Università
Ca'Foscari
Venezia

Master's Degree
in Management

Final Thesis

Review of AI-based marketing approaches

Supervisor
Ch. Prof. Giovanni Fasano

Assistant supervisor
Ch. Prof. Giacomo Di Tollo

Graduand
Giovanni Quarantotto Vittori
Matriculation number
873216

Academic Year
2019 / 2020

THIS PAGE IS INTENTIONALLY LEFT BLANK

“Any sufficiently advanced technology is indistinguishable from magic”

(Arthur C. Clarke, 1973)

Abstract

Artificial intelligence is the shiny new frontier in marketing. The increasing computational capabilities of computers and algorithms, and the increasing availability of huge volumes of data, may pave the way to new and unexplored venues in the whole marketing sector. Every single piece of the customer journey is going to change, and the process has already started.

ACKNOWLEDGEMENTS

First and foremost, I want to thank the incredible people that I have encountered during my interviews: Ester Liquori and Lorenzo Viscanti. Thank you for taking the time to let me conduct interviews, part of this work would not be possible without your help.

Secondly I want to thank my dissertation professor, Giovanni Fasano, for taking me under his tutelage and accepting my topic.

Lastly, I want to thank my mother who always wanted to graduate but never had the opportunity and my sister for helping me focus on the work when I did not feel like it.

Table of Contents

Chapter 1 – Introduction

Human beings hate making mistakes. That is the reason why every single customer journey begins by seeking information about the product that we want to purchase. We want to reduce our perceived risk and we try to accomplish this by learning more information about a certain item, so that we can be sure that our money is not going to be wasted in some poor decision. In doing so there are several aspects that may limit, or completely remove, our critical thinking: advertisements, friends' recommendations, or personal preferences. But what if we could, objectively, decide which product is the best for us? What if we could have a personal assistant that can help us find the perfect gift for mothers' day, one impeccable pair of shoes and an excellent car? This is not a job for a human. As a species we have countless biases and, most importantly, we are motivated by monetary or emotional rewards. Even if we could try to find the perfect item for a relative, a selfless objective good deed would be overtaken by the pleasure that we would gather by donating that perfect gift.

Finding a machine capable of being our very personal shopping assistant, which does not betray us for money, or works with companies interests in mind, is quite the paradox at this time. We are talking about a level of machine superintelligence which will hardly be reached for a couple of decades at least, but companies are nonetheless trying to make this possible in some other, very limited, areas. Granted they are working with their interests in mind but, is still remarkable how much a couple of codes lines can change the shopping experience.

In order to answer these questions, we need to dig deeper into the very definition of AI itself: learning what Artificial Intelligence means, which are its core components and what is an algorithm. Once the foundation is laid down, we may analyze the current state of the art of AI based marketing applications, what are the costs in terms of money and people and what the future may hold for these applications and for human labor.

Definitions and working principles behind Artificial Intelligence

Before defining what the state of the art is, in terms of Artificial Intelligence development in the field of marketing, few words are necessary to explain the working principles behind AI.

The problem starts with the definition of artificial intelligence. As Jerry Kaplan enounces in his book (Kaplan, 2016), there are many different explanations of AI, but most are roughly aligned around the concept of creating computer programs or machines capable of a behavior we would regard as intelligent if exhibited by humans.

This approach of defining machine intelligence is deeply flawed. First, it is difficult to define human intelligence since our love for simply visualized data guides us to numeric systems like Intelligence Quotient points; second this is not enough to measure comprehensively the level of intelligence of one machine with respect to a human being. Nonetheless, it is meaningful to say that one person is smarter than another, at least within many contexts. And there are certain markers of intelligence that are widely accepted and highly correlated with other indicators. For instance, how quickly and accurately students can add and subtract lists of numbers is extensively used as a measure of logical and quantitative abilities, not to mention attention to detail. But does it make any sense to apply this standard to a machine? A $1 calculator will beat any human being at this task hands down, even without hands.

Further complicating this matter, is the difference between how to solve a problem and whether you solve it at all. Consider a program that plays the simple game of tic tac toe. The only function that this program needs to carry out, in order to win the match against its human opponent, is to calculate the best combination of moves, among the 255.168 possible options that this game offers. Most people would not accept this kind of program as "intelligent". Keeping this in mind we can imagine a different approach. A program that learns how to play the game just by observing a human play. In this scenario, is not just a matter of finding the right combination of moves among hundreds of thousands of possibilities, but it is about learning strategies that are successful, like blocking the opponent once she does two in a row, or the fact that occupying the corners and leaving an empty square between them leads frequently to a win.

The last difference that needs to have our attention is the approach at making mistakes. Humans and smart machines alike learn by making mistakes. The difference lies in the response that humans have for the two. If the little Francesca is learning how to speak, it is just natural that she will make grammar and syntax mistakes and we will not get angry or frustrated with her because she is just learning something new. Instead when the Google Voice assistant is not capable of sounding "human" enough, we ridicule it and dismiss it with little empathy and consideration.

One of the most accurate and simple definition of artificial intelligence is given by Kaplan in his book (Kaplan, 2016): "*The essence of AI - indeed, the essence of intelligence - is the ability to make appropriate generalizations in a timely fashion based on limited data. The broader the domain of application, the quicker conclusions are drawn*". In just a few words the author is capable of defining the real use definition for artificial intelligence: drawing conclusions from unorganized data. Moving deeper into the definition itself, another description, given by Haenlein (Haenlein, 2018) starts to show the difficulty to define AI: "*…it is surprisingly difficult to define what AI is and what it is not. Or, to put it differently, there are about as many different definitions of AI as there are ways to describe Snow White's beauty…*". This later approach, can be also found in other books (Liquori, 2020). Here the author specifically stresses how the term is often abused and rendered empty of its meaning.

Society always fantasize about artificial intelligence as a sentient being that will replace every single job that we carry out and leave us with nothing, lot of this has to do with reality tv shows and sci-fi literature. In reality AI is a tool that helps us in our daily tasks. The more difficult the task, the greater will be the appreciation of having a swiss army like tool, which is capable of assisting us in many different ways in a single, compact, form factor.

Core components of AI

The term Artificial Intelligence can be viewed as an umbrella (Sterne, 2017), a broad field of study that encompasses different fields, theories, and technologies. Although it is beyond the scope of this work to define in detail all the different components of AI, a brief focus on the most important ones, with respect to the marketing application, must be outlined. For the sake of comprehension of future interaction among existing technologies in real world implementations, these are the core components of artificial intelligence.

Machine Learning. This is perhaps the first concept that comes to mind when dealing with artificial intelligence. This subset of AI has its roots in the idea of self-learning. As Oliver Theobald accurately describes in his book (Theobald, 2017), this refers to the application of statistical modeling to detect patterns and improve performance based on data and empirical information, all without direct programming commands. That being said, in order to function properly, machine learning, still requires a lot of human inputs: first the programmers need to outline two sets of data: the training data, which will be used to test the program; and the test data, which is more similar to real life information. Second the programmers need to select an appropriate algorithm, in order to reduce prediction errors. If the results with test data are satisfying, the developers can then deploy the program in a live setting.

Since **Neural Networks** and **Deep Learning** are so interconnected we will group them together in this paragraph. Neural networks, more precisely artificial neural networks, mimic the human brain through a series of algorithms (Kavlakoglu, 2020) and they are made up by at least four layers: input, weights, a bias, and an output. Deep learning instead describes a particular kind of artificial neural networks which is comprised by numerous layers, hence the depth in the name. The main difference with respect to machine learning stands in the way of how the algorithm learns. In the original machine learning algorithms, we need the human intervention to teach the machine by feeding it the structured test data set. Instead with deep learning a labeled data set is not strictly necessary, although we need a much larger data sample.

Lingering around the concept of human brain simulation, to be fair we are still talking about artificial intelligence, we find the field of **Computer Vision**. The scope of this technology is to create systems that are capable of recognizing images and videos like humans do. It is based on the popular hypothesis that image recognition in humans is carried out by pattern recognition, and this is exactly what we are trying to teach the machines. As an example, machines divide every single picture in different sets of pixels, each with a different color and brightness, then this data is given to the computer vision algorithm that will be responsible for the decision making.

The last core components that will be useful to better understand the real-world application of AI, is **Natural Language Processing**. The idea of talking to a computer is as old as the computer itself. Mostly popularized by movies like 2001 A Space Odyssey, the real-world applications for this technology are far more disappointing with respect to sci-fi literature. Natural language processing is the parsing and semantic interpretation of text, allowing computers to learn, analyze, and understand human language (Sodha, 2019). This is perhaps the broadest of the disciplines that have been previously mentioned, since it draws inspiration from many different fields, including computer science and computational linguistics, to fill the gap between human and machine communication.

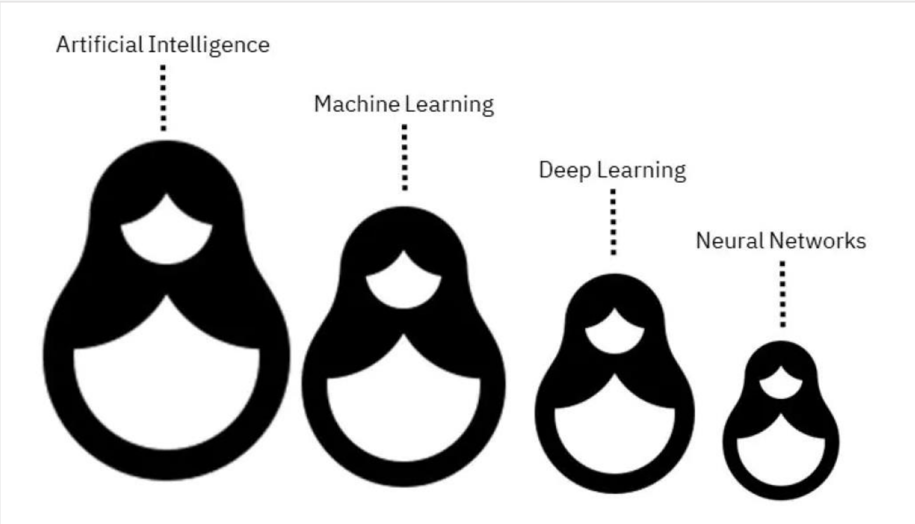

Figure 1. Kavlakoglu Eda AI vs. Machine Learning vs. Deep Learning vs. Neural Networks: What's the Difference? [Online] // IBM Cloud. - IBM, May 26, 2020. - October 3, 2020. - https://www.ibm.com/cloud/blog/ai-vs-machine-learning-vs-deep-learning-vs-neural-networks.

The supporting technological role, for the aforementioned core components, is carried out by a few worth-mentioning technologies:

- Graphical Processing Units. A lot of computational power is required to push these technologies, and the perfect candidates for this job are graphical processing chips, or graphics card in simpler terms.
- The Internet of Things. IOT helps generate incredible amounts of data, which will be then fed to algorithms.
- Advanced Algorithms. Trying to define a simple algorithm is fairly straightforward: a set of rules for the computer to follow in order to carry out tasks. Instead trying

to define an Advanced Algorithm proves to be a difficult task. This is because they are specifically designed for the situation that needs to be analyzed and they are combined in a way that enables them to analyze more data more efficiently (Li, 2017). Few examples of machine learning algorithms are: Supervised learning, Semi-supervised learning, Unsupervised learning, and Reinforcement learning.

Turing test

One basic benchmark to consider "intelligent" a machine is, and has been for the last decades, the Turing test. Designed in the mid-1900s, is still today, more than 60 years later, considered a valid methodology to test the capability of a machine. A lot of critics are forwarded by major subject matter experts, like Searle (Cole, 2020): we will go over a few of them in the next paragraph. Right now the focus should be pointed into the idea of the Turing test itself: what it wanted to accomplish and what it wanted to demonstrate.

Even though the birth date of Artificial intelligence officially dates back to 1956 (Coperland, 2003), when the John McCarthy's Dartmouth Summer Research Project on Artificial Intelligence took place, Alan Turing was thinking about machine intelligence as early as 1941.

Historically speaking we are at the very beginning of Artificial Intelligence. The technologies at that time were nowhere near as capable as the ones that we have at our disposal today. That being said, Alan was already envisioning a future where distinguishing between man and machine would be difficult. With a very optimistic attitude Turing, in the 1950s, believed that intelligent machines capable of sustaining a conversation in English where just around the corner (Cole, 2020). As we can see, even though we have some basic systems capable of interacting with humans, the line is still blurry at best, and attracts negative feedbacks by users most of the time. The key takeaway that we need to consider is the concept of learning. Much like infants, these systems are capable of a greater degree of autonomy and language understanding over time. This will assure that in the future we will have systems capable of carrying out conversations in the same way as we, humans, do.

Alan was very forward-looking and decided to develop a method that could grant us the ability to distinguish between human and machine. Not only this, but correctly pointed out in his paper (Moor, 1976), Moor proposes the following concept "*The proponents and critics of the imitation game have misunderstood its significance. The real value of the imitation game lies not in treating it as the basis for an operational definition but in considering it as a potential source of good inductive evidence for the hypothesis that machines think*". This is the concept that needs to receive attention: not to find a perfect test that is capable of finding the human, or machine, 100% of the time, with a straight and simple yes or no. The solution is to develop a test that is capable of suggesting, with a certain degree of confidence, if the answers come from a real human or a machine.

Described by the word of Turing itself (Coperland, 2003), the first iteration of the test consists in the following. A jury, who is not familiar with the concept and workarounds of machines, need to pose a series of question. The jury must not see the machine, otherwise the test would be far too easy. The intent of the system is to trick the jury into thinking that they are receiving answers from a real human. Tools to deceive the jury are given to the systems. Making spelling mistakes, or waiting a few seconds before answering the question, making seem like a real hesitation from a human being are permitted. This second tool is already a reality inside the google virtual assistant platform (Google Keynote I/O - 8th May 2018, 2018).

Weak versus Strong AI

Closely connected to the idea of machine intelligence, it is important to outline the difference between weak and strong AI.

Every artificial intelligence concept that is going to be described in this work, including the previous and future systems, belongs to the category of weak AI. **Weak AI**, also known as narrow AI, is a system capable of performing one single task at a time. For example, a machine learning algorithm is able to play chess or answer questions regarding your sales data, not both (IBM Cloud Education, 2020). At the opposite side of the spectrum we have the theoretical concept of **Strong AI**. These systems are capable of carrying out different tasks at the same time, much like humans can talk and walk simultaneously, without having the need of the human input.

The term "theoretical", was used on purpose. Up to this point in time, there are not any existing systems capable of carrying out strong AI capabilities. The extant research is very divisive on this topic: some subject matter experts believe that strong AI is the next evolutionary steps of the human race (Bostrom, 2014), the last problem that we, as a species, will ever need to solve, and it is right around the corner. Other academics, like John Searle (Cole, 2020), believe that strong artificial intelligence will never be achieved, and this is not a matter of time or technological advancements, it has to do with the fact that computers can simulate intelligence without actively being smart.

The practical concept of the Chinese Room Argument will help us understand this matter. Proposed by John Searle in 1980 (Cole, 2020), can be summarized in his quote: "Computation is defined purely formally or syntactically, whereas minds have actual mental or semantic contents, and we cannot get from syntactical to the semantic just by having the syntactical operations and nothing else… A system, me, for example, would not acquire an understanding of Chinese just by going through the steps of a computer program that simulated the behavior of a Chinese speaker". Digging deeper. The Chinese Room Argument experiment is carried out as follows: a person, who does not speak Chinese, is left in a room with a book covering the basics Chinese language rules. Another person, who is fluent in Chinese passes notes to the person inside the room. The person inside the room, with the help of the phrasebook, is able to translate the notes that she has been given but this does not mean that she is able to speak Chinese. It was just a simulation of understanding. This experiment was designed to show how flawed the Turing test is.

On the opposite view of the spectrum we can find the words of Nick Bostrom (Bostrom, 2014), which poses that strong Artificial Intelligence will replace humans in the dominating species in the planet, much in the same way as humans are in the dominant position with respect to gorillas, or any other animal species.

Even though we are not at the point of strong autonomous AI yet; the world is getting ready for the imminent impact that artificial intelligence will cause in the next few decades. More and more everyday tasks are carried out by virtual agents: AI bots accounted for 85% of customer service interaction this year according to Gartner (Tonner, 2019), the AI software spending reached two billion dollars in 2020 according to IDC research (Brown, 2017) and while governments are still trying to grasp the possible consequences of this phenomenon, the creation of agencies and custom built working groups is taking place in every country.

