MASTER OF SCIENCE IN MANAGEMENT ENGINEERING

POLITECNICO
MILANO 1863

MASTER THESIS

Artificial Intelligence Adoption: State of the Art and Collection of Enterprise Applications

Supervisor:
Prof. Giovanni Miragliotta

Co-Supervisors:
Eng. Stefano Garavaglia
Eng. Carlo Negri

Master Thesis by:
Arturo Missaglia - 913430
Luca Pallavicini - 914001

Academic Year 2020/2021

Acknowledgements

First of all, we would like to express our gratitude to our supervisor, Professor Giovanni Miragliotta, for giving us the opportunity to work on this topic. Particularly, we want to thank our co-supervisors, Stefano Garavaglia e Carlo Negri: your guidance, competence and valuable advices during the thesis were essential for the quality of the final outcome.

We want to thank our families for giving us the possibility to pursue our dreams. Everything you did for us made us the people that we are today and we will be forever grateful to you.

Last but not least, we want to thank all our friends: thank you for the great memories and experiences shared together, we feel blessed to have you by our side.

“The best companies in the world never get nostalgic; They evolve”

Shawn Kanungo

Index

Index of Figures

Abstract (English)

Despite Artificial Intelligence has been a research field for a long time, practical applications in business have emerged only in recent years. Although the presence of an extensive body of literature related to Artificial Intelligence techniques, information about the adoption of Artificial Intelligence in enterprises is limited and mainly referred to past years. However, Artificial Intelligence is a rapidly growing technology and its enterprise adoption is rapidly evolving, implying the necessity to provide a constantly updated picture of its evolution. Therefore, the primary contribution of this thesis consists in investigating the current state of adoption of Artificial Intelligence solutions at international level. For this purpose, a database including Artificial Intelligence initiatives has been created, considering the first 235 firms of the Forbes Global 2000 list and identifying 1089 projects for 165 companies. Results show how Artificial Intelligence is rapidly spreading in a business context, with a variety of emerging applications in several industries and domains. Moreover, the broad applicability of the technology enables companies in different sectors to adopt it at advantage of their specific business needs. In this regard, the second contribution of this thesis consists in presenting a focus on how Artificial Intelligence solutions are currently being adopted in four industries: Food & Beverage; Manufacturing; Retail; Banking & Finance. For this purpose, a qualitative framework has been developed to support the analysis and presentation of conclusions. Results show how Artificial Intelligence is transforming and revolutionizing each of these sectors, supporting each aspect of their business with different solutions.

Abstract (Italian)

Nonostante l'Intelligenza Artificiale sia da tempo oggetto di ricerca, applicazioni pratiche in ambito aziendale sono emerse solo negli anni recenti. A dispetto di una vasta letteratura riguardante le tecniche di Intelligenza Artificiale, le informazioni sull'adozione di questa tecnologia nelle imprese sono limitate e si riferiscono prevalentemente ad anni precedenti. Tuttavia, l'Intelligenza Artificiale è una tecnologia in rapida crescita e la sua adozione a livello aziendale è in continua evoluzione. Questo implica la necessità di fornire una panoramica costantemente aggiornata sulla sua adozione. Di conseguenza, il contributo primario di questa tesi consiste nell''indagare lo stato attuale dell'adozione di soluzioni di Intelligenza Artificiale a livello globale. A tal fine, andando a considerare le prime 235 aziende della lista Forbes Global 2000, è stato creato un database con 1089 progetti di Intelligenza Artificiale relativi a 165 aziende. I risultati mostrano come l'Intelligenza Artificiale si stia rapidamente diffondendo in un contesto aziendale, con varie applicazioni in diversi settori e ambiti. Inoltre, la vasta applicabilità dell'Intelligenza Artificiale consente ad aziende operanti in diversi settori di adottarla a supporto delle proprie esigenze. In tal senso, il secondo contributo di questa tesi consiste nel presentare una panoramica di come l'Intelligenza Artificiale stia venendo adottata in quattro settori: Food & Beverage; Manufacturing; Retail; Banking & Finance. A tal fine, è stato appositamente sviluppato un modello qualitativo per supportare l'analisi e la presentazione delle conclusioni. I risultati mostrano come l'Intelligenza Artificiale stia trasformando e rivoluzionando ciascuno di questi settori, supportando ogni aspetto del loro business con soluzioni diverse.

Executive Summary

Artificial Intelligence has been a field of research for a long time, with the birth of the discipline dating back to 1950, with Alan Turing's seminal paper "Computing Machinery and Intelligence". However, practical applications in businesses have emerged only in recent years, generating great interest by public and private companies.
Although the presence of an extensive body of literature related to Artificial Intelligence techniques, information about the adoption of Artificial Intelligence in enterprises is limited and mainly referred to past years. Moreover, the technology is rapidly growing and spreading. Therefore, its current state of adoption changes year by year, implying the necessity to provide a constantly updated picture about its evolution.
As a result, the primary contribution of this thesis consists in investigating the current adoption of Artificial Intelligence by enterprises at a global level. For this purpose, a database including Artificial Intelligence initiatives has been created, considering the first 235 firms of the Forbes Global 2000 list, and identifying 1089 projects for 165 companies.
Results show how Artificial Intelligence is rapidly spreading in a business context, with a variety of emerging applications in several industries and domains. Particularly, detailed results have been provided for each of the nine Classes of Solutions in Artificial Intelligence proposed by the Artificial Intelligence Observatory of Politecnico of Milan. The research highlights how companies are already benefiting from the advantage of implementing Artificial Intelligence solutions, with benefits so significant that no company can ignore and stay out of the ongoing Artificial Intelligence revolution.
Moreover, the broad applicability of the technology enables companies in different sectors to adopt it at advantage of their specific business needs. In this regard, the second contribution of this thesis consists in providing a focus on how Artificial Intelligence is currently being adopted in four different industries: Food & Beverage; Manufacturing; Retail; Banking & Finance. For this purpose, a qualitative framework has been developed to support the analysis and presentation of conclusions. Results show how Artificial

Intelligence is transforming and revolutionizing each of these sectors, supporting each aspect of their business with different solutions.

Starting from Chapter 1, it includes a Literature Review about Artificial Intelligence, with an explanation of the main theoretical concepts behind the technology, the presentation of the nine Classes of Solutions, and information about their adoption, benefits and challenges when available.

