed 28.3.2019 https://www.i-scoop.eu/artificial-intelligence-cognitive-computing/

KDnuggets, Alex, C. (2016) AI and deep learning, neural networks explained. Accessed 30.3. 2019 https://www.kdnuggets.com/2016/10/artificial-intelligence-deep-learning-neural-networks-explained.html

Kidscodecs, Tim, S. (2014) Lisp. Accessed 5.4.2019 https://www.kidscodecs.com/lisp/

Kolabtree Blog, Maya, R. (2018) 5 Real-World Examples of AI in Healthcare. Accessed 8.10.2019 https://blog.kolabtree.com/5-real-world-examples-of-ai-in-healthcare/

Lasse Rouhiainen, Lasse, R. (2019) 13 Commonly Asked Questions on Artificial Intelligence for 2019. Accessed 24.10.2019 (https://www.lasserouhiainen.com/7-questions-artificial-intelligence/

Medium Corporation, Bhagirath, Kumar, L. (2018) Some Thoughts on Artificial Intelligence Singularity. Accessed 16.4.2019 https://medium.com/datadriveninvestor/some-thoughts-on-artificial-intelligence-singularity-3f16db2ae8ae

Medium, Farhad, M. (2018) Machine Learning Algorithms Comparison. Accessed 22.9.2019 https://medium.com/fintechexplained/machine-learning-algorithm-comparison-f14ce372b855

Neoteric, Kaja, P. (2019) Tinder algorithms: How the matchmaking happens. Accessed 7.10.2019 https://neoteric.eu/blog/tinder-algorithms-how-the-matchmaking-happens/

Nick, B. (2014) Superintelligence: Paths, Dangers, Strategies. (p.22-51) Accessed 17.9.2019

O'REILLY, Mike, B. Jon, B. (2016) What Are Conversational Bots? Accessed 7.10.2019 https://www.oreilly.com/library/view/what-are-conversational/9781492042808/

Oracle, The Definition of Big Data. Accessed 30.9.2019 https://www.oracle.com/big-data/guide/what-is-big-data.html

Quora, Max, T. (2019) Accessed 18.9.2019 https://www.quora.com/What-are-some-cool-things-about-Lisp

SAS, (2019) Analytic insights, Machine learning. Accessed 9.9.2019

https://www.sas.com/en_ae/insights/analytics/machine-learning.html

Science Direct, Peter, W. (2014) Unsupervised learning. Accessed 11.9.2019

https://www.sciencedirect.com/topics/computer-science/unsupervised-learning

Search Data Management, Margaret, R. (2019) Data analytics. Accessed 12.4.2019 https://searchdatamanagement.techtarget.com/definition/data-analytics

Seikku, E. (2018) What on Earth is AI, Machine Learning and Predictive Analytics? Accessed 12.4.2019 https://www.tivi.fi/kumppaniblogit/hewlett_packard_enterprise/mika-ihmeen-tekoaly-koneoppiminen-ja-ennakoiva-analytiikka/2a0453e1-43d6-3532-b206-274741d258ea

The Verge, Nick, S. (2019) Amazon says fully automated shipping warehouses are at least a decade away. Accessed 7.10.2019 https://www.theverge.com/2019/5/1/18526092/amazon-warehouse-robotics-automation-ai-10-years-away

Towards Data Science, Shaan, R. 2018 History of AI. Accessed 5.4.2019 https://towardsdatascience.com/history-of-ai-484a86fc16ef

Towards Data Science, Elliott, S. (2018) Unsupervised Machine Learning. Accessed 9.9.2019

https://towardsdatascience.com/unsupervised-machine-learning-9329c97d6d9f

Towards Data Science, Ryan, W. (2018) Reinforcement Learning: An Introduction to the Concepts, Applications and Code. Accessed 22.9.2019 https://towardsdatascience.com/reinforcement-learning-an-introduction-to-the-concepts-applications-and-code-ced6fbfd882d

YouTube, Amazon News. (2019) AI Keeps Amazon warehouses humming. Accessed 7.10.2019 https://www.youtube.com/watch?v=B2Humr181Qw