Data

We spent a lot of time describing algorithms and working applications, but there is still a missing component that powers everything that we outlined up to this point. The biggest asset and biggest roadblock for artificial intelligence: data.

In order for artificial intelligence to make sense, both practically and financially, a lot of data is required: most systems are not useful if they are not fed the correct amount of data. The tailoring process of personal assistants and learning algorithms requires an unprecedented level of consumer knowledge and although we will go deeper in later chapters about use and regulations of personal information, now it is time for us to dig into the last missing piece of artificial intelligence: big data.

Correctly pointed out in his book (Sterne, 2017), Sterne forwards a very interesting idea. The role of data is the same today, as it was coal a couple hundred years ago in the first industrial revolution. And much like in the past, the access to this kind of resource granted prosperity to European nations, in the same way today companies are in the need for more and more data. This is helpful for segmentation, recommendations, or simple reporting.

That being said, the quality, and quantity, of data necessary for machine learning or deep learning is very hard to achieve for a single company, especially for smaller ones. The diversity and the quantity of data, varies exponentially depending on the use case that needs to be achieved. A general rule of thumb is forwarded by Courville and Chui (Courville, 2016) (Michael Chui, 2018), where they propose an estimate: acceptable performance of supervised deep-learning algorithms, will be achieved with 5.000 labeled examples per category and will match human levels of performance with at least 10 million labeled examples. To better understand this let us use a more realistic scenario. Consider the example of an industrial machine carried out by Panchack and Sunblad (Panchak, 2018). "*...gathering only one variable about revolutions per minute of your machine is not going to be enough to tell you why a failure happened. However, if you add vibration, temperatures, and data about many conditions that contribute to machine failure, you can begin to build models and algorithms to predict failure. In addition, as more data is collected, you can create accuracy requirements, such as This algorithm will be able to predict this failure within one day's time, with 90% accuracy*".

This has created the opportunity for companies to do exactly this: find and prepare data. One small example of this is carried out by Bloomreach (bloomreach, 2020). Thanks to AI, this company is able to detect pattern in product browsing without actually needing the personal information of a prospective customer. For example if someone searched the web for a watch, and a couple of days later the same search is carried out, it is very likely that behind these searches there is the same person who is looking to buy a new watch.

Another interesting case is presented by the companies that are not in the business of data but end up acquiring a lot of them anyways. Under Armor (Sterne, 2017) is one example of this. When we consider their app usage, which is generally employed for monitoring physical activities like hearth rate, running pace, and time of the workouts, we consequently have access to a lot of data which can be extremely useful in the recommendation of the next purchase, i.e. a shoe or pair of shorts.

The key concept is the right amount of data, because there exists such thing as too much information. Feeding an algorithm with a bigger than needed dataset will cause overfitting (Michael Chui, 2018), making the "noise" of the data too impactful and resulting in a lack of accuracy in performance. On the opposite spectrum, having too little information is also detrimental: the model fails to capture relevant data because of the sample size was too limited.

The variety of data plays a pivotal role in these systems. Traditional excel files, caches and cookies information are not enough to feed these colossal datasets. More and more Artificial Intelligence systems require fresh and diverse data. Other sources for this kind of information are pictures, audio, and videos. Aided by the help Neural Networks and Computer Vision (Michael Chui, 2018), these complex information can be unpacked into more manageable subsets.

Additionally the data needs to be updated very frequently, some examples show that a monthly refresh of the data is pivotal for the correct functioning of the algorithms, this is especially the case for supply chain management.

We spoke to a great length about dimensions of dataset, quality, variety, and freshness, but there are a few problems that every firm needs to consider when dealing with data enabled artificial intelligence (Michael Chui, 2018).

We have already seen the first problem regarding the volume of data. The cost and difficulty to procure the right dataset is a very complicated task, even for big enterprises, especially when dealing with ad-hoc solutions, perfectly customized for niche needs.

Another important aspect is time. Most of the algorithms used today are based on supervised learning and the amount of time required to label the example dataset, which will be then fed to the algorithm, is very time consuming. Especially if we need to update the data on a monthly, or weekly, bases.

Lastly, the role of biases in the datasets can be detrimental when used improperly. This issue is more related to social disputes. When feeding a learning dataset to an algorithm, one needs to be careful about the involuntary biases introduced. Multiple example of improper use of facial recognition have been observed in the past and continue to make headlines in the newspapers (Julia Angwin, 2016).

The current and future implications of Artificial Intelligence are astonishing. Not only the pattern-finding machine learning algorithms are capable of discovering information where no human would have been able to find in the first place, but also we are experiencing artificial systems capable of working alongside humans and starting to decide the next best move based on petabytes of data.

The world is changing, and it might not be ready for the next evolution of machines. The decision process of consumers, the customer journey, is being disrupted. The traditional linearity of the process, and the touchpoints along the way, is being modified to accommodate a more machine oriented process. How often do we empty our on-line cart because something else is being recommended right at the last second? Consumers are changing their habits, the way in which they shop, the way in which they look at news and the way in which products are being recommended to them.

Business models are moving to a data centric approach. The new gold of the 21st century will be the access to vast databases of consumer information (William D. Eggers, 2013) (DEPARTMENT OF ECONOMIC AND SOCIAL AFFAIRS - UNITED NATIONS, 2007) (THE ECONOMIST, 2017). At the same time regulations need to take place as soon as possible. History and extant literature alike, points in the opposite direction in terms of fitness of government, but nonetheless progresses are made. From a privacy point of view, the level of consumer knowledge, required to execute such complex systems is not only unprecedented but also problematic.

The goal of this research is to outline the existing state of the art applications in the vast realm of Artificial Intelligence applied to marketing. Looking at costs, both in term of monetary expenses and time expenses, looking at difficulties, and then moving to what the real world case studies have to offer. Lastly, the final point that this work is going to address is the future. How many working positions will be filled up by machines? The answer is: we do not know yet. From an economical point of view, "hiring" a machine is always more beneficial than a human. Machines do not get sick, no wages need to be paid and no physical offices are required. What will happen when there will only be machines that are communicating? What will be the place of the humans in all of this? This research is not pointing towards the direction where AI is the only way of solving every possible problem that a firm may encounter, quite the opposite actually. One of the biggest problems that today's organizations face when dealing with Artificial Intelligence, is the lack of clear implementation strategy and the misunderstandings carried out with it.

Chapter 2 – State of the Art and Case Studies

Now that we are familiar with the core components and technologies of artificial intelligence, it is time for us to dig deeper into the fertile grounds of AI based marketing applications. Not only we are going to analyze state-of-the-art and successful real-world implementation of these technologies, but also, we are going to dig deeper into the consequences, whether they are economical, legal, practical, and even connected with human interactions.

Letting the machine do the work

The first question that we want to answer is the following: "why would I let a machine analyze my data and formulate a strategy for me?" The simple answer is time. The sheer volume of data that our companies gather and create on daily bases is physically impossible to analyze by a single person. That is why more companies are beginning to adopt AI based technologies. We can think about these systems as an additional employee, an additional pair of hands, that person in the office who has all the right answers for the boring questions like "where is that file from 3 years ago?" or "do you mind helping me sort these 500 pages mailing list?". Granted these examples are very naïve and can be sorted out without the help of AI, but what if one day your superior asks for a deep analysis of the past 5 years sales volumes and compare that to the marketing campaigns of said years? That is an incredible amount of tedious work which may not carry out concrete and reliable results. With the help of artificial intelligence, scouring trough the endless amount of data would be easier and much more cost-effective.

Let us consider another example: the art of recommending the next product. Until very recently this role of predicting the needs of customers was carried out by sales professionals, clerks, and retail employees. Your clientele was limited in numbers and after a while, you acquired knowledge about your customers: you would talk about her children or grandchildren, about her husband, her interests, and the list goes on. The point is, the level of intimacy and knowledge, required to serve her the best products, that you were able to gather from a long period of time and experience is no longer an option nowadays. Think about the millions of purchases that are carried out on Amazon, eBay, or Alibaba every single second. No human would be able to sit through all of these transactions and advice customers on their next purchase. That is when artificial intelligence comes in handy. With the help of machine learning, a system is able, after a couple of purchases, to point the customer in the right direction to the next product. Nowadays this feature seems really trivial, and it is taken for granted in all of our online experiences, but it carries a depth of information which is nothing but impressive.

Here I have intentionally omitted the construct "online shopping". In fact, this paradigm would not only limit the extent to which this technology is applied and adopted, but also insert a wrong image, that AI is solely responsible for these kinds of transactions, where in reality is much more.

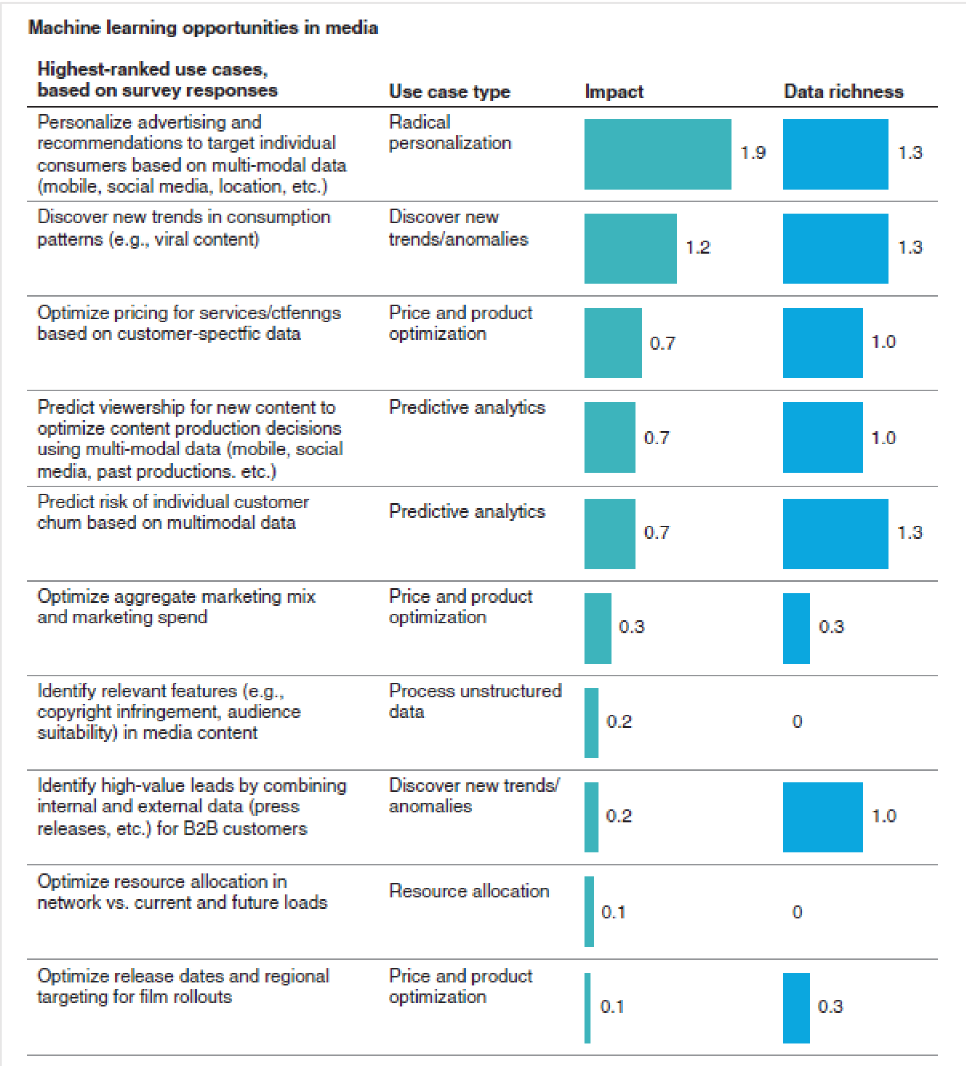

Figure 2. Sterne Jim Artificial Intelligence for Marketing - Practical Applications [Book]. - Hoboken, New Jersey : John Wiley & Sons, Inc., 2017. – page 8

To corroborate the last sentence, confirming that shopping recommendation is only a drop in the vast sea of artificial intelligence applications, we will start by analyzing few examples, the first one comes from one of the biggest video streaming players in the world: **Netflix.**

In this article by Medium (Yu, 2019), which outlines the uses that the tech giant found for artificial intelligence, the highlight is not only a system that is capable to serve a recommendation purpose, like the ones that we just described, but it is implemented in a way that makes pre-production, post-production, and market segmentation easier. Starting with the recommendation section, not only Netflix is able to suggest you your next movie or tv series based on what others like you enjoyed and your past viewing history, but it is even capable of auto - generating thumbnails. This application may go unnoticed by the average viewer, to be fair we do not really think about what is behind the picture that we see on the catalogue. Turns out there is quite a lot. They systems is based on the A/B testing methodology: as a starting point Netflix extrapolates thousands of video frames from its different movies. Then the AI, based on what other people with similar interests have watched, and your past watch history, ranks each image, in order to recommend the most appealing image. The idea behind this is that, if we like a certain actor/actress, then it is more likely that we are going to click on that specific thumbnail which includes that precise actor/actress.

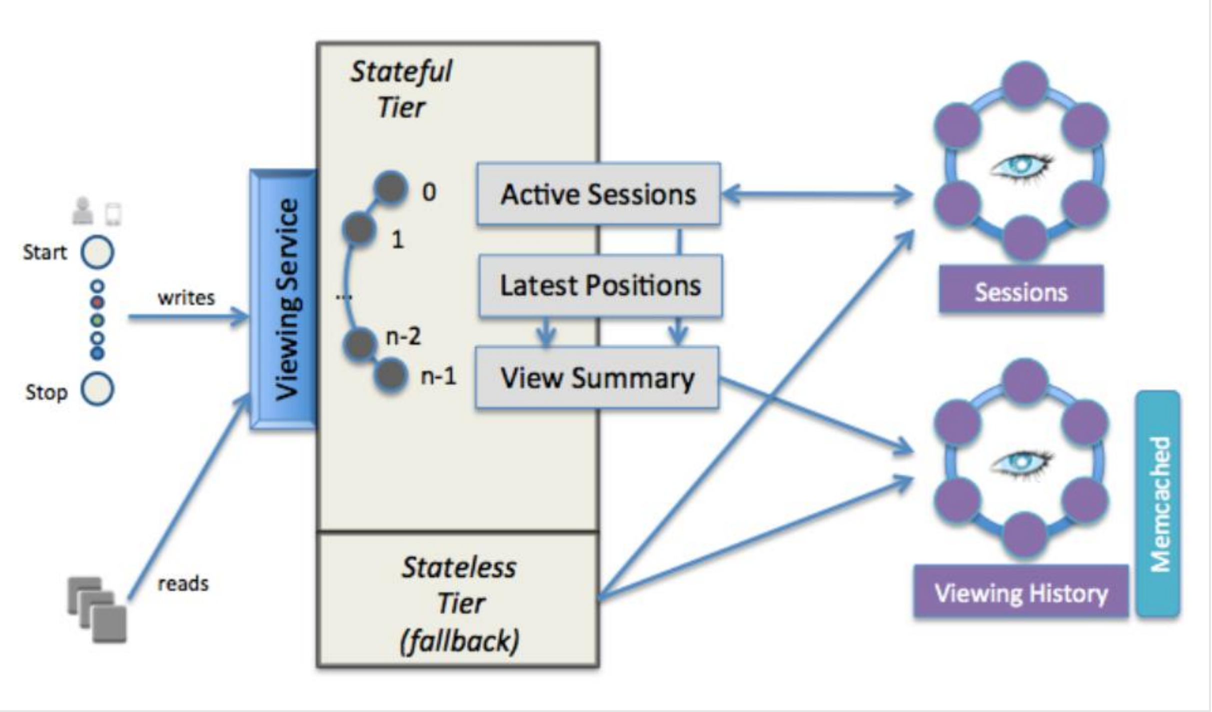

Figure 3. How Netflix uses AI and Data to conquer the world - LinkedIn - https://www.linkedin.com/pulse/how-netflix-uses-ai-data-conquer-world-mario-gavira/

Continuing on the streak of automation inside Netflix, artificial intelligence is starting to erode the very human role of location scouting. The California based company is starting to let data science administer the difficult job of scheduling, budgeting and production scene requirements (Yu, 2019). The important thing to keep in mind is that these tasks are not decided by the system alone. The algorithms are deployed to help the employees with these very difficult and complicated matters. This is perhaps one of the clearest applications of artificial intelligence and data science as it was meant to be: an additional pair of hands, a piece of software whose only job is the one to make your life easier.