According to the definition provided by the Artificial Intelligence Observatory of the Politecnico of Milan, particularly useful since combining a technological and business perspective, Artificial Intelligence can be defined as "*the branch of computer science that studies the development of hardware and software systems with specific capabilities typical of humans, able to autonomously pursue defined objectives, making decisions that previously were only made by humans*" (Perego, et al., 2019).
Particularly, the most relevant subfield of Artificial Intelligence, and preferred approach in it, is Machine Learning, addressing the question of how to create computer systems able to automatically improve through experience (Jordan & Michell, 2020). Thanks to Machine Learning techniques, computers are provided with the ability to automatically learn from experience without being explicitly programmed. (Liu, et al., 2017) Three different approaches can be used to train a machine and make it learn through experience in Machine Learning: Supervised Learning (SL), Unsupervised Learning (UL) and Reinforcement Learning (RL).
To conclude, Deep Learning is the most common class of techniques used in Machine Learning applications. Deep Learning algorithms generally consist of a hierarchical architecture composed of multiple layers, used to construct a progressively more abstract representation of data. These deep architectures, since composed of multiple layers, are generally referred as Deep Artificial Neural Networks. (Liu, et al., 2017) Two of the most common types of Deep Neural Networks, achieving outstanding results in recent years, are Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs). While CNNs have been a breakthrough in image and video processing and have been widely adopted by the Computer Vision community, RNNs have achieved many practical

successes in processing sequential data such as text and speech. (LeCun, Bengio, & Hinton, 2015)

Once provided an historical contextualization of these concepts and methods, the role that the technology can have in a business context has been considered: according to the Artificial Intelligence Observatory of the Politecnico of Milan, when researchers talk about Artificial Intelligence, they implicitly refer to a multiplicity of "Classes of Solutions". These Classes of Solutions have been explicitly defined by the Observatory to consider the possible applications of Artificial Intelligence in business, helping to understand the different areas and scope that companies want to achieve using this technology.
Each Class of Solutions has been so reviewed in the Literature, reporting information, when possible, about practical applications in business, adoption and diffusion, benefits, challenges and concerns.

Particularly, nine different Classes of Solutions, each furtherly organized into different Specifications, have been defined:

a. Intelligent Data Processing, including all those analytic solutions using Artificial Intelligence on structured and unstructured data to extract information for supporting the decision-making process. (Perego, et al., 2020) This is particularly relevant for companies, since due to the Big Data explosion they are more and more overwhelmed by data: therefore, Artificial Intelligence analytics algorithms are typically used to process them. 5 Specifications for this Class of Solutions can be identified: Forecasting (solutions to predict the value that one or more variables may assume in the future); Classification & Clustering (solutions to assign a category to each record from a predefined set, or to divide input data into clusters); Identification (solutions to identify anomalies or patterns within large datasets); Optimization (solutions to achieve the optimal value for a target variable, by modifying the variables of the system to achieve the optimum); Content/Design Creation (solutions analyzing data to create new contents or plan new services/products).

b. Natural Language Processing, including solutions to analyse and represents naturally occurring text at different levels of linguistic analysis, for a variety of tasks or applications. (Liddy, 2001) Natural Language Processing can be divided into two main components: Natural Language Understanding, focused on the understanding of human language, and Natural Language Generation, dealing with the creation of a human response. The main Specifications for the Natural Language Processing Class of Solutions are: Information Retrieval (solutions to automatically extract relevant information from a multitude of unstructured data); Information Filtering (solutions to filter and screen documents and unstructured text based on specific criteria); Text Generation (solutions leveraging Natural Language Generation for a variety of purposes); Language Modelling (solutions to correct errors regarding the use of syntax, lexicon or morphology).

c. Virtual Assistants/Chatbots, including software agents capable of performing actions and providing services to a human user, based on commands and requests received through written or spoken interaction. (Perego, et al., 2020) These are software applications conversing with a human using natural language text, or speech, to achieve a result. At present, the most familiar achievements in this Class of Solutions are voice-driven Virtual Assistants such as Siri, Alexa, Cortana and Google Assistant, but behind them thousands of text-based chatbots, operating in a variety of channels, exist, and are used to help with specific tasks in narrower domains. (Dale, 2016) While past Chatbots were simple rule-based solutions, nowadays they are typically endowed with Artificial Intelligence capabilities, so that companies are more and more interested in them, especially in the customer service area. (Nimavat & Champaneria, 2017)

d. Computer Vision, including solutions for image and video analysis, biometric recognition, and the general extraction of information from images or videos. Particularly, Computer Vision can be defined as the branch of Artificial Intelligence focused on the use of algorithms and optical sensors to simulate human visualization, enabling machines to understand and process visual data, such as images and videos, in the same way that humans do. (Wiley & Lucas, 2018) (Burger & Wheelock, 2015) The most common Computer Vision tasks are Image Classification, Object Detection, Image Retrieval, Face Recognition, Action and Activity Recognition, Human Pose Estimation and Semantic Segmentation, while 4 different Specifications exist: Image Analysis; Video Analysis; Biometric Recognition; Image & Video Editing.

e. iRPA (Intelligent Robotic Process Automation) refers to the execution of recurring tasks by a software integrated with Artificial Intelligence capabilities. (Perego, et al., 2019) Consequently, iRPA is able to achieve flexible and intelligent automation by combining traditional Robotic Process Automation (RPA), a business process automation technology allowing the automation of high-volume routines, and Artificial Intelligence. (Zhang C. , 2018) Particularly, an interpretative framework developed by the Artificial Intelligence Observatory of Politecnico of Milan defines three possible levels of integration between Artificial Intelligence and RPA: Programmed RPA, AI Assisted RPA, AI Driven RPA.

f. Recommendation, including solutions using information about the preferences, interests or decisions of a user, with the aim to deliver personalized recommendations at different points of the customer journey or the decision-making process. (Perego, et al., 2019) Recommendation Systems are an area of huge interest, since helping users to deal with information overload by suggesting products, services or contents tailored to their interests. Some well-known applications of these systems are products recommendations of Amazon.com and films recommendations for Netflix. (Adomavicius & Tuzhilin, 2005) (Burke, Felfernig, & Göker, 2011) (Jannach & Jugovac, 2019) The use of Artificial Intelligence has greatly contributed to advance the state-of-the-art in Recommendation technologies, in contrast with early systems based on heuristics. According to the used recommendation method, three main categories of Recommendation Systems can be defined: Content-Based Recommendation Systems; Collaborative Recommendation Systems; Hybrid Recommendation Systems. (Adomavicius & Tuzhilin, 2005) The main Specifications for

Recommendation are: Purchasing Recommendation (to suggest a user the purchase of products or services aligned with his interests); Content Recommendation (for the suggestion of contents); Online Advertising; Dynamic Pricing.

g. Autonomous Robot, including robots endowed with Artificial Intelligence capabilities able to move themselves or some of their parts and perform various actions without human intervention, gathering information from the surrounding environment and adapting to unforeseen or coded events. (Perego, et al., 2019) Solutions belonging to this class can find application in a multitude of contexts and are extremely different in terms of characteristics, functionalities, and aspects, ranging from humanoid robots, to collaborative robots, to in-store assistants. Autonomous Robots can be used both by corporations, to automate and speed up processes and creating competitive advantage, and by consumers, to receive support in their daily routine. (Hinds, Roberts, & Jones, 2011)

h. Autonomous Vehicle, including self-driving vehicles for transportation of people, animals or things, circulating on the road, intended for navigation or flight and capable of perceiving the external environment and identifying manoeuvres to adapt to it. (Perego, et al., 2019) Artificial Intelligence algorithms are widely applied in Autonomous Vehicles for a variety of purposes, including perception of a driving scene, identification of the appropriate navigation path, detection and classification of objects, behaviour arbitration and motion control. (Grigorescu, Trasnea, & Macesanu, 2019) (Bimbraw, 2015)
Currently, the most evident application of Autonomous Vehicles are self-driving cars, with the classification model of the Society of Automotive Engineers (SAE) defining six possible levels of autonomous driving. The main Specifications for Autonomous Vehicles are: Advanced Driving Assistance Systems, Autonomous system along a defined path, autonomous system along a non-defined path.

i. Intelligent Object, including objects able to perform actions and take decisions without human intervention, interacting with the environment through sensors and actuators and learning from actions of people interacting with them. (Perego, et al., 2019) Specifically, this Class of Solutions includes all those applications in which Artificial Intelligence algorithms run directly on the hardware of smart objects and IoT devices, avoiding the elaboration of data on cloud platforms. The main reason for having a dedicated Class of Solution is that nowadays the IoT represents a very hot topic, due to the huge number of devices with sensing capabilities gathering or generating sensory data for a variety of aims and applications. Particularly, a new research area has emerged, focused on bringing analytics to resource-constrained hardware to enable real-time analytics. (Mohammadi, Al-Fuqaha, Sorour, & Guizani, 2018)

In Chapter 2, research questions and methodologies followed to answer them are presented. While an extensive body of literature focuses on Artificial Intelligence techniques, information about its adoption in enterprises is limited and mainly referred to past years. However, Artificial Intelligence is a rapidly growing technology, and its enterprise adoption is rapidly evolving, implying the necessity to provide a constantly updated picture of its evolution. Therefore, the primary contribution of this thesis consists in answering the following research question:

RQ1: What is the present state about the international adoption of Artificial Intelligence solutions?

Moreover, the broad applicability of this technology is currently enabling companies in different sectors to adopt it at advantage of their specific business needs. In this regard, the secondary aim of this thesis work is to provide a focus on how Artificial Intelligence is currently being adopted in some specific industries, with a sectoral analysis to answer the following research question:

RQ2: How Artificial Intelligence adoption changes from industry to industry and what is the potential contribution for the sector?

To answer these research questions several research methodologies have been adopted. Firstly, the existing Artificial Intelligence literature has been reviewed, allowing to achieve an appropriate level of comprehension of the theoretical concepts behind the technology and to gain knowledge about the nine Classes of Solutions introduced by Observatory, at the basis of Artificial Intelligence's adoption at a business level. Then, an analysis of secondary sources on the web has been carried out, to identify active Artificial Intelligence initiatives of firms, following a well-established methodological scheme. 1089 Artificial Intelligence projects of 165 companies analysed have been identified by applying this methodology. The identified projects have been used for a census of projects, feeding a database specifically created for the purpose of the research. Each of the initiatives has been categorized according to specific criteria. The database has been later used as a basis for analysis, providing answers about the identified research questions. To conclude, a qualitative framework has been specifically developed to support sectoral analysis, helping to answer *RQ2.* The model has been used to support both the analysis of how Artificial Intelligence is being applied within a certain industry and a structured and orderly presentation of results.

To perform this analysis, a sample including the first 235 public companies rated in the 2020 Forbes Global 2000 list has been considered, identifying their active Artificial Intelligence projects. A cap of maximum 20 companies analysed per industry has been placed, to obtain a homogeneous overview of the sectors and to avoid biases in the database. By excluding vendors of Artificial Intelligence solutions and companies exceeding the maximum number of 20 firms per industry, 169 companies have been analyzed.

Chapter 3 is intended to answer the research question *RQ1: What is the present state of the international adoption of Artificial Intelligence solutions?* Results coming from the analysis of the database are presented and discussed, providing an overview about the current adoption of Artificial Intelligence at a global level. 235 companies have been considered and, following the research methodology, 169 of them have been actually analyzed. No Artificial Intelligence initiatives have been found for 13 of them. Therefore, the created database contains 156 companies and 1089 related projects.

A high-level analysis of the results shows how practical applications of Artificial Intelligence technology have emerged only recently, but the number of Artificial Intelligence projects launched in the last few years (2018; 2019; 2020) is impressive, and the technology is rapidly diffusing in a business context. Moreover, it is spreading in all the sectors, with the most various industries currently adopting, or experimenting with, solutions based on Artificial Intelligence. At the same time, it is also being applied in the most various contexts, both within the boundaries of firms and as an enabler of products and services sold to B2B, B2C and B2G markets. These results highlight the great applicability of the technology to the most different scenarios.

Furthermore, the research shows that a large number of initiatives launched in recent years have now become Operative (57,9% of the projects) or are currently in Implementation at a large scale (8,4%), suggesting positive results from early testing with the solutions, so that they are being implemented at scale. Meanwhile, the number of Pilot projects (21,9%) and Project Proposals (11,8%) suggests that many companies are starting only now to consider this technology, while other ones are exploring the opportunities offered by this technology by experimenting with new applications.

Considering the diffusion of the different Classes of Solutions, results are shown in Figure 1. As shown, solutions based on Intelligent Data Processing are currently the most diffused ones (37,2% of the initiatives), underlying the relevance of Artificial Intelligence algorithms in making sense of the impressive amount of data companies are overwhelmed with. Computer Vision, Chatbot/Virtual Assistant and Natural Language Processing solutions are spreading in companies too, with respectively the 17,7%, 14,0% and 11,3% of the initiatives. The percentage of Virtual Assistant/Chatbot solutions should not be interpreted as a signal of intermediate diffusion of this class: by analyzing the collected projects, it results that the 67,3% of the companies have at least one project in the Virtual Assistant/Chatbot Class of Solutions, suggesting that Virtual Assistants and Chatbots have a high level of adoption and have become a widespread reality within businesses. Recommendation represents just the 5,1% of the projects, also due to an applicability of Recommendation Systems particularly valuable for just some industries.

Autonomous Vehicles and Autonomous Robots represent the 5,7% and the 4,0% of the initiatives. Their limited diffusion is mainly because these solutions are still at an embryonic stage and further advancements are needed, because of the complex Artificial Intelligence capabilities required for their purposes. Therefore, as shown in Figure 2, these two Classes of Solutions are the ones with the lowest maturity level, with most of the projects at a Pilot stage or still under development (Project Proposals). On the contrary, it is shown how Recommendation Systems and Virtual Assistant/Chatbot are the most mature classes and their adoption is being consolidated in recent years, with most of the projects already Operative or in Implementation.
iRPA solutions are just the 3,8% of initiatives, suggesting that many companies are still adopting solutions based on traditional Programmed RPA. To conclude, the low percentage of Intelligent Object (1,3%) projects points out how Artificial Intelligence algorithms are still typically running on the cloud, with a limited number of devices directly elaborating information on their hardware.