It is very easy to waste resources and focus on the technology itself without aiming it at a concrete problem. This is true not only for artificial intelligence but for every real-world application of "futuristic" technologies. We do not need machine learning to provide personalization just because it is a cool piece of technology, we need to link it to a real business problem, in this case the watch time and engagement of viewers.

Our second example, digging deeper inside the concept of recommendation agents, come from **Spotify**.

Correctly pointed out in this deep analysis of Spotify recommendation algorithm (Boam, 2019), an interesting aspect emerges: the suggestion on what to listen to next is very similar to what you are already listening. This is not only what is perceived by the author but also what the reality is, data in hand. Information points to the fact that not only Spotify prefers to suggest something you are familiar with, in order to keep you listening, but also the user is discouraged to check out new and undiscovered artist due to the fact that automatically generated playlists are providing content performed by known artist, where the user would feel more at ease. At a first glance this problem seems fairly straightforward, and of minor importance, but this is not the case. When speaking with friends or relatives and talking about music, there is always the human recommendation: "I really liked that song, you should check it out". This type of social pressure is a powerful factor that invites you to try new things. This human element is missing completely from artificial intelligence and it is one of the reasons why when listening to Spotify, watching Netflix, or browsing YouTube it seems that we are always been fed the same kinds of content.

When comparing these two real world applications of machine learning, the winning scenario goes to Netflix. The reason for this decision goes beyond the implementation of the technology. Both of them are serving a purpose of recommending something new, but because Netflix is able to utilize AI in a more concrete way, provides a better example fro a real world case.

One last example that I want to stress and might help answer the question: "why should I let a machine do the work for me?" is the use of Artificial Intelligence by the logistics giant **UPS**. Their proprietary system is called Orion (Stevens, 2015), and consists of a 1.000-page long algorithm. The goal for this system is to help the company save money and time, two key components in the low margin world of logistics. With these two factors in mind, the senior engineering team decided to let the machine do the work. On an average day of work, a UPS driver can make up to 120 stops per day, for a single human it would be impossible to calculate the perfect route every day. Instead Orion is capable of producing a suggestion in 8 seconds, to help the driver decide the best route. The cost reductions, in these modifications provided by the AI may amount to as low as 1,99$ per ride, but counting the 55.000 daily routes in the United States alone, this number can easily scale up and produce a fairly substantial margin for the company.

But even this seemingly fantastic system has its drawbacks. The engineers at UPS worked on this problem for more than 10 years and in some instances they were very close at abandoning the project. Not only this but the estimated costs for the development of Orion are in the neighborhoods of hundreds of millions of dollars (Stevens, 2015), not a lot of companies can afford this amount of money.

Shortcomings aside, this is a great example of real-life case where artificial intelligence is designed to work closely with a human. A coworker with the incredible suggesting capabilities and, most importantly, is an example of a real business problem solution.

Ease of use

Sometimes the difficulty of applying AI not only stands in the mere technological application or finding people on board with the idea, but the system also needs to have a real problem to solve and not be adopted just for "the fashion" of technology.

This gently guides me into the next point, the ease of use or deployment difficulty. Artificial intelligence has an inherited difficulty of being adopted. The reasons for this are plentiful and in part were already seen, but the focus should be on two aspects: the technology behind it and the people.

We can start with the people and more specifically the gap between two very different professional figures: data scientists and marketing professionals. Sterne tries to bridge this gap in his book (Sterne, 2017), correctly outlining the idea that one does not need to be well prepared into the other field of work to make Artificial Intelligence work. Instead he points out how the correct cooperation could spring some fantastic real world applications.

These difficulties come from the fact that marketing professionals lack the technical backgrounds that data scientists possess, making it very difficult to implement an AI solution on their own. In the opposite way it is very difficult for data scientist to develop a system that produces the wanted results if they do not have any experience in the marketing field. This process of teamwork not only is very long, but also requires a company that has these two kind of professional figures inside, which is rare and precludes the vast majority of already existing enterprises.

Another thing that needs to be considered is the job security issue. Mostly present in the existing literature (Sterne, 2017) (Liquori, 2020), this issue aims at the psychological side of things. More often than not, people, in this case marketing professionals, are afraid of losing their job to a machine. The psychological background that this concern opens up is beyond the scope of this work, but it is worth pointing out that this fear is unfounded, at least for the foreseeable future. The idea that a machine will be capable of replacing a marketing professional entirely is not true, the same as it is not true for most of the existing professions. The balance stands in the degree of autonomy that we want to delegate. A lot of our everyday work consists in repetitive actions that are taken without even thinking about it, like sorting emails or categorizing customers, and this is where AI shines. A computer is capable of repeating their task endlessly, pattern recognition is literally child's play, and we would be more than happy to delegate these boring functions. The real star of the show, where humans do shine, is the decision-making part. We are able to see the bigger picture (Sterne, 2017), and infer the correct solution, basing our decision on facts and past experience. Past experience especially is something that needs time and work to be achieved, but data is not. Relegating the data-reporting part of the work, is going to be largely accepted regardless of your fear to lose the job.

This paragraph would not be complete if we did not mention the problem of taking decisions. As much as a machine is cold and repetitive, we, as species, are always at the mercy of our emotions. When taking decision we can be influenced by positive and negative externalities, not always the analytical/expert side of us is the one having the final word, a lot of the times we are influenced by hunger, angry, envy or politics.

The word technology is purposely maintained vague. This paragraph wants to dig a little deeper into the technology behind the implementation of AI: data infrastructure, data storage and labeling. Not only a correct understanding of the deployment of AI would require an engineering background but also the correct application is considered very difficult even for mature project engineers. Nonetheless some working principles can be extrapolated from the extant literature.

Costs

The third issue that I want to address regards costs. Both in terms of time and money artificial intelligence deployment is a paramount investment and should not be taken lightly.

Put in the simplest commercial terms: if a machine can do the job as bad as a human does, it will always be better than hiring an actual person (Sterne, 2017). It does not need any office, wage, pension, and it will not run to the nearest competitor when things get sour. Even more than this, narrow machine intelligence is already better than humans at routine tasks, which makes it objectively preferable for these kinds of jobs. Additionally, if we stick to the concept that AI is getting better, a software update would be a minimal investment with respect to hiring a new human professional.

Should we be afraid of machines taking our jobs? As a species this is not the first time that we pose ourselves this question. We have already experienced this fight against technology back in the 1800s where luddites were destroying looms which were, supposedly, taking their jobs. This rather pointless point of view was unfounded at best and totally wrong in the long run. Not only the looms granted a bigger production and increase revenues, but also required a lot of more skilled people to maintain and, consequently, to employ. Experts argue that we will experience a much similar situation with respect to Artificial Intelligence (Gartner, 2017). Starting in 2020 AI will be net positive job motivator, eliminating 1.8 million positions but, at the same time, creating 2.3 million (Gartner, 2017). Different industries and business functions are being adjusting at different paces: service operations, product development and marketing are the main business functions that attract AI investments, while telecom, high tech and financial services are the main affected industries (Michael Chui, 2018).

It is undoubtably clear that artificial intelligence brings values when deployed. Depending on the business function, McKinsey found that only 1% of the surveyed people who adopted AI reported insignificant or negative value added from its use (Michael Chui, 2018). Manufacturing showed the greatest value added while human resources provided the least. More specifically, in terms of physical adoption, the simplest sources that create business value are (Gartner, 2018):

- New revenues, increasing sales of existing products or helping create new ones
- Cost reduction, optimize production and delivery to allow for an increase in margins
- Customer experiences, creating niche solutions perfectly suited for every need thanks to personalized customizations

With all of this in mind, how much are firms spending on artificial intelligence? Gartner research describes an increasing amount of global spending in AI, pointing to a mainstream adoption in the 2020 to 2025 time frame (Gartner, 2016). Estimates are indicating a 29$ billion technology spending in 2021 (eMarketer, 2018) with investments not slowing down even during the times of the pandemic, where 66% of the firms which were adopting AI, increased, or did not change their artificial intelligence expenditures (Gartner, 2020). The majority of these investments is carried out by the big players of the technological landscapes: Google, Baidu and IBM are accounting for roughly 70% of the total investments in AI, with the remaining 30% taken by startups. Considering these amounts, and our previously mentioned UPS example (Stevens, 2015), it can be inferred that a single investment, which is able to pay its dividends in the long-term, is well beyond the possibilities of most of the firms.

Time

Now that we have talked about monetary costs and value we can move to the concept of time. Both in terms of investment and results, time expenditures and time savings are a big part of artificial intelligence.

Starting with the biggest apparent drawback, time invested in learning how to use AI and get the systems of the ground and fully operational requires a lot of resources. Not only you need to teach your workforce how to use AI, with the following period of getting used to, but also the time required for the system to learn how to behave is worth considering (Sterne, 2017). But we may start to consider this time tradeoff as a zero sum game, where the time spent implementing artificial intelligence, can be recuperated by the system advantages. Time saved in doing repetitive tasks, like running SQL queries all day, can be spent in more productive activities such as determining brand engagement over time (eMarketer, 2019).

One way to convince the more hesitating stakeholders, is to take a step-by-step approach (eMarketer, 2019). Raj Balasundaram, vice president of solutions at Emarsys, a customer engagement platform (Emarsys), focuses on this approach: he highlights to prospective customers the bottlenecks in their journey and then proceed to educate them on how much time they can save using their algorithms and machine learning.

We have previously seen how money and time can be, at first, perceived as the major barriers of entry when it comes to artificial intelligence. But this is not the case. Even though these are two major contributing factors in the development and implementation of AI, we have highlighted how abundant data needs to satisfy certain criterion and be present inside the organization. Very closely related to this is the concept of strategy: what is my strategy for implementing AI, what are the issues that artificial intelligence will help me overcome?

Time and money can solve a lot of problem in short-term, but if the company does not have a clear plan to implement, these resources are better placed elsewhere. This progress is different from something like moving from paper to digital (eMarketer, 2019): good AI strategies are built upon multidisciplinary internal teams. In the lamest terms: tech people must work with business people and vice versa (Sterne, 2017). A clear strategy will provide trust, help overcome the psychological barriers and pour founding from the right places (eMarketer, 2019).

State of the art applications

Before digging deeper into the real world case studies, the focus should be put on a couple of state of the art applications which will be beneficial in order to better understand future examples.

We have previously taken a broader approach when talking about AI, a sort of prelude to the concepts that will be outlined from here on out. These artificial intelligence applications, not always are implemented in a real world use, some of them are still proof of concept or mere ideas. That being said, the implementations of these possibilities are fascinating and worth mentioning.

The applications shall be divided into three main categories, moving from "simpler" systems to more advanced ones. The thematic aggregation of these systems focused on the concept of consumer, and its many facets.

Market segmentation

We can begin with the first broad application: market segmentation. Although the concept of slicing the market in tiny fragments is nothing new, applying AI makes things much more interesting. Nicely described in his article (Granville, 2015), Granville defines market segmentation as the following: "market segmentation, also called customer profiling, is a marketing strategy which involves dividing a broad target market into subsets of consumers, businesses, or countries that have, or are perceived to have, common needs, interests, and priorities, and then designing and implementing strategies to target them. Market segmentation strategies are generally used to identify and further define the target customers and provide supporting data for marketing plan elements such as positioning to achieve certain marketing plan objectives. Businesses may develop product differentiation strategies, or an undifferentiated approach, involving specific products or product lines depending on the specific demand and attributes of the target segment".

This highly challenging process is usually done manually by marketing professionals. Combining different spread sheets and other resources, making it most of the times an ad-hoc solution, often containing misleading false positives around anomaly detection as well as human errors (Brown, 2017).

Machines can do this job much faster. Comparing different segments between each other can be done at a fraction of the cost and without involving human help. Then once these data is available, we can compare metrics such as percentage adoption of our products and how different clusters of consumer are comparing to each other.

However these capabilities are not limited to existing segments (Sterne, 2017). Systems can be able to divide existing clusters to new data sets with different characteristics: untapped potential which might be difficult to spot, for a human operator, in the myriads of data. For example, what if the machine learning algorithm finds a sub-segment inside a region, which performs, in terms of sales, much worse than the other ones? This issue opens up a debate for a special customized product or a different pricing strategy. Furthermore what if could find more of that 1% customer, which is so loyal to our product that it does not even consider other brands? Adobe's Target seems to point in that direction (Brighton, 2016). Adobe's Target not only aggregates consumers using existing sources, both online and offline, but it also determines which variables are the most predictive of conversion, eliminating redundant clutter.

The last leap forward that I want to address is something dear to a good chunk of marketers: the concept of Total Addressable Markets (TAM). Up to this point the costs and manual work usually done to carry out this task was prohibitive for most companies. Even for the ones that had the resources to do so, it was more of an exercise about slicing up predefined segments already existing in some already present clusters. On top of this the digital tools apt to this kind of task are inadequate and lack the needed accuracy (Brown, 2017). Machine learning can help define statistically valid TAM of digital brand engagement at any specific point in time thanks to the continual real-time feedback of the marketplace.

Artificial intelligence has the potential to reinvent the art of market segmentation and democratize its use, which could make it the benchmark, and model, for future applications.

Physical application: out of home advertising to raise awareness and attract customers

The previous, and future, paragraphs are and will be more focused on the business analytics, rather than how AI can help us understand data. Let us now consider a more physical approach: the out of home advertising. Albeit, compared to other business functions, being the lesser impacted area, for now at least, it is still interesting to see what machines can bring to the table.

By this I mean the use of different areas of artificial intelligence in a way that helps raise awareness and attracts customers in the real world, not behind a computer. With the help of computer vision and deep learning we can expect store fronts, as well as brick and mortar store in general, to drastically change in the next few decades.

Let us begin by saying that the number of channels in which consumers can access the product has increased exponentially in the past decades, so are the ways in which information can be accessed. The previously well-structured customer journey is no longer straight, information gathering, and the actual purchase can happen at different points in time (Eric T. Bradlow, 2017). Firms have picked up on this trend and are starting to create hybrid models where brick and mortar stores are acting more and more as showrooms and then the actual purchases are done online, even Amazon is going offline (eMarketer, 2018).

The key to this transition is the increasing amount of available data. Not only "classical" data needed to operate loyalty programs or coupons, which is outdated most of the time (Eric T. Bradlow, 2017), but also geofencing mechanisms and visual recognition that are able to modify advertisement in the store fronts to attract prospective customers inside the physical building. Keeping this in mind we can start working on two fronts: better data and consequent increase of AI use.

Common knowledge points into the direction that having available past data is important to understand our year over year position. This is true only to a certain degree (Eric T. Bradlow, 2017). Having a lot of data does not mean having good data. If our past records are not relevant, then they will not be helpful when taking decision regarding pricing and products rotations. What do we mean by relevant? Relevant data is that information which can be linked together, spanning across different resources, to a single customer, making a complete picture of her past purchases, web presence, loyalty data and location data. Having these kinds of information enables us to work with in a different manner, not only we are able to see our year over improvement, or worsening, but also prepare tailor-made promotion for our clientele. For instance, a customer past purchase history, promotional response history and click-stream history can come together for ad-hoc micro segmentation, with dynamic pricing and personalized promotions grating a level of unprecedented customization (Eric T. Bradlow, 2017).

Keeping in mind the concept of ad-hoc personalization as the final goal, we can start working on the other notion previously mentioned: the increased use of AI and related technologies. Tools such as magic mirrors, have been in place for quite some time now: Neiman Marcus, an American chain of luxury department stores (Neiman Marcus), was one of the early adopters of the magic mirrors since 2015 (eMarketer, 2018), Saks Fifth Avenue has been using facial recognition since 2016 to identify VIPs and shoplifters (eMarketer, 2019). Moving even more to the outside world it is worth mentioning the real life case of the company Outernets. This firm works with retail brands like McDonald's to create personalized window displays (Outernets). In the offline world there is a profound difference with respect to personalization, where it is not a demand anymore, it is something expected (eMarketer, 2018). With the bulk of shopping still happening offline, brick and mortar store do not need to fear the big online retailer, but instead they should focus on creating a compelling offer which makes them valuable and at the same time differentiate them from ecommerce.

We can think about the store front as the home page of a website: where we need to understand what the customer is looking for and how to increase their basket size. Machine learning and computer vision are our best allies in this task: leveraging different sets of data, they can make the displays in the store front change depending on who is watching it.