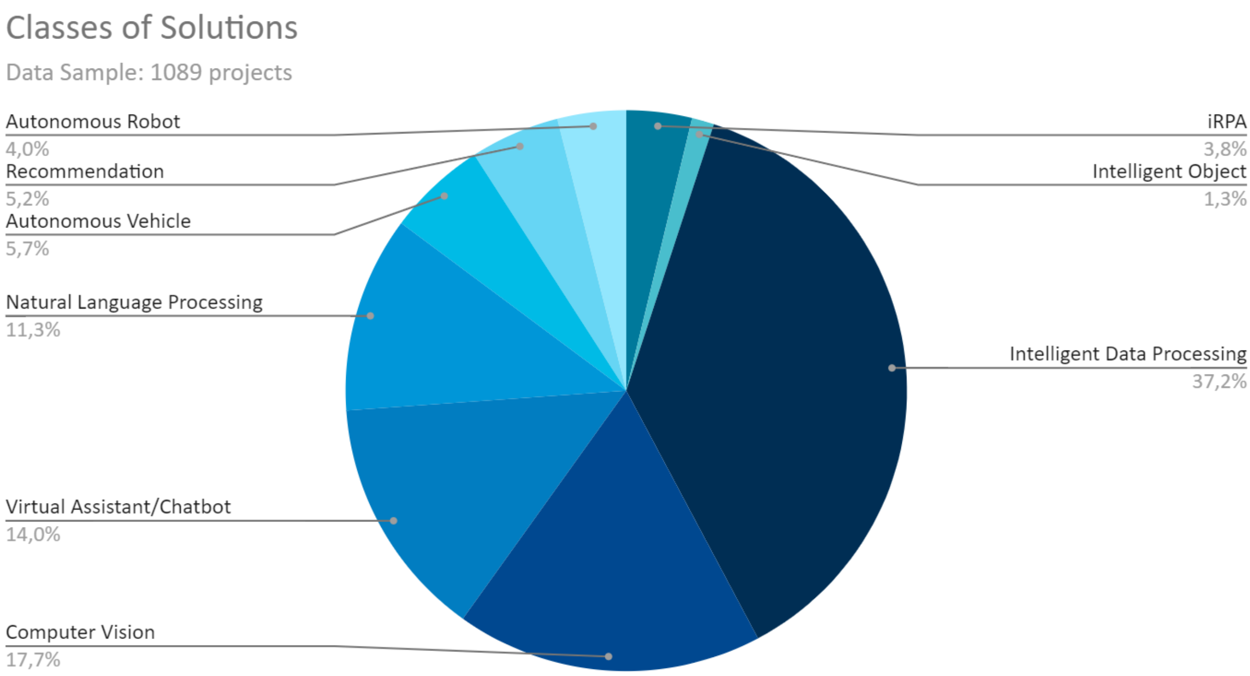

Figure 1: Distribution of Classes of Solutions

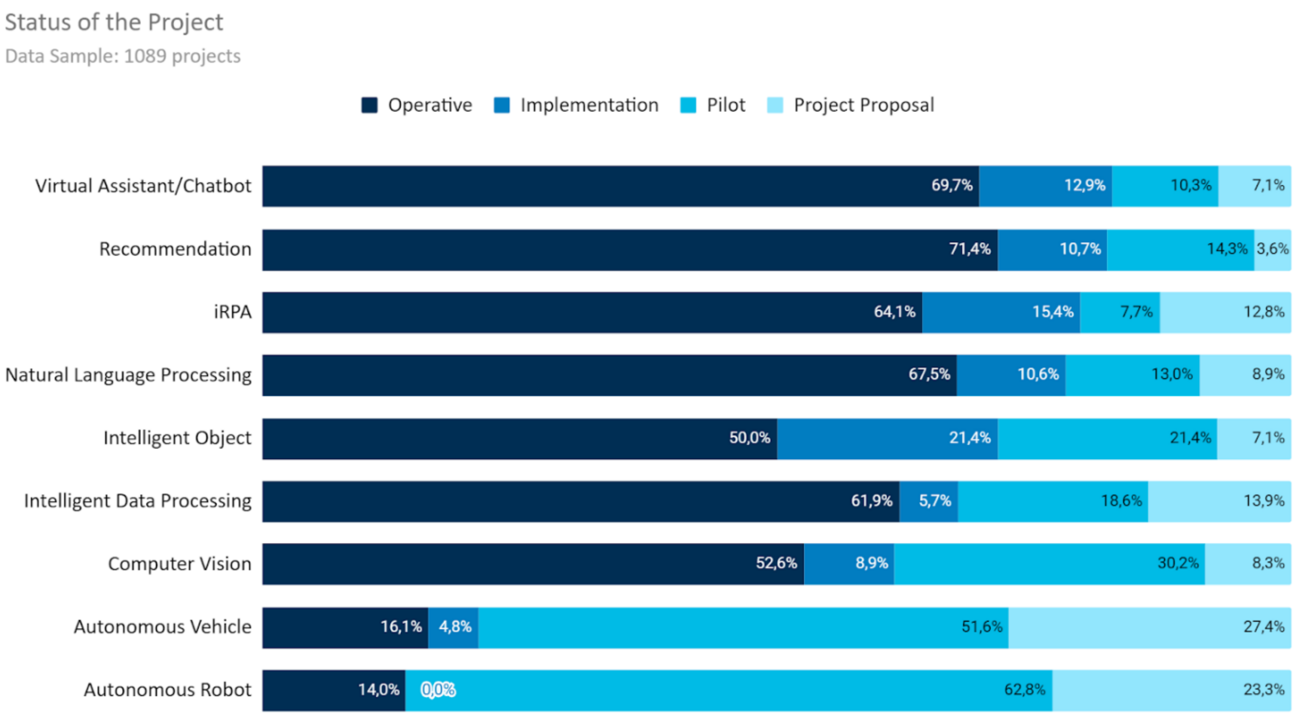

Figure 2: Status of the Project per Class of Solutions

A focus on each Class of Solutions and the corresponding Specifications has been afterward provided, presenting the most interested sectors, applications, and relevant business cases. The main results for each Class of Solutions are here summarized:

a. In Intelligent Data Processing, the most diffused Specifications are Identification, Forecasting and Optimization. Identification is typically at the basis of Artificial Intelligence solutions for Cybersecurity and Fraud Detection, with the Banking & Finance and Insurance industries widely adopting these solutions. Another relevant application field is Maintenance, with the use of Artificial Intelligence algorithms based on anomaly detection to prevent failures in assets and equipment. The Energy, Resource & Utility industry has launched several initiatives for this purpose. To conclude, algorithms for pattern identification have found widespread adoption in the Pharmaceutical industry for new drug development, supporting the analysis of data to develop new drugs with notable savings in time, cost, and risk. When it comes to Forecasting, Artificial Intelligence is being applied to improve the accuracy in forecasting demand for products and services, especially in the Retail and Telco sectors. Meanwhile, several services enabled by Artificial Intelligence are emerging, such as investment advisory services in Banking & Finance, leveraging Artificial Intelligence to predict the movements of the markets. To conclude, Optimization algorithms are at the basis of several services, and are also increasingly used for external logistics and production planning. In External Logistics, they are applied to optimize the transportation of goods, decreasing shipping costs and wasted time for movements. This kind of solutions is finding widespread adoption in a variety of industries, with the Transportation & Logistics sector particularly active. On the other side, optimization algorithms are also being used for production planning in manufacturing plants, with solutions to optimize processes through the adjustment of production parameters and production/workforce schedule optimization.

b. In Natural Language Processing, Information Retrieval is the most diffused Specification, followed by Information Filtering and Text Generator. In Information Retrieval, Artificial Intelligence is typically used for sentiment analysis, enabling businesses to gain information about the perception of the brand by customers and their satisfaction along the customer journey. Meanwhile, it is also finding application in various sectors to automate the analysis and extraction of information from unstructured text. Particularly, insurance companies are applying Artificial Intelligence to analyze text for policy reviews and to extract information from insurance documents of interest. Information Filtering is being internally used by companies to support employees in the execution of their tasks, with smart screening solutions reducing the time spent in searching for relevant documents. These solutions are also diffused in the Pharma Industry, to help researchers to quickly go through thousands of research papers and articles when searching for relevant publications. Other relevant applications regard Operations Automations, with the adoption of solutions to automate the categorization of mails, documents, and other textual resources. To conclude, Text Generator solutions are at the basis of several services addressed to a B2C market, such as Google Translate or Gmail with automatic email reply. Meanwhile, it is increasingly applied for marketing optimization purposes, enabling the creation of more effective marketing messages.