Another problem that can be mitigated is the measurements of the campaign: the Return on Investment for example. Previously the out-of-home advertising was one of the least appealing medium at the disposal of marketers, due to the fact that no one had any idea if it worked. Making the core of the strategy brand awareness instead of actual sales. With AI there will be a shift allowing marketing professionals to actually see the sales resulting from their ads (eMarketer, 2018).

As magical as all of this may seem, this technology will never take off if companies are misleading regarding the use of the personal data. Customers will have to face an unprecedented decision: either to receive personalized offers by giving up a lot of their personal data or maintain a certain level of privacy and forfeit potential deals. Surveys are pointing in the former direction (eMarketer, 2018) with around 50% of the surveyed population being in favor of a better personalization at the expense of personal data.

Recommendation systems & virtual assistants

The decision to group together these two examples comes from the their deployment and what they can offer. We can think of the recommendation systems as a starting point for virtual assistant, in this way it is helpful to look at the former first.

When speaking about recommender systems, we are not at the point of fully autonomous entities, which are able to help us regardless of the task. That being said, the implications that arise from the theoretical application are important for both professional and consumers alike.

The core idea that shines through the extant literature is that Artificial Intelligence needs to cover a support role inside companies and everyday life: an additional hand which is present to make your life easier, leaving the decision making process to the human, which is exactly what is happening nowadays. Corroborating this idea are the survey results provided by Gartner (Gartner, 2018), where 58% of the surveyed people are convinced that AI will help them save time by automating tasks, and 53% believe that AI will help them save money, thanks to better deals discovered by automated systems.

Let us start with the **recommendation systems**. As previously mentioned throughout this work, these kinds of systems are already in place and work very well most of the times. From what to look next on Netflix or what is the next best podcast for me provided by Spotify, we experience AI without even noticing. In order to better understand the idea behind these processes, we are going to dig deeper into the **Amazon** recommendation system. Although their proprietary algorithm is secret, and much too complicated for this research, Amazon patent "application Personalized recommendations of items represented within a database" (Linden, 2001) can provide fascinating insights about the ins and outs of this system. Obviously, different recommendation systems work differently, we have previously seen Netflix and Spotify (Boam, 2019) (Yu, 2019), but analyzing the "secret sauce" that power one of the biggest ecommerce platform in the world will be a good snapshot of the industry.

McKinsey reports that 35% of all of Amazon transactions come from this system, as well as 75% of what we watch on Netflix (Ian MacKenzie, 2013). Again, as previously mentioned countless times, the first ingredient that powers this whole system is data. More concretely Amazon recommendation system is based on a series of information scattered around the shopping experience: past purchases, product placed in the shopping cart, items that has been rated and what other customers look at (Gavira, 2018). Consequently two main approaches have emerged in Amazon, and in the ecommerce in general:

- Collaborative filtering, where the same products are suggested to consumers with past similar transaction history.
- Content based filtering, focusing on the products and not the customer, suggesting items similar to what have been previously seen

To determine how relevant the products are for the customer, the algorithm looks at customer ratings for each product and filters out items that have already been bought by the user (Gavira, 2018). All of this is powered by Amazon proprietary Neural Networks, which process millions of products and customer in real time.

Albeit being powerful, this recommendation system faces an important difficulty: the cold start problem (Blerina Lika, 2014). This issue refers to the situation in which the systems have troubles getting started due to the lack of information availability. This can happen both for new customers and new products: new customers do not have a history of purchases with which the algorithm can work, and new products do not have reviews which makes it difficult to be recommended. This dilemma can be partially solved by turning to third party websites, where, hopefully, the customer has a browsing history, and it can provide preferences regarding ads and past purchases.

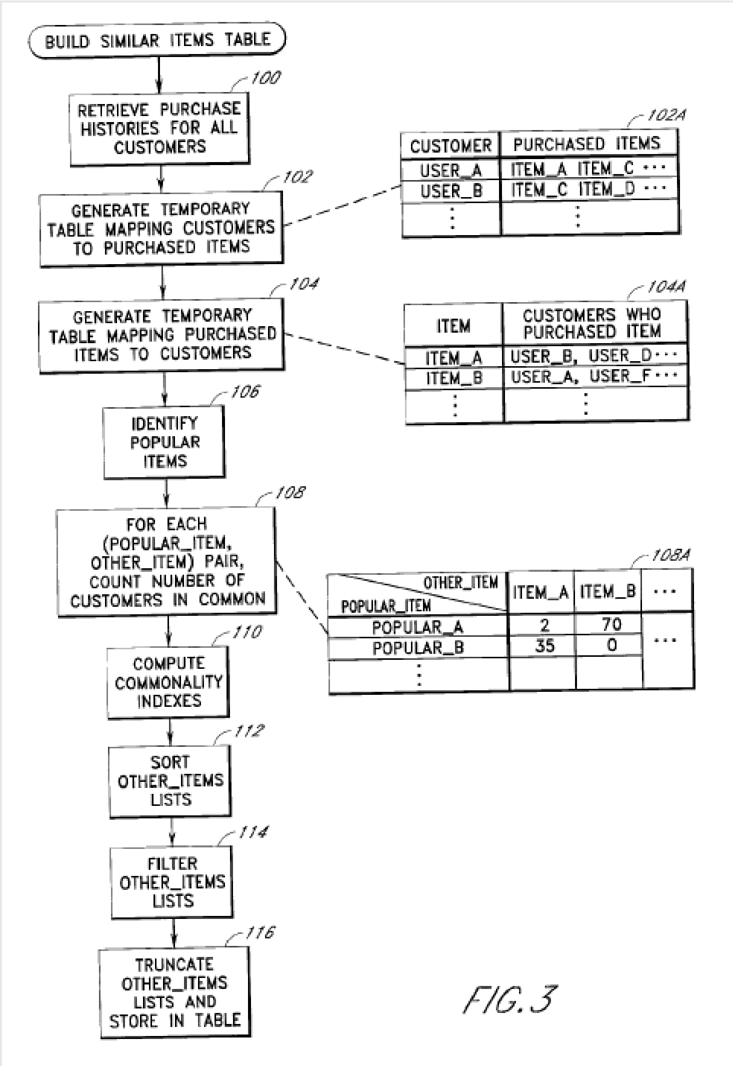

Figure 4. Linden Jennifer Jacobi Eric Benson Gregory Personalized recommendations of items represented within a database [Patent]. - United States, May 7, 2001.- page 4

The extant literature points to a single direction regarding the future of recommendation systems: the ability to offer new products, or services, before the prospective customer is considering looking for something. Perfectly summarized by Bren Smith "*We're convinced the future of recommendations will further build on intelligent computer algorithms leveraging collective human intelligence. The future will continue to be computers helping people help other people*" (Brent Smith, 2017).

Let us continue on the trend of future innovations by looking at **virtual assistants**. Building a functional service platform that is not just a flashy showcase of AI is very challenging (eMarketer, 2019), nonetheless progresses both in consumer and enterprise grade systems are being carried out, with Gartner predicting that 25% of digital workers will be using a Virtual Employee Assistant (VEA) by 2021 (Omale, 2019).

We can begin our analysis with the consumer centric systems first. Right now the most common example of assistant is the chat bot: interacting with customer and answering simple questions but still being very specific to fixed pages and predetermined situations with the most advanced examples being able to carry out natural language processing both as input and outputs.

That being said, creating a concierge-like system is a completely different task. To better understand this difficult process, it is worth mentioning the company Flybits which focuses on creating personalized recommendations in the banking sector (Flybits). The reason why this company focused its attention in the financial sector is straightforward: access to data (eMarketer, 2019). In this business area organizations have a lot of valuable data that can be used to provide personalized solutions to customers into areas like retail banking, asset management and insurance. The primary example of concierge-like application is being carried out by TD Bank (eMarketer, 2019), where Flybits provided the software development kit and then the bank implemented it in their proprietary applications. In this way existing customers did not have to download another app, and the advices portion of the program is perfectly integrated. The biggest issue, aside from developing the system, was imposed by existing regulations, which are numerous in every field, and especially strict in the banking sector. Different countries have different regulation and developing a single solution which is able to fit across jurisdictions poses a paramount challenge.

We can now move to enterprise grade systems. Even though at first glance there seems to be not much difference between these two categories, the real world application are quite dissimilar, even if they are overlapping most of the times. Due to the research purposes, trying to find the best and newest application of artificial intelligence, in the enterprise area, the role of chat-bots will be omitted for two reasons. Number one: we have already seen it in the portion of this work related to customers, number two: even though their deployment proved to be positively received by consumers, enterprises, and academia, they offer fewer and fewer starting points for cutting edge innovations and application. The most that we can expect in the following years are interfaces getting much more user-friendly and human-like, while providing the same services.

Furthermore, one important concept that I want to highlight, before looking at firm use cases, is how consumer grade systems might be good enough for enterprise use but not vice versa. The theoretical principles behind a recommendation engine can be translated from one audience to another, but a pattern-behavior recognition software, would not be of much use by the average user. This is exactly what the firm Pointillist is trying to achieve (Pointillist Home Page). Their marketing platform claims to automatically discover patterns of behavior across multiple touchpoints and predict what the consumer will do next (Sterne, 2017). On the same note, Thunderhead's ONE is able to understand customer intent and personalize further actions, offering the best response at any touchpoint (Thunderhead Home Page).

Building upon this we are now going to look at two examples of virtual assistants in the enterprise world. Moving to progressively more complex systems, is going to help us get a clear view of what is possible at the moment, and what may happen in the near future. We are now going to look at Salesforces' Einstein and Equals3 Lucy (Lucy Website - About Us) (Salesfoce Einstein Overview - Website). Both of these systems are very similar, but it can be argued that Lucy is the most advanced out of the two. The reason for the inclusion of these two services is their completeness. All of the previously mentioned technologies such as machine learning, predictions, and automation, are present here in a real use case scenario.

From the company's website: "***Salesforce*** *is a customer relationship management solution that brings companies and customers together. It's one integrated CRM platform that gives all your departments — including marketing, sales, commerce, and service — a single, shared view of every customer*" (Salesforce Website). With over 150 thousands customers served, this CRM platform is using AI for four main reasons: discover insights, predict outcomes, recommend next steps, and automate workflows (Salesforce Einstein Overview).

Discover insights. This first point is about the data already present inside the company. Einstein can help to identify significant patterns, that were previously missed, and better understand your customer. All of this process is carried out by analyzing the touchpoints between company and customer, without needing the human input. The aim of this process is to understand which channels, messages, and contents, resonate better with your clientele.

Predict outcomes. This second step is about knowing what is going to happen in the future before it actually happens. From predicting which sales lead will convert to actual purchases, to how to write marketing emails that will actually be read, Einstein help the company while at the same time providing predictive scoring. For example, the predictive lead scoring system gives each lead a score representing the likelihood it will convert into an opportunity (Sterne, 2017).

Recommend next steps. We have already seen a plethora of recommender systems in the consumer grade space. Much in the same way, AI can help employees increase their efficiency and productivity by suggesting what to do next.

Automate workflows. One of the pivotal points of Artificial Intelligence is the automation. Not doing boring tasks, such as login in customer data, including contract information, email threads, calendar events and social data, is the reason why many companies chose to deploy AI in the first place.

Now we can move our attention to the second real world case, **Lucy**. Based on IBM Watson technologies, Lucy is perhaps one of the most innovative virtual assistant/coworker. Defined as "AI-powered enterprise knowledge management system" (Lucy), Lucy is capable of utilizing most of the technologies that we have seen up to this point. Whether we are talking about natural language processing, natural language understanding, segmentation and planning, Lucy is capable of doing all of the above.

1. Listening. The key difference with respect to previously mentioned Einstein is the natural language processing and understanding capabilities. The system is able to understand regular questions, much like "Which media types are my competitors using?" or "What is the buyer profile for my service?" (Sterne, 2017), effectively making it an answer engine rather than a search engine (Lucy).

2. Research. Once you input the question, for which you need an answer, the result page will provide snippets of information: instead of links, like the results page from Google, you will receive charts, graphs or paragraphs which are most likely to contain the answer to the question, ranked by likelihood (Sterne, 2017).

Buyer Profile	
Male	83.9%
Female	16.1%
Income under $50,000	5.7%
Income $50,000–$99,999	17.2%
Income over $100,000	77.3%
18–44 yrs. old	33.2%
45–64 yrs. old	50.6%
65+ yrs. old	16.2%

Figure 5. *Sterne Jim Artificial Intelligence for Marketing - Practical Applications [Book]. - Hoboken, New Jersey : John Wiley & Sons, Inc., 2017. - page 228*

3. Segmentation. Now that research has been done, we need to know to whom send our messages. Starting from our existing audience, Lucy identifies their common traits, along spectra of 52 Jungian archetype traits (Sterne, 2017) (Lucy), and then classifies them into different personas that fit the promotional message. This is extremely helpful when prospecting customers.

4. With our prospects ready we need to thinker on different campaigns and how the public will react to certain media with respect to others. Lucy can provide different scenarios made up by different combinations of media outlets and point to the more promising one.

5. Machine learning colleague. Much in the same way as consumer grade systems, constantly ask "how did we do?" "Was this result helpful?", Lucy is getting help by the rating of its responses. The more you communicate with her, and the more time you spend together, the better the system will perform. Lucy's website claims that after just a month of use we can start to see promising results (Lucy).

We have separately spoken about these systems and we can notice how similar they look and how similar their functionalities are. But this does not mean that they are closed. In fact Salesforce is using Watson on top of its very performing Einstein (Levine, 2017). The reason is that Einstein knows consumer data, but Watson knows the rest of the world (Sterne, 2017). This integration is fundamental for a deeper customization, tailoring services to an unprecedented degree.

Legal implications

We previously mentioned how AI needs data to work. To ensure smooth operations AI systems often need to ingest millions of data. That being said we need to cover one important aspect that affects our everyday life: the legal implications that arise from the use of these systems.

Generally speaking, regulations, regardless of the area in which they are being deployed, have always been slower than the adoption of new technologies. We often see how tragic events lead to a wakeup calls and consequent regulations of a certain market area which was previously unregulated. The recent scandal of Cambridge Analytica is the perfect example of how unregulated sectors will cause harm to citizens and it showed how companies may behave in absence of a solid legal framework (Chan, 2019).

The questions are the following: how do we regulate AI and access to our private data? How do we regulate something that keeps evolving so fast? To better understand what governments are trying to achieve, this paragraph will be sub-divided into three smaller sections: one covering Europe, one covering the United States and the last one covering China.

In the previous paragraphs we hinted on how National and Supranational bodies are getting ready for the imminent impact of AI, and one of the best examples to date is carried out by the European Union. Leading the world when dealing with privacy concern, the EU provided in the past year a comprehensive guide where each point of the recent GDPR directive was associated and analyzed with respect to AI (Sartor, 2020).

On the other side of the ocean there is one of the biggest player in the world with respect to AI , the United States. The most liberal way of doing business in the country finds fertile grounds here, where a general laissez faire is tolerated. The federal nature of the nation allows for different interpretation of privacy depending on the different States. That being said, California, the birthplace of many of the biggest players in the world in terms of AI, recently passed a law to contain the problem, the California Consumer Privacy Act (CCPA). This directive gives consumers more control over the personal information that business collects about them and increases the privacy rights for California consumers. The attention here needs to be pointed to the fact that this is a California specific regulation, the other federal states do not have similar acts in place, making it easy for companies to circumvent this law.

When it comes to privacy and data regulation China has a different view with respect to the rest of the world. Their social scoring systems and constant surveillance makes it, at the same time, the best and worst country to analyze.

European Union

Several aspects of the GDPR are interlacing with AI, and although all of them posit interesting concepts worth mentioning, for this review we are going to focus on 3 of them: personal data, profiling, and consent. The reason for this decision stands in the functioning of AI itself. The interest in machines taking the place of humans in trivial actions will need access to personal data, used for profiling of the individual for which consent was asked.

Personal data represents any information relating to an identified or identifiable natural person (Sartor, 2020), meaning that information that do not regards humans are excluded, for example data on natural phenomena. Here the GDPR does not cover only practical information (e.g. name, date of birth, etc.) which can be directly linked to the person, but also anonymized data. If these latter type of information can be translated, by technical developments, into personal data, they should be considered too as such and fall under the regulation.

The importance of considering personal data in the GDPR stands in the utilization of the data itself. For example, a large dataset containing medical records of patients can be beneficial for all of the participants, because it provides a statistical valid sample for better treatments and cures. That being said, if these anonymized data can be traced back to the single individual, it will need to be treated according to the GDPR provisions. This concept, known as re-identification of data, can be seen as a specific kind of inference of personal data (Sartor, 2020).

Furthermore AI is capable, through algorithmic learning, to infer personal data. The important fact issued here is if the new inferred data needs to be considered in the same way as personal data. There are different results coming from cases of the European Court of Justice (ECJ). There have been cases in which inferred data was considered personal data and other cases where this was not the case (Sartor, 2020). Ultimately the European Union as an entity does not care about the inferring process in itself, but if personal data is misused.

Closely related to the concept of **inference of new data**, we find the practice of profiling. Profiling is in principle prohibited, but there are ample exemptions such as contract law or consent (Sartor, 2020). The uncertainties arise when the data of an individual is used for automatic decision making and if this process should be explained to the individual. Data inferred through profiling should be considered personal data and data protection principles should apply (Sartor, 2020).

To better understand this problematic situation let us consider a simple example. A dataset tasked to find the likelihood of heart diseases of applicants for insurance purposes, will not only consider the base health records of individuals, but it will also count for their life habits such as eating and physical exercise. With this additional data the algorithmic model no longer contains personal data since it links any possible predictors to a corresponding target. Moreover, the correlation inside the model is not considered as personal data, but the starting set of information and the inferred outcome are. This is important because the inferred data can be challenged by the individual.

The concept of **consent** is continuously challenged by individuals and institutions alike. Individuals can perceive consent as abstract and difficult (Sartor, 2020) and firms believe that GDPR regulation purposely omits certain uses of data, like inferred data. Both these implications can be disputed by observing the GDPR itself. In fact this European regulation does not limit the use of data, and its relative consent. Instead its implementation can be adopted in ways that are both beneficial for the data use itself and not hinder the trust of the people (Calo, 2013).

The user-friendliness of how these issues are perceived by individuals plays a big role in the overall perception of the problems. Both scholars and institutions points to a more direct approach, a simple opt-out option which does not require convoluted links and agreements (Calo, 2013) (Sartor, 2020). The CCPA is well prepared on this matter: by requiring companies to include a single "do not sell my data" link on their websites, CCPA enables users to exclude the transfer of data to third parties (Sartor, 2020) (CCPA, 2018).

Moreover three key issues are concerning European legal scholars: specificity, granularity, and freedom.

The specificity of consent is key. If I, as a user, have agreed upon the treatment of my data for sales and promotions purposes, this does not mean that the organization that obtained my data can freely utilize it. They need to strictly use it in the confined space that is the one of sales and promotion, nothing else (Sartor, 2020).

Closely related to specificity is the concept of granularity. For instance, is a general consent to any kind of analytics and profiling sufficient to authorize the AI-based sending of targeted commercial or political advertising? (Sartor, 2020) (GDPR, 2016). Recital 43 of the GDPR states that consent is presumed to be not freely given, therefore users should not, in theory, be dependent on agreeing on profiling practices in order to access services.

Lastly, we look into the concept of freedom. Recital 42 and 43 of the GDPR infer that consent is not free to use under clear situation of imbalance, whether they appear when dealing with the private sector or institutions (GDPR, 2016). This means that in situations of imbalance, consent cannot be assumed, even more consent should not be valid when refusal of consent is linked to a detriment of the service (e.g. patients are told that in order to obtain a medical treatment they must consent that their medical data is used for purposes not needed for that treatment) (Sartor, 2020).

Much work still needs to be done when dealing with AI and personal data usage, but the GDPR provides a solid framework for governments and firms alike. We have seen how European law focuses on the differences between data learning sets and operative data sets. We have seen how pseudonymization plays a central role, especially in training sets and how transparency is difficult to achieve. Now we are going to move from the right-based approach of the GDPR to the risk-based approach that emerges in US law.

United States

Being a confederation of states, the US presents itself with a myriad of diverse regulations, much like Europe. That being said, we are going to focus solely on the Californian CCPA for two main reasons: it is the most comprehensive piece of law with respect to personal data use in the US and because it shares some similarities with the GPDR, so that comparisons will be easier to carry out.

Much in the same way as GDPR, the CCPA focuses on personal data rather than AI itself. That being said, since the two concepts are so closely related, the same concepts that we have seen for the GDPR should be applicable.

Similarly to how we decomposed the GDPR, we are going to focus on four main implications of how the CCPA affects everyday consumer: right to be informed, right to deletion, right to opt out and right of access.

Right to be informed. Much in the same way of the GDPR, this part of the Californian provision focuses on the user's right to know when their data are being used (Paka, 2020). Furthermore it requires people to be informed on how model recommendations work (CCPA, 2018). This posits an important issue which already emerged in the GDPR literature: there are some gray areas when dealing with AI. We have previously seen how it is difficult, even for data scientists, to explain the decision process of neural networks systems (Taulli, 2019). Trying to easily explain these issues to the average user might not be within reach of everyone.

Right to deletion. This part of the legislation is very similar to the European one, but it differs in one key aspect. Not only it provides the user with the right of deletion of her data, both personal and inferred, but it would also require companies and institutions to delete their models if they can lead back to the single user (Paka, 2020). This is a critical point. Companies need to spend a lot of time, not to mention money, in order to implement a sound system which could be erased in a matter of days.

Right to opt out. This guidance dictates that users have the option to request the removal of their data at any time (CCPA, 2018). The only exception to this provision is the primary business reason for the service itself. The biggest impact, in the day to day use that the average user might experience, is the loss of personalization (Taulli, 2019) (Paka, 2020).

Let us say that I have given consent to access my information to receive a better suggestion for my next watch. The removal of data of another user might lead to a different model and different final recommendations.

Right of access. All of the user data, as well as all of the recommendation inferred by the systems should be saved and provided to the user if she requests so (CCPA, 2018). This means that companies and institutions will need to save both present and past information.

The biggest issue that stems from the existing literature seems to be the impact of current business models of many service providers (Taulli, 2019) (Ghosh, 2018). Especially in areas like California, the birthplace of many tech giants, having the ability to sell users data to third party, powered hundreds of business models while at the same time keeping the services free for the final user. A lot of companies will need to change the way in which they do business and at the same time, try to find a clearer approach when dealing with data.

China

Unlike from the previous subsections, in this paragraph we are not going to decompose and analyze Chinese regulations for two reasons. First it is impossible for me to translate and comprehend the original documents provided by the Chinese government. This will open the work to criticism and translation errors. Secondly Chinese companies working in the EU need to adhere to European regulations. Therefore we are going to see how the GDPR impacts Chinese companies even on their national soil.

Expanding on the idea of the nationality blind protection, the GDPR tries to extend its high level or personal data security outside the EU (Chen, 2019). Even though the best intentions are being carried out by the European Union, the actual application of this provision is difficult and not always respected, due to conflicting legislations and the list of countries where data is treated fairly according to the EU (European Commission).

The criteria upon which Chinese firms will need to comply with the GDPR are very stringent and require very little, in terms of physical or digital, presence to be triggered. That being said, the concept of establishment, is purposely omitted by the GDPR (Chen, 2019), which means that there is no single description of what constitutes an established presence in the EU. Even though some characteristics have been laid down by existing cases, this omission opens up a possible grey area in which it is difficult to determine the presence of a Chinese company (Jelinek, 2018)

Keeping this in mind, the most affected entities are the digital firms such as Huawei, Alibaba, Xiaomi, and other business entities such as banks based in China which have subsidiaries in the European Union (Chen, 2019). Since all of these companies will process a lot of data in order to operate, they will need to be GDPR compliant, otherwise they may face administrative sanctions and litigations.

The concept of establishment triggers the application of GDPR but it is not the sole contributor. For example if a Chinese company does not have an entity in the EU, but it processes the data of European citizens (e.g. a marketplace website hosted on a Chinese server which offers its product in Europe), this will also count as worthy of GDPR application. The reason being that it relies on the fact that it processes personal data of European citizens.

These two different situations trigger different outcomes. For companies with established entities it is mandatory to be GDPR compliant, whether they are big or small. Instead for companies without established entities the situation is more difficult. The decision to be GDPR compliant may impose a financial burden that only big, structured enterprises can afford (Chen, 2019). Furthermore the actual ability of a European court of laws to impact decision in mainland China has been questioned multiple times, and this reluctance to implement foreign decisions may remain for a long time (Chen, 2019).

Case studies

To complete my work, I decided to conduct interviews with companies working with AI in the marketing sector. Though real world examples have already been used throughout this work to corroborate the literature, I felt that a more detailed approach required a series of personally developed case studies.

While searching for suitable candidates to investigate, I quickly realized how vast and disperse this world can be. Therefore I decided to limit the focus on the following set of criteria:

- The company has to provide a service related to marketing and/or advertising;
- The company has to provide a service related to AI that used one of the previously mentioned tools of AI (i.e. machine learning, deep learning and/or artificial neural networks, natural language processing, natural language understanding and computer vision);
- The head office needed to be located in Italy. Office branches outside Italy were not considered;
- The dimensions of the company were not considered (i.e. the number of employees did not matter, the amount of annual revenues did not matter);
- The categorization of the company was not considered (i.e. I did not considered in the analysis if the company was considered a start-up or an established entity).

The reasons why I decided to adopt these criteria are the following: limit the otherwise extremely wide sample size and demonstrate that Italian companies are as capable as their international counterparts of providing compelling and interesting services.

The necessity of limiting the sample size was made necessary by the sheer volume of companies operating in this sector. Even a simple Google research will produce millions of results and although most of these pages are empty shells, focusing on a niche of the market seemed to be the right step towards a more detailed work. The final crumb of help was provided by the Italian national agency AgID (Agenzia per l'Italia Digitale). Tasked with implementing the technological procedures of the public administration, the agency keeps track of numerous entities dealing with artificial intelligence within Italian borders

(Task Force IA - AgID website). Whether they are firms, startups or university programs, this database was my first starting point.

The second intent, i.e. to demonstrate that Italian companies are capable of providing competitive services, was both a personal decision and a consequence of the first intent. First of all, I am a firm believer that not enough credit is given to our fellow countrymen where credit is due. We always admire enterprises coming from the United States or China or Germany and very little is done to advertise our successful endeavors. This work is not the place to explain the reasons behind my beliefs, but it seemed a good criterion to include and highlight certain enterprises instead of others. Secondly, since I had limited myself to Italian companies, I felt the need to further prove their capabilities in machine intelligence.

Once these criteria were in place, I started to dig into the list provided by AgID and the research material gathered through the process of this work. I specifically looked for companies in the marketing sector providing AI services (i.e. software to be applied) or using machine intelligence (i.e. companies that replaced part of their marketing jobs with AI). I ended up sending fifteen requests of interviews to companies I considered worthy of an interview, either in the form of direct emails and contact forms present in companies' websites. Out of the fifteen requests, I received three replies, and among these three replies two companies agreed to be interviewed: GhostwriterAI and Mapendo. Although I would have preferred to receive a wider feedback, my case studies' sample size was determined by the poor response I received.

Due to the ongoing pandemia of COVID-19 which makes traveling impossible, all of the interviews were conducted either via phone call or internet application (i.e. Google meet, Skype, Microsoft teams etc.). All of the conversations were recorded and transcribed (APPENDIX B and APPENDIX C).

Since my resulting sample size was very limited, and a quantitative analysis was out of the question, I decided to opt for a semi-structured interview (APPENDIX A).

The first question of the interview aimed at understanding the overall structure of the company while easing the interviewee in, allowing him/her to talk about the company freely. The second question targeted the competitive space of the companies. My intent was to grasp a well-defined round of players which could then be included in the last portion of the work: the guidelines for companies. Questions three to six aimed at understanding the application of the technology itself: seeing if it yielded benefits, which where the difficulties, and how long it took to complete the process. The last question was intended to capture a jack of all trade index: a measurement upon which companies interested in pursuing AI could go to have a solid feedback for their actions.

Despite conducting different interviews, I tried to stay as close as possible to the original script in order to avoid biases and impairing the results.

The final goal of the questionnaire was to see if any overarching theme emerged. Obviously, this ideal goal was hindered by the lack of companies willing to take part in the interview.

Case study 1 – Ghostwriter AI

The first company that I have considered for interviewing was Ghostwriter AI (GhostwriterAI Homepage). The conversation was conducted over the phone and my interviewee was the CEO of the company Ester Liquori.

This Turin based project has been founded in 2014 within its parent company YouAreMyGuide. The services provided by this firm focus on utilizing AI to analyze text-based conversations. The goal is to increase efficiency, improve target marketing and call center response time.

Text analysis has always been the focus of the company. Previous iterations of the project focused on small areas of interest, for example conversation analysis about travel. Now that the algorithm is more mature, it can provide helpful insights regardless of the area in which it is set up.

Even though, during the interview, the interviewee was reluctant in giving me access to their case studies and see what the algorithm was capable of, snippets of information from the rest of the conversation helped to paint a clearer picture.

First of all, depending on where the AI is applied, different data is required. That being said, the most common are: analytics from social media, insight from campaigns and help desk interactions between the call center and the caller.

Secondly let us take a look at how much time is needed for the algorithm to produce results autonomously. Depending on the kind of services that the algorithm is tasked to provide, from deployment to self-produced insights, at least a month is required. That being said, the interviewee suggested a lengthier timeframe of 3 months. This to ensure that thy algorithm is exposed to a much wider size of real life interactions.

Lastly, the difficulties encountered while working are mostly related to the expectations and what these kind of service can offer. There is a lot of confusion when talking with potential customers: most people expect a plug-and-play solution that does not require any external work. Deployment is actually not that trivial. It requires a commitment both in terms of resources and time. Instead, the interviewee stated that people generally expect to achieve results immediately, whereas the timeframe can be very long.

This example of AI is the correct starting point. An application that makes sure that the tedious part of the job of the marketer, i.e. gathering data and analyzing it, is done by the machine, leaving the human employee free to make decisions. This is exactly what we have outlined in the first paragraphs of this work, a machine freeing time for the human.

Keeping this in mind, this practice lacks, in my opinion, a core strategy implementation. It is true that you can use this service to speed up the process and remove mind-numbing work, while at the same time improving efficiency and reducing costs. But this can be considered as merely an optimization process rather than an overhaul which justifies expensive investments in AI. If we compare this service to what Netflix offers, with AI guiding not only the recommendations but also being a core component of their strategy for location scouting and targeting, we can realize how ample areas of improvement can be seen.

Case study 2 – Mapendo

The second company I have considered for interviewing was Mapendo. Based in Bologna, this firm focuses on using algorithms to improve advertising's efficiency. My interview was conducted over Google Meet, and I had the opportunity to talk with the CEO of the company Lorenzo Viscanti. Founded in 2015, the focus of this business is on the mobile advertising and 80% of their revenues comes from the American market.

In terms of data, the company is proud to state that it does not need any proprietary data from the company, and it relies on customer data to feed its algorithms. That being said we did not get into the specifics on how this process exactly works.

There are a lot of factors that contribute to the changes of the time variable. There is a general time frame of around 3 months, after which the algorithm is fully operational, but it can vary. For example it was pointed out, during the interview, that in the months of October to December, there is an overall push in the advertising and marketing sector, and this benefits the whole industry and makes the process easier for their algorithms. In the opposite case, during the months of January till March, a slowdown can be seen in the industry and this does not help the learning process.

The contribution of external factors to the learning speed of AI was not present in my literature review. The only way I was able to gather this insight was possible thanks to the interview.

Lastly, let us consider the difficulties encountered when dealing with customers and prospective clients. The approach taken by Mapendo is different than GhostwriterAI. Where Ghostwriter tries to explain everything to the customer, how the algorithm works and what data it needs, Mapendo instead tries to leave out these details and only presents results. The reason for this decision is that most of their clientele belongs to marketing departments, which are not familiar to the core concepts and inner workings of AI. This division of labor can be encountered throughout the literature but one of the more prominent figure against this approach is Dr. Sterne (Sterne, 2017). Sterne firmly believes that in order to achieve the best results when dealing with AI applied to marketing, a constant communication needs to be carried out, and both the parties, data scientists and marketeers, need to understand what the other wants and how is done.

Drawing conclusions, guidelines for other companies

The last matter in the Mapendo case study showed how theory and practical application can be different. Albeit the sample size being small, we can still see traces of discrepancies between literature and practice in both cases and interesting starting points to build on.