c. Virtual Assistant/Chatbot solutions are characterized by a high maturity rate and, despite constituting just the 14,2% of initiatives, the 67,3% of the analyzed companies have at least one active project in this Class of Solutions. Therefore, Virtual Assistants and, especially, Chatbots have become a widespread and consolidated reality in business. Chatbots with Artificial Intelligence capabilities are typically used to assist and support customers, handling their queries and providing them useful information in a variety of channels. They have been introduced in almost all the sectors, but most of all the Banking & Finance and Insurance industries have adopted these solutions. Despite the focus on customer service, internal applications for Chatbots are also progressively taking hold, exploiting them to support employees in the execution of their tasks through internal chat systems, supporting the early stages of the recruiting process, and for HR Management. On the other side, Virtual Assistants, or Voicebots, are typically applied in services, with solutions that can be grouped in three clusters: the popular and well-known Virtual Assistants, such as Siri or Google Assistant, available in smartphones or connected with devices, able to understand spoken queries and to accordingly answer or take actions; virtual assistants that automobile companies are starting to integrate into the operating system of their vehicles, as additional onboard service; voice-based services on the Virtual Assistants of the first group, that companies in the most various sectors are starting to introduce. To conclude, benefits resulting from the adoption of Chatbots and Virtual Assistant solutions to support the customers or employees have been identified, considering both the customer and company perspective.

d. In Computer Vision, the most diffused solutions regard Image Analysis, with a variety of application fields. For instance, solutions leveraging Artificial Intelligence for Image Analysis are at the basis of several services, such as image-reading software systems for early disease detection that companies in various sectors are developing and providing to medical institutions. Meanwhile, companies in the Manufacturing and Automotive sectors are increasingly using inspection systems leveraging Computer Vision in quality control: recognizing manufacturing defects and anomalies along production and assembly lines, they enable the automatic inspection of products and parts and the optimization of quality controls. To conclude, solutions based on Image Analysis are used to automate several industry-specific tasks. On the other side, Biometric Recognition is finding diffusion especially in Asia, with Chinese banks that have extensively adopted it for authentication at ATMs or in mobile phone banking applications. Video Analysis's solutions are being typically adopted in Safety and Security. In both cases, Artificial Intelligence is used for the analysis of videos, identifying unusual and abnormal behaviours, and triggering alerts. Safety solutions are typically applied in a business context, with companies in the Energy, Resources & Utility and Construction sectors using them to monitor safety conditions in plants and construction sites, detecting and reporting eventual safety risks. On the other side, Security solutions are finding adoption also in B2C and B2G scenarios. To conclude, Image & Video Editing typically regards mobile applications or specific functionalities of the services developed by Big Tech companies, allowing users to edit photos and videos. In conclusion, some critical issues specific for Computer Vision projects are presented, including: huge amounts of training data needed; images and videos affected by noise or ambient factors; reliability of solutions; privacy concerns; ethics concerns.

e. As regards the iRPA Class of Solutions, the limited diffusion of initiatives suggests that most of the companies are still using traditional Programmed RPA, and the shift to Artificial Intelligence Supported, or Driven, RPA is still far from being reality. iRPA solutions are typically applied for the intelligent automation of repetitive tasks and processes in different industries, but two sectors are showing particular interest in the adoption of iRPA solutions: Banking & Finance, and Insurance. Indeed, being some of the most data-intensive sectors, they can particularly benefit from the adoption of these solutions.

f. When it comes to Recommendation, it includes just the 5,1% of the analyzed initiatives. However, its applicability generally limited to just some industries should be considered: if considering the Retail and Media & Entertainment industries, in which Recommendation Systems have a great applicability, the 75% of Retail companies and the 71,5% of Media & Entertainment companies have at least one active project in this Class of Solution, underlying a high diffusion in those sectors that can benefits the most from their implementation. Moreover, this is the Class of Solutions with the highest maturity, along with Virtual Assistant/Chatbot. The most diffused Specifications are Purchasing Recommendation, with Artificial Intelligence increasingly applied by retailers to recommend users with products and services they are interested with, and Content Recommendation, referring to the use by companies in the Media & Entertainment industry to suggest users of their services tailored contents.

g. As regards the Autonomous Robot Class of Solutions, its limited diffusion can be attributed to its low maturity level. Consequently, further advancements are needed before a wider adoption of this class. However, the opportunities offered by Autonomous Robots are tremendous, and they can find application in a multitude of industries. Several industries, above all the Automotive and Manufacturing sectors, are piloting Autonomous Robots to automate operations in production plants. Robots can be used to autonomously complete entire processes, to substitute the workers in just some tasks, rather than to collaborate with them in production. Meanwhile, Autonomous Robots are starting to generate interest also in the Banking & Finance and Energy, Resources & Utility sectors, with respectively the exploration of humanoid robots to interact with customers in branches and the adoption of robots for inspections. To conclude, in recent years the interest for Autonomous Robots in a domestic environment has increased, with several solutions for a variety of tasks addressing consumers. However, most of these initiatives are still at an embryonic stage.

h. Autonomous Vehicles, similarly, are characterized by a limited diffusion and maturity, with further advancements required to reach a wider diffusion. The

Automotive industry is, among all, the one more active in launching initiatives for Autonomous Vehicles. Alongside self-driving cars and ADAS as additional on-board services, internal applications for internal logistics and external logistics are emerging, underlying how Autonomous Vehicles can also play a role within the boundaries of the firms. Particularly, in external logistics most of the projects regard the exploration of self-driving trucks to automate the transportation of goods and packages delivery. Above all, the shipping companies operating in the Transportation & Logistics sector are interested in these solutions. Autonomous Systems along a not defined path represent the majority of initiatives, because of the several projects that Automotive companies are launching to develop the next generation of self-driving cars. Advanced Driving Assistance Systems follow, with Advanced Driving Assistance Systems integrated in self-driving cars or trucks. To conclude, Autonomous Systems along a defined path are typically finding adoption for Internal Logistics in a variety of industries. In this case, Autonomous Vehicles typically consist of automated material handling systems, such as forklifts and pallet drones, operating in predetermined areas such as warehouses, production centers, and distribution centers to automate movements of goods.

i. Intelligent Objects, with just 14 projects out of 1089, represent the 1,2% of the analyzed initiatives, suggesting that Artificial Intelligence algorithms rarely run directly on the devices' hardware, because of limitations in computing, memory and storage capacity. Therefore, data are typically still elaborated on dedicated cloud platforms and the number of projects in this Class of Solutions is limited. Anyway, it is expected that, once these limitations will be overcome, the amount of applications with Deep Learning algorithms directly hosted on the device will rise.