Keeping this in mind, I wanted to extract the best practices, seen in either case, to help building a small framework, or guidelines, for companies interested in the application of AI. The structure is going to focus on the 4 main areas that emerged during the interviews: data, timeframe, expectations, and competitions. The reason why I decided to focus on these topics during the interviews was the following: I believed these are the most relevant factors in determining the overall situation of the companies and grasp a snapshot of the reality in which they are operating.

Data. An entire paragraph of this work was devoted to describing how data is pivotal for the correct application of AI. This was also corroborated by the interview with Ester Liquori, GhostwriterAI, which highlighted how it is necessary to expose the system to a good amount of real case studies in order to make it perform autonomously and correctly.

This approach has inherited difficulties both in terms of how the data needs to be organized (Panchak, 2018) (Michael Chui, 2018), and at a privacy level (Sartor, 2020). Another possibility emerged during the Mapendo interview: utilizing user data instead of company data. This opposite way of feeding data to algorithms, which did not emerge during my literature review, can be beneficial for companies and it respects the previously seen European and American regulations (GDPR, 2016) (CCPA, 2018). To summarize, we can state that depending on the service that you require, data may, or may not, be necessary.

Timeframe. This area is the most unpredictable among the four points that we are describing. From the decision to adopt AI, to the results provided by the system autonomously, different factors influence this timeframe:

- Decision-making process: this concept emerged in the Mapendo interview. Specifically, the interviewee stated how different is the decision process in Europe with respect to the United Sates. While in the Anglo-Saxon country the time-frame can last one or two days, from the first meeting to the beginning of the project. On the contrary in the Old Continent this process can take weeks.
- The internal structure of the company. More and more firms which are deciding to dip their toes in the waters of AI are small or medium entities. This means that they are not equipped with data science departments, so no specialized internal resources are present, and there is no clear problem in which AI could help. This lack of specialized human resources emerged during the interview with GhostwriterAI and addressing this missing assets takes time.
- Deployment to full operation. From the learning process to actual deployment we need to consider a couple of months before having a fully operative system. In both interviews the need emerged to let the algorithm learn for a prolonged period of time. Especially when talking with Ester Liquori (GhostwriterAI), this extra time seemed necessary in the application of AI to the analysis of call center interactions. This lengthy procedure should be applied to ensure that the algorithm gets exposed to all of the possible outcomes it may encounter when deployed. Having a bigger sample size can improve the results of the AI. During the interviews also emerged how external conditions may speed-up, or slow-down, this process. We

can say that a couple of months, from two to four, is a good benchmark against which future applications may be tested and autonomous results expected.

Expectations. This is perhaps the most fragile aspect when it comes to what AI can and cannot do. As perceived in the interview with Ester Liquori, GhostwriterAI, the confusion coming from this new technology is both an impediment and a myth to be debunked. An impediment because people might be discouraged by the level of complexity that this process may take. Furthermore we have seen in the interview with Lorenzo Viscanti (Mapendo) how they prefer to present only the results and omit, if possible, to explain how their algorithm works. A myth to be debunked because it rises doubts and makes people believe that the reality is more difficult, or easier, than what is really like. To help companies reduce this fear of adopting AI, my suggestion would be to start with a small project like developing an algorithm to categorize emails autonomously or segment the customer base. We need to consider AI as any other new activity that we are not familiar with, starting with something smaller will help speed up the process and let the company understand if AI is something that could be implemented at higher level or something beyond their current needs.

Competitions. The situation that emerged from the interviews was cloudy at best. Neither interviewee had a clear response when asked "*which is your closest competitor on either Italian or European scale?*". Some examples showed up, but they were referring to similar services provided by other companies not in the field of AI. For example Mapendo highlighted how the portion of their work, in particular the connection between companies seeking help in AI and firms providing these services, was done by an external agency M&C Saatchi. The reason for this lack in transparency could be attributed to a decision of omitting the answer during the interview, or the absence of a clear situation in the overall plethora of companies working with AI. This second point could be better understood recalling that this field is relatively young, and despite being under the spotlight, it is possibly not receiving the attention it deserves.

Chapter 3 – Conclusions

We started this work focusing on how the advent of AI could be beneficial for everyone, both for consumers and professionals alike. We talked about how an unbiased agent would help in our day-to-day tasks and decisions, alleviating us from repetitive and boring errands.

We have learned how AI algorithms are being developed and how groundbreaking discoveries are taking place very rapidly. We have learned how human and machine interactions are increasing exponentially and how sometimes it is difficult to distinguish between the two. We have seen how governments and legislations are trying to keep up with the extremely difficult job of regulating AI and its related instances, which are changing so rapidly, and how the correct firm application of AI can be extremely difficult.

Through our research we discovered that ample resources need to be devoted to the correct application of AI, and even thought the barriers to entry can be reduced, it is difficult to implement it at a core strategy level, even for big international entities. We have seen how it took UPS years to develop Orion (Stevens, 2015) and how the firm almost gave up even though they had an incredible amount of human and economical resources at their disposal.

We have analyzed how the first uses for AI are trivial and repetitive actions: looking through vast amounts of data and analyzing costs are the first examples that come to mind (Sterne, 2017). Then moving to more complex operations, we have seen use cases of companies devoting algorithms to guide their work routes (UPS) and implementing scouting strategies for their new tv-series (Netflix).

Finally, to better understand the role of AI inside companies and its current application, I decided to dig into the real world and interview Italian firms. This helped me better understand differences between theory and practice while providing interesting insights on what the real world had to offer.

During these interviews I had the opportunity to learn hidden aspects that influence the learning processes of algorithms. For example, in some specific real applications, the learning abilities of an algorithm depend on the month in which it is deployed: if it is deployed at the end of the year, where a big push in advertising, can be seen, this process is shorter than if it were deployed at the beginning of the year. Another interesting insight is how the different areas inside a country can play a role in the learning process: depending on where the AI is deployed the learning process shapes differently.

Another interesting aspect that emerged during the interviews was the difficulty in defining competitors, as well as the difficulties when dealing with clients and how different countries react differently to these new services. The difficulties when determining a competitive scene can be attributed to the fact that AI algorithms cover a huge variety of possibilities and it is difficult to pinpoint the exact set of competitors. The difficulties of dealing with clients, although being an inherited characteristics of working with people, is attributed both to the ignorance of companies and possibly the extreme difficulty of AI. It is for this reason that Mapendo prefers to show only results of the algorithm rather than explaining how it works. The reason for this is that “marketing departments understand KPIs better than anything else”.

Last to all, the reason why I have decided to embark on the topic of AI comes from my personal interest in technology as well as from the nature of my degree. Technology has always fascinated me, and I firmly believe that it needs to make the life of the human species easier, that is a tool to enhance our abilities and help us in a time of need: whether it is going through emails or comparing X-ray scans to find traces of cancer. The other reason lies in the title of my degree course: innovation and marketing. Finding an innovative topic seemed a perfect conclusion for my final dissertation.

After concluding this thesis I have a clearer idea on the reality of AI and how most of the times both its capabilities and terminologies are misinterpreted. I would describe the state of this technology as immature and uneven. **Immature** because even though a lot of material is available, it might be difficult to dig beyond the first basic concepts and apply the newly received knowledge. I believe the reason for this difficulties deals with how AI works: one thing is to talk about algorithms, another is trying to correctly explain how they are developed. That being said, I believe that more efforts need to be placed in this matter. It can also be considered **uneven** because of the degree of disparity in resources available to smaller companies and international corporations. Most of the case studies that I have analyzed are multi-billion dollar conglomerates with extremely deep pockets and access to international assets and labor. I feel that smaller players experience severe drawbacks when it comes to developing their own solutions. Nonetheless possibilities exist. The interviewees that I have talked to belong to relatively small companies which have found their niece and are trying to promote themselves as competitive services providers.

Finally, even though we can never be fully prepared for what the future may hold, the direction in which AI is heading is the one of full autonomy without any human input (Bostrom, 2014). The precise moment in which this will happen is not clear yet, but scientists indicate that we are at least a couple of decades away before hitting ultimate AI (Bostrom, 2014) (Sterne, 2017).

Appendix A – Interview questions

Questionnaire:

1. Potrebbe gentilmente fornirmi alcune informazioni sulla struttura dell'azienda e sui servizi che essa offre?
2. Potrebbe indicarmi quali sono i principali competitors della sua azienda a livello italiano e/o europeo?
3. Quali sono i requisiti minimi in termini di dati necessari al funzionamento del vostro algoritmo?
4. Avete dei particolari requisiti hardware perché il vostro algoritmo sia operativo?
5. Quali sono le maggiori difficoltà che i clienti della vostra azienda devono affrontare nell'adottare sistemi AI?
6. Di quanto tempo necessita il vostro algoritmo per produrre i primi risultati?
7. La vostra azienda come misura il ROI per questo tipo di tecnologie?

Appendix B – Interview Ghostwriter AI

Interviewer: Giovanni Quarantotto Vittori

Interviewee: Ester Liquori

QUESTION: Innanzitutto se mi puoi dare qualche dettaglio in più per quanto riguarda la struttura di GhostwriterAI, quante persone siete, da quanto tempo andate avanti, dal 2018 se non sbaglio, intanto partiamo con questo.

ANSWER: Ok allora l'azienda è Youaremyguide ed è nata fine 2014. Nel 2018 è nato il servizio Ghostwriter, ed è da quella data che siamo presenti sui portali AgID. Come i servizi precedenti anche Ghostwriter si basa su intelligenze artificiale. In realtà [questi servizi] sono stati il primo seme di quello che poi è stato, ed è, Ghostwriter. Ovvero un'analisi delle conversazioni, di attimi di conversazione a livello statistico sui comportamenti di viaggio. Per quanto riguarda Ghostwriter invece, Ghostwriter analizza a prescindere da quello che è il tema, il verticale del feedback. Quindi a seconda di quale azienda richiede il servizio, vengono analizzati quel tipo di dati.

QUESTION: Poi volevo chiederti chi sono i vostri competitors a livello italiano. Se c'è qualcuno simile a voi, se sapete un po' il mercato, se sentite una stretta come dire, una corsa alle armi, se mi passi il termine. Oppure è più tranquilla la situazione?

ANSWER: Mah allora, considera che, come avrai studiato e avrai visto dai vari servizi di intelligenza artificiale, in realtà intelligenza artificiale vuol dire tutto e vuol dire niente, perché sono algoritmi che vengono applicati a seconda dei casi. Allora i nostri sono applicati al testo e quindi un'analisi di tipo testuale. Competitors locali e competitors all'estero ce ne sono vari, ma anche in Italia ce ne sono di competitors. Il fatto è che a seconda di che, diciamo, dato di inizio e con quale obbiettivo si parte [sono differenti]. Noi ci muoviamo su 3, diciamo, filoni principali: che è quello dell'analisi conversazionale cioè del marketing di mercato, analisi conversazionale marketing per customer experience oppure a supporto dell'azienda che deve capire come sta andando con i KPIs o a supporto della parte d'interesse e poi l'ultimo oggetto, che è la più completa, da parte del marketing è quella [analisi] delle vendite. Quindi capisco, supporto, vendo, di fatto. Sulla parte di vendita abbiamo delle analisi, per esempio, del customer advertising, capire quale testo, quale campagna di advertising ha avuto più successo oppure [analisi di] attività di inbound e outbound del call center, quello che trovo nelle interazioni tra chi vende e il customer. Quindi a livello di competitors ce ne possono essere vari, nel senso che, potrei citarti Tigro, per quanto riguarda uno dei pezzetti che facciamo su quello che è la customer experience perché aiutiamo ad individuare i testi per il supporto tecnico. Lato più marketing potrei indicarti a Dada, ultimamente dicono che fa AI, prima non la facevano, non sappiamo poi dove sta la realtà. Poi altri, di sicuro non siamo competitors di aziende come IndigoAI o di visual bot, perché di fatto noi non facciamo chatbot, non facciamo virtual assistant di quel tipo. La nostra è più un analisi a, diciamo, posteriori, in genere stessa tecnologia, quindi è lo stesso tipo, diciamo, di tecnologia proprietaria, che possiamo sempre definire sotto il NLP, che però noi applichiamo in modo diverso.

QUESTION: Poi volevo chiederti, quando arriva un cliente che dice "ok sono convinto dalla vostra business proposition, implementiamola". A livello di dati cosa necessita fisicamente per partire?

ANSWER: Bah allora dipende da quello che è appunto l'obbiettivo. Cioè mediamente alla fine le informazioni si riducono a 3 elementi: che sono i dati CRN, dati che vengono dai canali social, quindi analytics dei social, o se devono analizzare le campagne gli insights della parte dei social delle campagne che fanno e se invece devono, diciamo, supportare, come dicevo, gli agenti che fanno l'help desk [guarderanno] le interazioni e gli scambi necessari per addestrare o dare dei boost agli agenti. Un dato che vogliono esporre agli agenti in modo che il sistema lo impari e ne comprenda il linguaggio naturale, sia capace di proporlo e poi invece si tratta proprio di dati analitici tutto quello che è l'interazione, gli scambi, i dialoghi, tutto quello che c'è dal CRM che sia utile per lo scopo.

QUESTION: Invece per quanto riguarda a livello di implementazione hardware, c'è qualche tipo di necessità? Oppure qualche infrastruttura particolare di cui hanno bisogno i clienti oppure è una cosa esterna che viene poi solo utilizzata?

ANSWER: No hardware assolutamente no, è tutto in cloud. C'è un livello base di software necessario ma non ce ne occupiamo noi. Noi lo forniamo, fondamentalmente, in cloud su licenza.

QUESTION: Sempre a livello di clienti, anche magari lavorando, non lo so, in un Paese che purtroppo noi vediamo spesso come retrogrado e l'hai scritto anche nel libro, le difficoltà più grandi che trovi tu, che trovate voi come azienda sia nel proporvi sia nell'implementare una cosa così innovativa.

ANSWER: Beh le difficoltà più grandi sono più che altro riconducibili al fatto che c'è tanto hype, attorno all'intelligenza artificiale, però c'è ancora poca applicazione. Quindi molto spesso le aziende o sono quelle molto grandi che allora ho hanno servizi di data science e magari sono degli organi internazionali allora si appoggiano a livello internazionale oppure sono quelle medio-piccole sono un po' confuse su quello che può fare e non può fare l'AI. Cercano magari di rivolgersi a big perché pensano che di sicuro avranno successo e alle volte restano un po' scocciati perché hanno a che fare con qualcosa di diverso. L'idea è che attacco l'AI a quello che sto facendo adesso e da domani wow ho potenziato l'intelligenza artificiale, no. Tra l'altro necessità di addestramento e conoscenza per portarti dei vantaggi che vedremo nel medio-lungo periodo. Nel breve [periodo]ci sono dei vantaggi sicuramente, però devono essere, ecco, considerati un primo seme su poi quello che è il vantaggio di maggior successo che deriva dall'addestramento costante del sistema che capisce come funzionare con la tua azienda. Capisce come va.

QUESTION: Infatti a proposito di tempo, mi pare che sul libro tu abbia menzionato un periodo di lavoro che va da qualche mese fino ad arrivare a degli anni, per creare un sistema, non voglio dire autonomo, ma con un buon grado di applicabilità.

ANSWER: È variabile, nel senso che i primi dati, i primi risultati, i primi elementi che puoi già applicare anche meno. Mediamente le aziende, già nel primo mese hanno già dei dati su cui sono già contenti perché non ce li avevano prima. Chiaro che se vuoi che il sistema sia totalmente autonomo, che abbia, che sia costante nel tempo nel livello di qualità di risposta che ti da eccetera, occorre più tempo. L'addestramento, nel caso per esempio di temi per il suggerimento del help desk, per quello che è l'analisi del call center e customer experience, suggeriamo almeno 3 mesi. Ma perché il sistema abbia a disposizione tutta la variabilità, almeno iniziale, necessaria per le varie tipologie di risposta perché naturalmente è difficile che un'azienda riceva abbastanza casi diversi, abbastanza varietà di casistica, affinché possa coprire una buona percentuale di risposta. Ovviamente più a lungo ti va, più vede storia, più accumulerà esperienza e quindi imparerà, quello non c'è dubbio.

QUESTION: Come misuriamo un ritorno d'investimento per un servizio come il vostro?

ANSWER: Il ritorno sull'investimento si calcola a seconda del tipo di servizio. Nel senso che, qual è il ritorno sull'investimento? Bisogna raggiungere l'obbiettivo. Se il mio obbiettivo è, per esempio, raggiungere con maggior successo i lead a cui trovo, che sto trovando con il mio call center di inbound telefonico, se prima dell'obbiettivo convertivo 100 su 1000 telefonate e poi dopo il servizio, utilizzando lo script, cambiando il modo in cui i miei collaboratori dialogano e così via, ne converto 120 ho ottimizzato [il processo] each. E anche là qual è l'obbiettivo?! Voglio fare questo? È una storia di successo? È come nel marketing, anche nel marketing devo stabile cosa voglio ottenere, per quanto tempo e con quanto budget, perché pure qualsiasi campagna di marketing posso dire "ah voglio più clienti" va bene, sei disposto ad investire un miliardo per avere 10% in più di clienti? Probabilmente no. Dipende che cosa hai a disposizione e che cosa vuoi ottenere.

QUESTION: Ultimissima domanda, poi ti lascio: se puoi dirmi quanti clienti gestite? Una media a portafoglio.

ANSWER: Sono variabili, questo dato è riservato perché tanti clienti non vogliono essere nominati quindi non te lo posso dire.

Appendix C – Interview Mapendo

Interviewer: Giovanni Quarantotto Vittori

Interviewee: Lorenzo Viscanti

QUESTION: Come prima domanda, volevo chiederti se potevi parlarmi della storia di Mapendo, di come vi è venuta l'idea, un po' del vostro background e della vostra crescita.

ANSWER: Allora Mapendo è stata fondata un po' per caso da me ed un mio compagno di studi di ingegneria. Ci eravamo sempre lasciati dicendo che avremo lavorato assieme, poi ad un certo punto si è concretizzata l'occasione di fare qualcosa assieme e dato che lavoravamo nel settore della pubblicità, pian piano è venuta fuori un'opportunità, abbiamo visto, su quello che era il mondo delle app e del mobile. 2015 quindi un po' prima, nello sviluppo diciamo del mobile, rispetto allo stato attuale. Una situazione un po'più vantaggiosa di quello attuale. In pratica noi facciamo degli algoritmi di ottimizzazione che permettono una gestione dell'advertising abbastanza efficiente, fuori dal mondo di Google e Facebook. I nostri clienti appunto vengono per comprare, noi ci assumiamo il rischio del pagare click ed impressioni, che è l'unità di pagamento della pubblicità e a loro vendiamo nuovi utenti, quindi un modello totalmente senza rischio, per loro.

QUESTION: questo discorso del rischio l'altra volta non era venuto fuori, quindi adesso volevo espandere un attimino qui. Come lo sentite questo rischio? Chiaramente c'è un mark-up immagino: voi vi assumete il rischio in cambio di un costo un po' più alto rispetto ai vostri competitors.

ANSWER: Noi abbiamo una serie di dati che ci consentono di lavorare con un'abbastanza tranquilla sicurezza di fare del margine.

QUESTION: Poi invece, per quanto riguarda i vostri piani per il futuro, i vostri target per il futuro per i prossimi anni a cosa state puntando?

ANSWER: Allora noi ci stiamo espandendo sul mercato americano, che rappresenta l'80% delle nostre revenues. Abbiamo in mente una dimensione, che è la dimensione target a cui vogliamo arrivare come attività gestita ogni anno. Dopodiché contiamo di aggregare attività simili a quelle che facciamo, quindi espanderci a settori simili per continuare a crescere. Quindi insomma vogliamo restare un'azienda indipendente grande e lavorare dall'Italia sul mercato americano, crescendo sempre di un 20-30% anno su anno.

QUESTION: E come mai questa decisione di rimanere fissi in Italia, based in Italy, e non andare oltre oceano, sebbene abbiate un sacco di clienti oltre oceano.

ANSWER: Posso dartene due di ragioni. Una sentimentale ed emozionale diciamo, ed una basata sul calcolo. Quella emozionale è abbastanza semplice: siamo di qua e quindi viviamo e lavoriamo qui pur avendo uno sguardo su tutto mondo perché adesso non puoi non avere uno sguardo globale no?! Nel mondo del lavoro. Però diciamo, soprattutto nel digitale, fare una cosa dall'Italia dandosi un confine che è quello dell'Italia, sarebbe ridicolo diciamo. D'altra parte, invece, l'opportunità è nel fatto che qui riusciamo ad avere dei costi molto più bassi rispetto agli Stati Uniti. Quindi questo ci consente anche di essere competitivi, è un altro motivo. C'è gente abbastanza brava mercato del lavoro, non è perfetto, soprattutto nelle materie stem. Ovviamente, come sappiamo tutti, non siamo al top però nonostante questo c'è gente molto molto brava quindi il livello medio è piuttosto alto, il livello medio delle persone che si trovano. Costi più bassi, livello medio di qualità delle persone piuttosto alto è un'opportunità. E poi viviamo qui quindi, nonostante tutto c'è una buona qualità della vita.

QUESTION: Poi volevo chiederti se a livello italiano, potevi indicarmi i vostri competitors, e se nel caso questi non ci sono a livello italiano, nelle nostre immediate vicinanze quindi livello europeo. O addirittura se anche qui non ci arriviamo a livello statunitense.

ANSWER: Allora esistono alcuni competitors in Europa, fammi considerare ad esempio ancora Europa non ovviamente l'unione Europea ma Europa e il Regno Unito. C'è la più grande agenzia di pubblicità e marketing indipendente che si chiama M&C Saatchi. E loro, una parte consistente della loro attività è simile a quella che facciamo noi. Però loro sono un'agenzia quindi non hanno tecnologia, semplicemente fanno da, non so come dire, connettore tra le piattaforme tecnologiche, che saremmo noi, ed il cliente finale. Loro sono un esempio. In Germania ci sono uno o due società, te ne cito una che si chiama Applift, App come applicazione. Questa ad esempio ecco, è una società che potrebbe essere considerata un nostro competitors circa. Fanno un po' più giochi mentre noi siamo arrivati, abbiamo cominciato a lavorare in un momento in cui il mondo dei giochi su mobile era già partito. Quindi avevamo un anno di ritardo circa, ci siamo indirizzati più sul mondo e-commerce e servizi digitali che sul mondo giochi, però adesso stiamo facendo qualcosa sui giochi ad esempio.

QUESTION: Collegandoci un po’ alla prima domanda che ti ho fatto, in cui stavi spiegando i vostri servizi, a livello di dati, in termini di dati da parte dell’azienda, avete necessità di qualcosa di particolare per iniziare i vostri progetti?

ANSWER: No perché noi, uno dei nostri punti di forza è quello di lavorare con il GDPR e quindi portare questo modello, basato sul GDPR anche negli stati uniti, dove non è legge, ovviamente perché è un regolamento europeo. Però questo fatto di rispettare un regolamento più stretto del loro sulla privacy, che è un tema sensibile per tutti, piace anche a loro, quindi commercialmente da questo punto di vista è vincente. Noi non usiamo dati della società stessa. In questo modo è tutto molto più pulito e tranquillo. Ci sono meno rischi legali e quindi è meglio.

QUESTION: Invece in termini un po’ più materiali, in termini di hardware avete qualche necessità particolare per lanciare i vostri progetti o si basa più su servizi cloud?

ANSWER: Allora l’hardware è un problema risolto nel senso che ci basiamo totalmente su servizi cloud. Quello che non è un problema risolto, invece è la questione della banda è della velocità. Noi abbiamo bisogno di raggiungere tutti i cellulari negli Stati Uniti, però soprattutto lì perché lavoriamo soprattutto lì, ma anche nel resto del mondo dobbiamo raggiungere tutto, qualsiasi cellullare deve raggiungersi in meno di 50 millisecondi come tempo di latenza di banda, quindi legato alla velocità della banda. Di conseguenza è veramente molto molto importante la questione banda per quanto ci riguarda.

QUESTION: Perché sono importanti questi 50 ms di latency?

ANSWER: Di solito è il tempo di latenza che l’utente non percepisce come un ritardo, viene usato come standard nel nostro settore. 50 o 100 millisecondi, però cambia poco cambia poco, diciamo, tra 50 e 100 millisecondi, l’ordine di grandezza è lo stesso. Di sicuro non puoi andare a 200 millisecondi, come tempo medio, perché altrimenti l’utente vede qualcosa che non va. E l’esperienza utente è meno fluida e quindi succede qualcosa: i tassi di conversione peggiorano, queste cose qua.

QUESTION: Volevo chiederti quali sono le maggiori difficoltà che i vostri clienti trovano nell’adottare i sistemi di intelligenza artificiale e quali sono le difficoltà che trovate voi quando avete a che fare con loro.

ANSWER: Allora questa è una cosa un po' particolare. Noi ragioniamo come una black box nei confronti del cliente, quindi il cliente ha dei parametri di misurazione che sono standard delle API di integrazione che sono standard, quindi il cliente è pronto a partire. Poi la parte di machine learning, che è un subset di intelligenza artificiale, che noi usiamo è invisibile al cliente tranne in questi parametri che sono i KPIs dell'attività che il cliente ovviamente visualizza. Però cerchiamo di astrarre da questo, dato che il cliente è un team marketing quello con cui lavoriamo, meno vedono di algoritmi meglio è. Mentre invece sanno leggere molto bene i KPIs. Quindi noi dobbiamo lavorare di machine learning e dar fuori dei risultati che vengono elaborati a livello statistico per vedere la performance che stiamo raggiungendo. Quello è molto molto importante invece. Quindi la lettura dei KPIs.

QUESTION: Invece a livello di avere a che fare con clienti tra il vecchio continente ed il nuovo continente, puoi dirmi qualcosa in merito?

ANSWER: Allora è tutto un po' più veloce negli Stati Uniti. L'integrazione, la partenza con un cliente nuovo, quindi il contatto il contatto commercial e l'integrazione è più veloce. In Europa è un po' più modello basato sulla conoscenza, su di loro [Stati uniti] è un modello basato sul test, risultato, scalare. Quindi mi convinci per quello che offri, che io sto cercando in quel momento, perché poi il timing nel contattare il cliente è la cosa più decisiva di tutte no?! Più della qualità di quello che fai. Se arrivi mentre il cliente sta cercando qualcuno nuovo, il cliente anglosassone lo fa un test con te. Se tu porti risultati il cliente anglosassone scala nei numeri con te, cresce nei numeri con te. E questo è molto molto veloce, il cliente anglosassone, il primo test, in 24-36 ore dal contatto, dal primo contatto, parte. Il cliente Europeo ha dei tempi totalmente diversi. L'italiano è nella media di quello Europeo, non sputo sul cliente italiano. Quando abbiamo avuto a che fare con clienti italiani sono un po' più lenti, come lo sono tutti i clienti europei rispetto a quelli anglosassoni.

QUESTION: Visto che mi hai detto 24-36 ore per gli anglosassoni, il cliente europeo in media?

ANSWER: La partenza di un fornitore nuovo è molto diversa. Ma nel nostro settore ci sono delle differenze, anche perché noi quando lavoriamo con il cliente americano, il cliente americano lavora direttamente con noi perché ha fatto fuori un'agenzia, una società di consulenza esterna che fungeva da intermediario, nella stragrande maggioranza dei casi.

Mentre il cliente Europeo ha un intermediario, che è un agenzia, il che rende più lento il tutto. Quindi ci sono ragioni sia culturali sia di organizzazione del business, fammi dire.

QUESTION: Poi volevo chiederti, in termini di tempo, quanto tempo necessita il vostro sistema per iniziare a produrre risultati.

ANSWER: Diciamo che la fase di learning, arriva al 90% del risultato in circa due settimane. Poi ci possono essere fino a 3 mesi per arrivare ad un risultato ottimo, però diciamo 2 settimane. Attenzione poi ci sono anche delle condizioni di mercato che cambiano, questo influisce sui tempi, ti faccio un esempio: novembre, dicembre sono mesi in cui tutto il mondo dell'advertising viene spinto molto anche con attività di branding. Quindi l'apprendimento [del sistema] è molto più facile e molto più veloce, perché ci sono una serie di attività di marketing complesse messe in piedi dalle aziende. Gennaio, Febbraio, Marzo sono momenti in cui le aziende investono di meno e quindi vuol dire che anche il learning è un po' più complicato perché in generale il cliente riceve meno messaggi. Ti faccio un esempio: se a Ottobre, Novembre, Dicembre un nostro cliente sta investendo in una campagna di affissione sulle strade di una città americana, questo aiuta anche noi e ci sono più canali di marketing attivi: c'è Google, c'è Facebook ci sono le affissioni sulle strade, ci sono le tivù e insistere su tanti canali aiuta tutti i canali a performare meglio. In un'altra fase dell'anno, in cui i budget sono minori, e ci sono meno canali attivi, come questo periodo dell'anno cui, Gennaio-Marzo, il tempo di learning è più lungo. Scusa ti ho dato una risposta un po' complicata però per spiegarti che ci sono mille fattori che contribuiscono, non è solo una cosa puramente algoritmica nostra, dipende dall'environnement diciamo.

QUESTION: Nessun problema. Infatti, ti avevo fatto questa domanda perché avevo visto sul vostro sito che c'erano i case studies, per esempio quello "photo printing" dove c'era un tempo indicato di 14 settimane, invece quello "public transportation" di 18 e allora volevo farmi un'idea.

ANSWER: Io, infatti, ti ho detto 3 mesi e siamo lì no?! 3 mesi, 3 mesi e mezzo come tempo. Sì però poi ci sono delle differenze.

QUESTION: Parliamo un po', anche se hai già risposto in parte nelle domande precedenti, a livello di risultati. Uno dei principali compiti del marketing è quello di generare lead e vendite, e tu mi hai parlato prima di come, quando hai a che fare con le persone del *marketing è meglio che vedano più risultati che processi di machine learning. Per misurare l'impatto positivo dei vostri servizi, oltre a KPIs classici come ROI etc, avete qualche altro metro per far vedere che avete portato un risultato positivo?*

ANSWER: Allora c'è il costo di acquisizione dell'utente, quanto mi costa un lead, un utente registrato, questo è il metodo standard da cui si calcola il ROI o il ROAS ovviamente. Più ci sono dei parametri tecnici nel mondo delle app che sono importanti. Ad esempio, un KPI sempre utilizzato è il tempo tra l'ultimo click su una campagna di advertising ed il termine del lead, si chiama click to lead time. È un parametro tecnico che però tutti conoscono nel nostro mondo ed è molto molto importante, deve essere tendenzialmente piuttosto basso.

QUESTION: Se potevi parlarmi di Jenga, del vostro algoritmo proprietario. In generale, non troppo specifico, so che siamo un po' oltre.

ANSWER: Molto semplicemente, Jenga viene dal nome del gioco con i bastoncini di legno no?! In cui si fa una torre e bisogna levare un pezzo alla volta. Perché Jenga si basa su un sistema di machine learning che si chiama pre-emptive learning e semplicemente è un modo per rendere più efficacie la campagna su 3 dimensioni: i costi, la riduzione dei costi per il nostro cliente advertiser, la qualità ini termini di KPIs come quelli che ti dicevo prima e il terzo punto è la nostra marginalità. Sono questi i 3 parametri esterni, diciamo, poi dentro la black box ci sono una pletora di parametri che noi utilizziamo e che servono a capire in che modo andare ad ottimizzare il cliente. Infatti, ti faccio un esempio, sembrano delle cose banali, ma ci sono delle cose strane, non banali: a noi, ad esempio, contano tantissimo la geolocalizzazione della città dell'utente. A seconda di ogni attività, il tasso di conversione di un utente iscritto ad un servizio a seconda che sia in una città sotto i 50mila e sopra i 50mila abitanti, ci sono delle differenze enormi. Quindi saper incrociare questo dato in modo veloce ha un fattore determinante sui margini di un'attività come la nostra. Te ne ho detto uno.

QUESTION: Fammi capire meglio quest'ultima. Che differenze ci sono, per esempio, tra un utente che abita in una città sotto i 50 e sopra i 50 [mila abitanti]?