To conclude, an overview about Critical Issues encountered in Artificial Intelligence projects, relevant challenges, and rising concerns that could limit a future wider adoption of the technology are provided, based on the analysis of results. The identified Critical Issues are: Huge Amount of Training Data; Quality of Training Data; Artificial Intelligence bias; The Black Box problem of Deep Neural Networks; Reliability of Solutions; Unfair Use

of Artificial Intelligence; Ethical Concerns; Privacy Concerns; Implications on Job; Necessity of Regulation.
Therefore, a further diffusion of Artificial Intelligence at a global level requires both overcoming a series of practical challenges and further developments in the technology. However, the benefits offered by these solutions are so evident that no company can ignore this and stay out of the ongoing Artificial Intelligence revolution. This becomes even more relevant by considering that, in some fields, companies have just started to explore the opportunities offered by this technology. This is the reason why Artificial Intelligence is expected to have a bright future, with an increasing adoption of solutions also in years to come.

While in Chapter 3, the focus is on the international adoption of Artificial Intelligence and the analysis of each Class of Solutions, Chapter 4 includes the second contribution of this thesis, consisting in investigating how Artificial Intelligence is applied by organizations operating in some specific industries, and how it is transforming their business. In this sense, the perspective shifts from the point of view of the Classes of Solutions to the one of the Industry, with the aim to answer the following research question:

RQ2: How Artificial Intelligence adoption changes from industry to industry and what is the potential contribution for the sector?

Particularly, the analysis has been carried out for companies operating in the Food & Beverage, Manufacturing, Retail, Banking & Finance industries. The investigation has been carried out by analysing the created database and focusing on companies operating in the industries under analysis. For this purpose, a qualitative framework has also been specially developed, to support the analysis of how Artificial Intelligence is being applied within a certain industry and as a structured and orderly presentation of results.

The developed framework is shown in Figure 3: as shown, it is organized in 3 different Levels, in which the possible Artificial Intelligence's application fields have been grouped based on their proximity with the specific customer of the industry. Particularly, the model can be represented with a pyramidal scheme that places on top the customer.

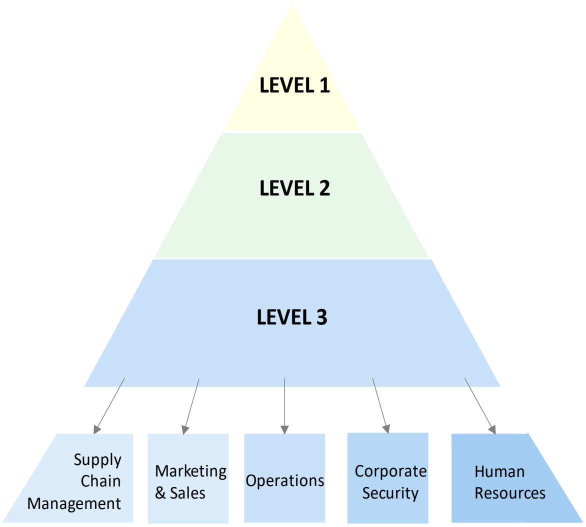

Level 1	AI - Enabled Products & Services and Customer Interaction	Service Product Customer Interaction
Level 2	Marketing, Sales and NPD	Customer Profiling Demand Forecasting Marketing Optimization New Product Development Price Recommendation Sentiment Analysis
Livel 3	Supply Chain Management	External Logistics Internal Logistics Supply Chain Improvement
	Business & Financial Control	Regulatory Compliance Risk Management
	Operations	Employees Support Engineering Asset Management Maintenance Operations Automation Production Planning Quality Control
	Corporate Security	Cybersecurity/Fraud Detection Safety Security
	Human Resources	HR Management Recruiting

Figure 3: Qualitative Framework

Particularly, the three Levels have been organized according to the following definitions, establishing criteria to associate possible Applications for Artificial Intelligence solutions with Levels.

Level 1 - AI Enabled Products & Services

At this level, Artificial Intelligence is used as an enabler of the Product and Services that companies offer to their clients (B2B, B2C or B2G market), or is at the basis of interactions between the clients and the company.

Level 2 – AI in Customer-Oriented Processes

At this level, Artificial intelligence is used internally by companies. Data elaborated by Artificial Intelligence algorithms are customer data, or the customer has a direct visibility on the outcomes of Artificial Intelligence solutions.

Level 3 – AI in Enterprise-Oriented Processes

At this level, Artificial intelligence is still used internally by companies, but generally data elaborated by Artificial Intelligence algorithms are not customer data, or the customer has not a direct visibility on the outcomes of Artificial Intelligence solutions (even if he could benefit from their implementation). Since Level 3 includes a variety of Applications, it has been furtherly organized into sub-Levels to enable a more structured analysis, grouping Applications based on their field of application. The following sub-Levels, each with its own Applications, have been defined: Supply Chain Management; Business and Financial Control; Operations; Corporate Security; Human Resources.

The developed framework has then been applied to each of the industries previously mentioned. The model has been firstly applied at a high-level to understand the distribution of Artificial Intelligence projects in the different levels and comment on it. Then, the analysis goes into detail to each Level, describing which are the most relevant sub-Levels and Applications and providing, when needed, relevant business cases.

Particularly, the application of the qualitative framework to the Food & Beverage industry highlights how Artificial Intelligence has a limited applicability and diffusion in the industry at Level 1 (AI Enabled Products & Services), with just 13,5% of initiatives. However, Artificial Intelligence finds space in interactions with B2B and B2C customers, with the use of Chatbots and Voicebots to handle customer requests and provide a better customer experience. At Level 2 (AI in Customer-Oriented Processes: 27,5% of the projects), Artificial Intelligence has found a good level of diffusion, with Food & Beverage companies using the technology to better know their customers through solutions for Customer Profiling and Sentiment Analysis. The results are then used to bring new food and beverage products to the market (supporting New Product Development) or for better marketing and promotion of products (supporting Marketing Optimization). These kinds of solutions, particularly, have found widespread and consolidated adoption in the Food & Beverage industry in recent years.

Level 3 (AI in Enterprise-Oriented Processes) is the one with the highest number of initiatives (58%), underlying the great opportunity for the sector to use Artificial Intelligence to improve internal activities. On one side, it is more and more applied at the factory level, to support maintenance, production planning and the quality control of food and beverage products. On the other side, several initiatives in which Food & Beverage companies are collaborating with their Supply Chain Partners (growers and retailers) at advantage of the entire Food & Beverage Supply Chain have been launched and are already operative, relying on Intelligent Data Processing and Computer Vision solutions.

To conclude, several initiatives for External and Internal Logistics have been launched, despite all these projects being just pilot projects or project proposals. Indeed, they mainly refer to the adoption of Autonomous Vehicles and Autonomous Robots for goods transportation and warehouse operations.

In the Manufacturing industry, by applying the qualitative framework, it emerges how most of the projects regard Level 1(AI Enabled Products & Services) and Level 3 (AI in Enterprise-Oriented Processes), with respectively the 49,3% and 31,5% of the initiatives in the industry.
On one side, the diffusion of solutions at Level 1 suggests that the Manufacturing Industry is widely using Artificial Intelligence as enabler of the functionalities of Products and Services sold by the firms to a variety of B2B markets and to a B2C market.
On the other side, most of the projects at Level 3 are related to Operations, suggesting the key role that Artificial Intelligence can play, and is playing, in the digital transformation of manufacturing plants. Several industrial applications of this technology are possible, with solutions mainly based on Intelligent Data Processing and Computer Vision.
In Quality Control several solutions based on Computer Vision are already operative in the industry, automating the inspection of products and components and increasing the accuracy of controls if compared with traditional human inspection. In Production Planning, Artificial Intelligence is applied to optimize production and improve the process efficiency, with Intelligent Data Processing solutions to autonomously adjust production parameters, cutting the time needed to set up optimal parameters manually, and to optimize production scheduling. For Operations Automation, companies are exploring solutions based on Autonomous Robots to automate operations in production plants, with projects still under development or at a pilot stage. Artificial Intelligence is relevant in Maintenance too: manufacturing companies are typically using it to analyze massive amounts of sensor data for anomaly detection, improving the productivity of factories and production plants by alerting in case of possible failures. Therefore, adoption of Artificial Intelligence solutions in the Manufacturing industry is particularly significant for Operations.

To conclude, despite projects mainly regard Level 1 and 3 of analysis, initiatives have been found at Level 2 (AI in Customer-Oriented Processes) too, supporting Marketing Optimization, Demand Forecasting and New Product Development.

In the Retail Industry, by applying the qualitative framework, it results that retailers are adopting Artificial Intelligence solutions at all the levels, with quite similar diffusion rates. The high diffusion of initiatives at Level 1(AI Enabled Products & Services) and 2 (AI in Customer-Oriented Processes) can be explained by the fact that Retailers are in direct contact with the ultimate consumers, representing the last step of the Supply Chain before end users. At Level 1 (28,3% of the projects in the industry), solutions for Customer Interaction are finding widespread adoption. These mainly refer to the use of customer service Chatbots and Virtual Assistants in the online channels of retailers, solutions that are even more relevant in the Retail sector because of the direct contact with consumers.
At Level 2 (33,3% of the projects) data about the customers and their behaviours are widely used for Marketing Optimization, important for retailers to attract customers and entice them to purchase products. Solutions typically regards the use of Recommendation Systems, that all the companies have introduced, or are introducing, in their online stores. Using Artificial Intelligence, these systems are used to show in online stores products based on the past behaviour and preferences of users, enabling personalization of the shopping experience.
To conclude, Level 3 (AI in Enterprise-Oriented Processes) is the one with the highest diffusion of projects, underlying how Artificial Intelligence is also being internally applied by Retailers for Internal and External Logistics, and Operations Automation in physical stores. The use of Artificial Intelligence in External Logistics mainly refers to the use of Intelligent Data Processing algorithms for optimization of the shipping process and Autonomous Vehicles. The latter are Pilots or just Project Proposals, and typically require collaboration with third parties.: terrestrial or aerial self-driving vehicles, such as cargo vans or drones, can be used for transportation between factories, warehouses and distribution centers, or deliveries to final consumers. In Internal Logistics, different solutions are under exploration by retailers, ranging from Intelligent Data Processing solutions to sequence the picking and packaging of items to Autonomous Robots to pack

items. When it comes to Operations Automation, several solutions, typically in a pilot stage, are currently being tested by retailers to automate basic activities and tasks in physical stores, often taking advantage of Computer Vision to automate the replenishment of shelves, and for other tasks.

In conclusion, the Banking & Finance Industry is presented. When considering the status of the projects in this industry, it emerges how the majority of solutions are already Operative (73% of the projects in the industry) or in Implementation (8,7%), while the remaining 18,3% of solutions are Pilot projects (11,3%) and Project Proposals (7%). This underlines how, generally speaking, the vast majority of Artificial Intelligence solutions in the Banking & Finance sector have a high maturity, suggesting that solutions in this industry are consolidated and have found widespread adoption in the recent years. By applying the developed qualitative framework, it emerges how most of the initiatives in the Banking and Finance Industry are related to Level 1 (AI Enabled Products & Services) and 3 (AI in Enterprise-Oriented Processes), with respectively the 41,7% and 47,8% of the projects in the industry.

At Level 1, Artificial Intelligence is at the basis of the Services that banks offer to their clients, or it is used for Customer Interaction. Solutions for Customer Interaction mainly refers to the use of Chatbots powered by Artificial Intelligence for customer service, used to provide information to customers, support them in managing and reviewing their bank accounts, fulfil requests, make recommendations and for several other purposes. This technology has become a standard in the industry. Meanwhile, a second group of solutions related to customer interaction regards the use of Biometric Recognition, that has found widespread adoption in recent years in China, with banks extensively adopting Artificial Intelligence to allow their clients authentication at ATMs or in mobile phone banking applications. As regards Services enabled by Artificial Intelligence, in the Banking & Finance industry they typically are investment advisory services, or robo-advisor services, using Artificial Intelligence to offer investment recommendations to customers. These kinds of solutions, as Chatbots for Customer Interaction, have spread in the industry in recent years.

On the other side, at Level 3 banks are using this technology especially for Automations Operations and Cybersecurity/Fraud Detection. Operations Automations includes a variety of solutions to automate processes or single tasks, ranging from Intelligent Data Processing and Natural Language Processing, to iRPA and Autonomous Robots. Particularly, iRPA systems are largely diffused in the industry to automate repetitive tasks. The adoption of these systems is extremely beneficial for the Banking & Finance industry, being generally considered as one of the most data intensive sectors. On the other side, Cybersecurity and Fraud Detection are vitally important in the Banking & Finance sector, to protect customer assets and keep their trust and to prevent financial losses. Particularly, Artificial Intelligence technology is widely used to support these activities, with the use of Intelligent Data Processing solutions for anomaly detection.

To conclude, Chapter 5 recap the identified research questions and summarize the answers provided. Moreover, limitations of the research and future directions of research are discussed.

The main limitations of the research are the following: first, the sample of companies present in the database is limited, not allowing a comparison between the investments of different industries. Second, the selection of companies is based on the Forbes Global 2000 list and the methodology described in Chapter 2. Consequently, the perspective considered is only the one of the largest public companies in the world, that have been the first to start investing in Artificial Intelligence. To conclude, the use of secondary sources for the creation of the database may have limited the visibility of certain projects and Classes of Solutions. For instance, the intercepted projects are just the ones made publicly available, and these typically are just the ones at the most advanced stages.

Suggestions for future research include continuing to map Artificial Intelligence initiatives also in future years, to have an always updated picture about its adoption and the possibility to analyze its evolution along years. Then, it could be interesting to introduce a higher level of detail for some applications, to provide a more detailed overview. To conclude, the developed qualitative framework could be adopted also in future research, to analyze different sectors or the same ones with larger sample volumes, resulting in a better picture of industries.

Introduction

Artificial Intelligence, a branch of computer science concerned with the creation of machines capable of performing tasks that typically require human intelligence, is permeating several aspects of modern society. From email filtering, to product recommendations in e-commerce platforms, to retrieval of relevant results in web search engines, Artificial Intelligence is rapidly expanding.

Despite the notion of General Artificial Intelligence, or "Strong" Artificial Intelligence, referring to intelligent systems able to handle multiple tasks like humans can do, remains at present just a theoretical concept and a long term objective for researchers, thanks to recent advancements in technology, computational power, Machine Learning techniques and availability of online data, progress has been made in the so-called Narrow Artificial Intelligence, or "Weak" Artificial Intelligence. Particularly, at present every existing intelligent machine is a Weak Artificial Intelligence.
This notion refers to intelligent systems able to handle just a single or limited tasks, such as only translating text for Google Translate, only driving a car for self-driving vehicles, only handling user queries and accordingly making actions or answering for Siri.