ANSWER: Sostanzialmente i budget di pubblicità si concentrano talmente nelle città più grandi, dove la battaglia è talmente più alta che i tassi di conversione crollano a causa della battaglia. Quindi se vai fuori hai dei tassi di conversione più alti per la stragrande maggioranza delle attività. Però invece ci sono attività che non sono interessanti fuori dalle città, quindi ci sono una serie di fattori socio-economico-demografici che ritornano un po' per tutti i servizi dei nostri clienti che determinano in che modo può essere efficacie il tutto. Sono cose abbastanza generiche, non è detto che siano sempre vere eh?! Ti sto dicendo dei dati così, però ad esempio se tu vuoi avere le campagne che parlano di sport, noi lavoriamo molto con dei servizi legati allo sport, funzionano molto meglio se tu analizzi il reddito medio del quartiere e delle città: nelle città a reddito medio-basso hanno un tasso d'interesse incredibilmente più alto per lo sport. Se tu tieni dentro alcuni fattori di questi, che sono validi per tutti i clienti che lavorano nello sport, sono un input agli algoritmi, poi all'algoritmo glielo dici in un altro modo, all'algoritmo glielo dici appunto dicendo che concentrandoci solo su certe province, tipo in Italia, o stati negli Stati Uniti, che però vedi che ci sono dei parametri che cambiano e questi [parametri] te li porti dietro. È, diciamo, il riflesso tecnologico di certi parametri socio-economici. Ti faccio un esempio: negli Stati Uniti gli operatori telefonici del cellulare no?! Nelle aree urbane è più diffuso T-mobile, nelle aree extra-urbane è più diffuso AT&T, lo sport converte meglio sugli utenti AT&T, perché sono tendenzialmente nelle aree rurali. Quindi ci sono un po' di trucchetti, che dopo un po' si imparano e ti permettono di ottimizzare le attività.

QUESTION: Il fatto di questa differenza tra gestori telefonici, è molto più forte, per esempio, in America o lo vedete anche, che ne so, qua in Italia qua in Europa?

ANSWER: È Uguale qua in Italia. C'è dappertutto questa differenza. Non sembra ma sono delle stratificazioni socioeconomiche della società, che si riflettono in delle cose come gli operatori telefonici. Ma poi sono banali [queste stratificazioni] cioè ad esempio, che ne so, ti dico una cretinata: indirizzi IP simili sono dati a connessioni in fibra ottica dello stesso quartiere e quindi presentano caratteristiche socio-demografiche simili, e quindi se tu ragioni sull'indirizzo IP hai beccato persone dello stesso quartiere che convertiranno in modo abbastanza simile tra di loro. Quindi partendo dalla tecnologia, partendo solo da un dato tecnologico, che è l'indirizzo IP, però riesci ad inferire della conoscenza che è socio-demografica che faresti fatica a mettere in mezzo.

QUESTION: Ultima cosa poi ti lascio stare. Quanto, influisce il vostro, l'input umano e quanto impatta il fatto che riesca ad imparare da solo?

ANSWER: Nono l'input umano è fondamentale nella capacità di razionare l'algoritmo. Ci sono delle cose, ad esempio dei dati di input che dai tu, ad esempio il fatto che qualcosa andrà meglio il fine settimana che durante la settimana che è una conoscenza umana che non devi dare in pasto all'algoritmo. Cioè gliela devi mettere tu. Anche perché è difficilmente parametrabile no?! Quindi non è che potremo fare un set di regole, è molto più semplice avere nel team 2 o 3 persone esperte che immediatamente cablano questo tipo di regole per l'algoritmo, senza che sia l'algoritmo a scoprirlo. Poi ogni tot, se hai un bravo team, lavori con un team bravo, ogni tot di tempo vai a rivedere le regole di inferenza che applichi ogni volta e dici "ma senti avevo sbagliato qualcosa?" e te le rigiochi.

List of figures

References

Blerina Lika Kostas Kolomvatsos, and Stathes Hadjiefthymiades Facing the cold start problem in recommender systems [Article] // Expert Systems with Applications. - Athens : Elsevier, 2014. - 41.

bloomreach [Online]. - 2020. - 6 December 2020. - https://www.bloomreach.com/en.

Boam Eric I Decoded the Spotify Recommendation Algorithm. Here's What I Found. [Online] // Medium. - January 14, 2019. - November 15, 2020. - https://medium.com/@ericboam/i-decoded-the-spotify-recommendation-algorithm-heres-what-i-found-4b0f3654035b.

Bostrom Nick Superintelligence [Book]. - Oxford : Oxford University Press, Reprinted edition (1st April 2016), 2014.

Brent Smith and Greg Linden Two Decades of Recommender Systems at Amazon.com [Article] // IEEE Internet Computing. - 2017. - Vol. 21.

Brighton Jamie Getting Started with Artificial Intelligence Today [Online] // Adobe Digital Experience Blog. - Adobe, August 22, 2016. - December 26, 2020. - https://blogs.adobe.com/digitaleurope/digital-marketing/getting-started-with-artificial-intelligence-today-its-easier-than-you-think/.

Brown Gerry Machine Learning Will Revolutionize Market Segmentation [Online] // Qubit. - IDC, January 2017. - 1.0. - 3 November 2020. - https://www.qubit.com/research/machine-learning-revolutionizes-segmentation-practices/.

Calo M. Ryan Against Notice Skepticism in Privacy (and Elsewhere) [Article] // Notre Dame Law Review. - 2013. - 87 : Vol. 3.

CCPA California Consumer Privacy Act (CCPA) [Online] // State of California Department of Justice. - 2018. - January 16, 2021. - https://oag.ca.gov/privacy/ccpa#top.

Chan Rosalie The Cambridge Analytica whistleblower explains how the firm used Facebook data to sway elections [Article] // Business Insider. - 2019.

Chen Bo Zhao & Weiquan Data Protection as a Fundamental Right: The European General Data Protection Regulation and Its Exterritorial Application in China [Article] // US-China Law Review. - Tilburg, Netherlands : DAVID PUBLISHING, 2019. - 3 : Vol. 6.

Cole David The Chinese Room Argument [Report]. - Stanford : The Metaphysics Research Lab, Center for theStudy of Language and Information (CSLI), Stanford University, 2020.

Coperland B. Jack The Turing Test, The Elusive Standard of Artificial Intelligence [Book]. - Dartmouth College, Hanover, U.S.A. : SPRINGER-8CIENCE+BUSINESS MEDIA, B.V., 2003.

Courville Ian Goodfellow and Yoshua Bengio and Aaron Deep Learning [Book]. - [s.l.] : MIT Press, 2016.

DEPARTMENT OF ECONOMIC AND SOCIAL AFFAIRS - UNITED NATIONS Industrial Development for the 21st Century: Sustainable Development Perspectives [Report]. - New York : [s.n.], 2007.

eMarketer Artificial Intelligence Roundup [Online]. - 2018.

eMarketer Artificial Intelligence Roundup [Online]. - 2019.

Emarsys Emarsys Home Page [Online]. - December 22, 2020. - https://emarsys.com/.

Eric T. Bradlow Manish Gangwar, Praveen Kopalle and Sudhir Voleti The Role of Big Data and Predictive Analytics in Retailing [Article] // Journal of Retailing. - New York : Elsevier Inc., 2017. - 93.

European Commission How the EU determines if a non-EU country has an adequate level of data protection. [Online] // European Commission Website. - https://ec.europa.eu/info/law/law-topic/data-protection/international-dimension-data-protection/adequacy-decisions_en.

Flybits Flybits home page [Online]. - 2020. - https://www.flybits.com/.

Gartner Gartner Says By 2020, Artificial Intelligence Will Create More Jobs Than It Eliminates [Online] // Gartner Says. - Gartner, December 13, 2017. - December 13, 2020. - https://www.gartner.com/en/newsroom/press-releases/2017-12-13-gartner-says-by-2020-artificial-intelligence-will-create-more-jobs-than-it-eliminates.

Gartner Gartner Says Global Artificial Intelligence Business Value to Reach $1.2 Trillion in 2018 [Online] // Gartner Says. - April 25, 2018. - December 13, 2020. - https://www.gartner.com/en/newsroom/press-releases/2018-04-25-gartner-says-global-artificial-intelligence-business-value-to-reach-1-point-2-trillion-in-2018.

Gartner Gartner Says Smart Machines Will Enter Mainstream Adoption By 2021 [Online] // Newsroom Press Release. - December 16, 2016. - December 13, 2020. - https://www.gartner.com/en/newsroom/press-releases/2016-12-15-gartner-says-smart-machines-will-enter-mainstream-adoption-by-2021.

Gartner Gartner Survey Finds Consumers Would Use AI to Save Time and Money [Online] // Gartner Newsroom. - Gartner, September 12, 2018. - December 13, 2020. - https://www.gartner.com/en/newsroom/press-releases/2018-09-12-gartner-survey-finds-consumers-would-use-ai-to-save-time-and-money.

Gartner Gartner Survey Reveals 66% of Organizations Increased or Did Not Change AI Investments Since the Onset of COVID-19 [Online]. - Gartner, 2020. - December 13, 2020. - https://www.gartner.com/en/newsroom/press-releases/2020-10-01-gartner-survey-revels-66-percent-of-orgnizations-increased-or-did-not-change-%E2%80%A61/3.

Gavira Mario Amazon´s Recommendation Engine Secret Sauce [Online] // LinkedIn. - July 13, 2018. - November 15, 2020. - https://www.linkedin.com/pulse/amazons-recommendation-engine-secret-sauce-mario-gavira/.

GDPR Official Journal of the European Union [Online] // REGULATION (EU) 2016/679 OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL of 27 April 2016. - 2016. - https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32016R0679.

Ghosh Dipayan Law, What You Need to Know About California's New Data Privacy [Article] // Harward Business Review. - 2018.

GhostwriterAI Homepage GhostwriterAI homepage [Online]. - https://ghostwriter.ai/.

Google Keynote I/O - 8th May 2018 Keynote (Google I/O '18). - [s.l.] : Google Developers, 2018. - Vol. https://www.youtube.com/watch?v=ogfYd705cRs.

Granville Vincent 24 Uses of Statistical Modeling [Online] // Data Science Central. - TechTarget, December 14, 2015. - December 26, 2020. - https://www.datasciencecentral.com/profiles/blogs/top-20-uses-of-statistical-modeling.

Haenlein Michael Siri, Siri, in my hand: Who's the fairest in the land? On the interpretations, illustrations, and implications of artificial intelligence [Article] // Business Horizon. - Paris : Elsevier, 2018.

Ian MacKenzie Chris Meyer, and Steve Noble How retailers can keep up with consumers [Online]. - McKinsey, October 1, 2013. - December 27, 2020. - https://www.mckinsey.com/industries/retail/our-insights/how-retailers-can-keep-up-with-consumers.

IBM Cloud Education What is Strong AI? [Online] // IBM Cloud Learn Hub. - IBM, August 31, 2020. - 1.0. - November 08, 2020. - https://www.ibm.com/cloud/learn/strong-ai#:~:text=Strong%20AI%20vs.,-weak%20AI&text=Weak%20AI%2C%20also%20known%20as,to%20solve%20for%20new%20problems..

Jelinek Andrea Guidelines 3/2018 on the territorial scope of the GDPR - (Article 3) - Version for public consultation [Report]. - [s.l.] : European Data Protection Board, 2018.

Julia Angwin Jeff Larson, Surya Mattu and Lauren Kirchner Machine Bias, There's software used across the country to predict future criminals. And it's biased against blacks. [Online] // ProPublica. - May 23, 2016. - December 6, 2020. - https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing.

Kaplan Jerry Artificial Intelligence, what everyone needs to know [Book]. - [s.l.] : Oxford University Press, 2016.

Kavlakoglu Eda AI vs. Machine Learning vs. Deep Learning vs. Neural Networks: What's the Difference? [Online] // IBM Cloud. - IBM, May 26, 2020. - October 3, 2020. -

https://www.ibm.com/cloud/blog/ai-vs-machine-learning-vs-deep-learning-vs-neural-networks.

Levine Barry Salesforce explains why, when it has Einstein, it needs Watson's intelligence [Article] // MarTech Today. - 2017.

Li Hui Which machine learning algorithm should I use? [Online] // The SAD Data Science Blog. - SAS, April 12, 2017. - November 03, 2020. - https://blogs.sas.com/content/subconsciousmusings/2017/04/12/machine-learning-algorithm-use/.

Linden Jennifer Jacobi Eric Benson Gregory Personalized recommendations of items represented within a database [Patent]. - United States, May 7, 2001.

Liquori Ester Marketing (Artificialmente) Intelligente [Book]. - [s.l.] : Flaccovio Dario, 2020.

Lucy Tour [Online] // Lucy Website. - https://www.lucy.ai/.

Lucy Website - About Us About Us [Online] // Lucy Website. - https://www.lucy.ai/about-us.

Michael Chui James Manyika , Mehdi Miremadi, Nicolaus Henke , Rita Chung, Pieter Nel, Sankalp Malhotra Notes from the AI frontier: Applications and value of deep learning [Report]. - [s.l.] : McKinsey Global Institute, 2018.

Michael Chui Sankalp Malhotra Adoption of AI advances, but foundational barriers remain [Online]. - McKinsey, November 3, 2018. - September 5, 2020. - https://www.mckinsey.com/featured-insights/artificial-intelligence/ai-adoption-advances-but-foundational-barriers-remain.

Moor James, H. An Analysis of the Turing Test [Journal]. - [s.l.] : D. Reidel Publishing Company, 1976. - Vol. Philosophical Studies.

Neiman Marcus Neiman Marcus Home Page [Online]. - December 2020. - https://www.neimanmarcus.com/en-it/.

Omale Gloria Gartner Predicts 25 Percent of Digital Workers Will Use Virtual Employee Assistants Daily by 2021 [Online] // Gartner Predicts. - Gartner, January 9, 2019. - December 13, 2020. - https://www.gartner.com/en/newsroom/press-releases/2019-01-09-gartner-predicts-25-percent-of-digital-workers-will-u.

Outernets Outernets Home Page [Online]. - 2020. - https://outernets.co/.

Paka Amit How Does The CCPA Impact Your AI? [Article] // Forbes Technology Council. - 2020.

Panchak Willem Sundblad & Patricia Data Is The Foundation For Artificial Intelligence And Machine Learning [Article] // Forbes. - 2018.

Pointillist Home Page Pointillist Home Page [Online] // Pointillist website. - December 29, 2020. - https://www.pointillist.com/.

Salesfoce Einstein Overview - Website Salesforce Einstein Overview [Online] // Salesfoce Website. - https://www.salesforce.com/eu/products/einstein/overview/.

Salesforce Website What is Salesforce [Online] // Salesforce Website. - Salesforce. - https://www.salesforce.com/eu/products/what-is-salesforce/.

Sartor Giovanni The impact of the General Data Protection Regulation (GDPR) on artificial intelligence [Report]. - [s.l.] : EPRS - European Parliamentary Research Service, 2020.

Sodha Sean How to get started with Natural Language Processing [Online] // IBM Watson Blog. - IBM, October 3, 2019. - November 3, 2020. - https://www.ibm.com/blogs/watson/2019/10/how-to-get-started-with-natural-language-processing/#:~:text=Natural%20language%20processing%20is%20the,analyze%2C%20and%20understand%20human%20language.&text=Watson%20Natural%20Language%20Understanding%20(NLU)%20i.

Sterne Jim Artificial Intelligence for Marketing - Practical Applications [Book]. - Hoboken, New Jersey : John Wiley & Sons, Inc., 2017.

Stevens Steven Rosenbush and Laura At UPS, the Algorithm Is the Driver [Article] // Wall Street Journal. - 2015.

Task Force IA - AgID website Ecosistema Intelligenza Artificiale [Online] // AgID. - 2020. - https://ia.italia.it/ia-in-italia/.

Taulli Tom CCPA: What Does It Mean For AI (Artificial Intelligence)? [Article] // Forbes. - 2019.

THE ECONOMIST The world's most valuable resource is no longer oil, but data [Article] // The Economist. - New York : [s.n.], 2017.

Theobald Oliver Machine Learning for Absolute Beginners: A Plain English Introduction [Book]. - [s.l.] : Scatterplot Press, 2017.

Thunderhead Home Page Thunderhead Our Product [Online] // Thunderhead Website. - December 29, 2020. - https://www.thunderhead.com/our-product/.

Tonner Andrew 9 Artificial Intelligence Stats That Will Blow You Away [Online] // FOX Business. - FOX, December 10, 2019. - 1.0. - November 3, 2020. - https://www.foxbusiness.com/markets/9-artificial-intelligence-stats-that-will-blow-you-away.

William D. Eggers Rob Hamill and Abed Ali Data as the new currency [Article] // Deloitte Review. - 2013. - Vol. 13.

Yu Allen How Netflix Uses AI, Data Science, and Machine Learning — From A Product Perspective [Online] // Medium. - February 27, 2019. - November 15, 2020. - https://becominghuman.ai/how-netflix-uses-ai-and-machine-learning-a087614630fe.