Thanks to recent progress, practical applications of Narrow Artificial Intelligence are becoming more and more common, allowing the technology to move from being just a research field to be a practical technology, with several applications possible. Solutions based on Artificial Intelligence range from music recommendations, to biometric recognition, to driverless vehicles. Particularly, Narrow Artificial Intelligence has also shown a wide business applicability, increasingly generating interest from public and private companies operating all over the world. Despite Artificial Intelligence being still early in adoption, it is rapidly expanding in a business context, and its state of adoption by enterprises is continuously evolving.

This thesis is so intended to investigate the current state of international adoption of Artificial Intelligence solutions by enterprises, providing an overview about the diffusion of the technology in a business context. Moreover, given that Artificial Intelligence is being used by different sectors at advantage of their specific business needs, the thesis investigates what is the current state of adoption of the technology, as well as its potential contribution, in four different industries: Food & Beverage, Manufacturing, Retail, Banking & Finance.

In Chapter 1, the existing literature is reviewed, describing the theoretical concepts behind the technology and presenting nine Classes of Solutions in Artificial Intelligence proposed by the Artificial Intelligent Observatory of Politecnico of Milan. In Chapter 2, research questions and methodologies followed to answer them are introduced. Chapter 3 discusses the results of the research, providing information about the enterprise adoption of Artificial Intelligence at international level. To conclude, Chapter 4 focuses on the state of adoption of the technology in the four industries under analysis. The concluding Chapter 5 summarizes the results, underlines limitations in the research and suggests directions for future research.

1. Literature Review

The aim of this Chapter 1 is to review the existing literature about Artificial Intelligence. The main theoretical concepts behind the technology and nine Classes of Solutions in Artificial Intelligence, introduced by the Artificial Intelligence Observatory of Politecnico of Milan to consider the applications of the technology in business, are presented.

1.1 Artificial Intelligence

Artificial Intelligence has been a field of research for a long time, with the birth of the discipline dating back to 1950, with Alan Turing's seminal paper "Computing Machinery and Intelligence". As a consequence, several definitions of Artificial Intelligence have been proposed over time by different authors; these ones consider different viewpoints and have different focuses, also depending on the historical period and the field of reference. The definition here adopted is the one proposed by the Artificial Intelligence Observatory of the Politecnico of Milan, particularly useful since combining a technological perspective and a business one.

Artificial Intelligence can be defined as "*the branch of computer science that studies the development of hardware and software systems with specific capabilities typical of humans, able to autonomously pursue defined objectives, making decisions that previously were only made by humans*" (Perego, et al., 2019)

By "specific capabilities typical of humans", the definition is referring to the human-like capabilities of Artificial Intelligence systems, such as processing natural language or images, learning and adapting, reasoning and planning, interacting with the environment.

Deeply grounded in the broader area of Computer Science, Artificial Intelligence includes systems, methods and tools able to perform complicated and intelligent computations similar to the ones routinely performed by the human brain, trying to simulate the way it acquires knowledge and reason to solve problems. (Agatonovic-Kustrin & Beresford, 1999)

1.2 Machine Learning

At the core of Artificial Intelligence and data science is the notion of Machine Learning. Machine Learning is a discipline addressing the question of how to construct computer systems able to automatically improve through experience. (Jordan & Michell, 2020).
Thanks to Machine Learning techniques, computers are provided with the ability to automatically learn without being explicitly programmed, learning from experience and making data-driven predictions or decisions (Liu, et al., 2017).

The study of Machine Learning becomes even more relevant by considering that nowadays, for a wide range of applications, it has been recognized that training a system by showing it examples of a desired input-output behaviour is easier and much more feasible than manually programming it, anticipating the desired answers for all the possible inputs. (Jordan & Michell, 2020)

Particularly, a "learning problem" can be defined as the problem of improving a certain measure of performance for the execution of a task through a training experience. For instance, consider learning to identify credit card frauds. In this case, the task consists in assigning an output label "fraud" or "not fraud" to credit card transactions in input. The performance metric to improve could be the accuracy of this fraud detection classifier, while the training experience could be a sample of historical transactions, already classified with the label "fraud" or "not fraud". We can therefore conceptually consider a Machine Learning algorithm as a search in a space of candidate programs, driven by training experience, to find the program able to optimize the performance metric.
To cover the wide diversity of data, problems, and applications, several Machine Learning algorithms and architectures have been developed, differing in the way candidate programs are represented or in the way the search through the space of possible candidate programs takes place, and providing different trade-offs between computational complexity, amount of needed data, performances. (Jordan & Michell, 2020)

Within Artificial Intelligence, Machine Learning represents one of the most rapidly growing technical fields and has emerged as the preferred method for developing software for Natural Language Processing, Computer Vision, Pattern Recognition and other areas. It has had a significant impact on our everyday lives over the past decades, with successful applications including Optical Character Recognition, web search, self-driving systems. (Liu, et al., 2017)

As a discipline, Machine Learning can be considered at the junction between computer science, statistics, and many other areas of research related to making decisions under uncertainty or automatically improving over time with experience. Other fields of study related to Machine Learning are, among the others, the psychological study of the way human beings learn, cognitive science, neuroscience, education practices. Despite an increasing interaction with these fields in the last decade, we are just starting to exploit the synergies among them and the experimental techniques they are using to study systems able to learn and improve through experience. (Jordan & Michell, 2020)

When dealing with Machine Learning, three different approaches through which a machine can be trained and can learn through experience exist: Supervised Learning (SL), Unsupervised Learning (UL) and Reinforcement Learning (RL).

1.2.1 Supervised Learning

At present, Supervised Learning is the most common method to train a machine. (LeCun, Bengio, & Hinton, 2015)

The aim of this method is to predict one or more target values from one or more variables in input, producing an output y, or a probability distribution for y, given an input x.
Predictions are done via a learned mapping f(x); many different methods to map f exist, such as Bayesian Classifiers, Decision Trees, Decision Forests, Support Vector Machines, Artificial Neural Networks. (Jordan & Michell, 2020)

In particular, in Supervised Learning the machine learns with a labeled training dataset, meaning that input data are already labeled with the correct answer, that is the target output. (Agatonovic-Kustrin & Beresford, 1999)
In other words, the training dataset could be represented as a set of (x,y) example pairs, and the goal is to make a prediction y* as a response to the input x*.

A classic example is the binary classification problem in which, given an e-mail x in input, y can assume one of the two values “spam” or “not spam”. The training dataset will consist of a collection of emails, each already labeled with the correct tag “spam” or “not spam”, enabling the machine to learn how to correctly label new emails in input. (Jordan & Michell, 2020)

In general, input data may be vectors or more complex inputs such as images, videos, documents, DNA sequences, graphs and more.

Practical applications of the Supervised Learning methods can be found, for instance, in spam classifiers of e-mail, face recognizers over images, medical systems for diagnostics. (Jordan & Michell, 2020)

1.2.2 Unsupervised Learning

Despite the successes achieved with the use of Supervised Learning methods, efforts have also been made to train machines without labeled training datasets: this refers to the notion of Unsupervised Learning. (Jordan & Michell, 2020)

Unsupervised Learning has been historically overshadowed by the excellent results of Supervised Learning, but it is expected to become more and more relevant in the future. (LeCun, Bengio, & Hinton, 2015)

In Unsupervised Learning, the machine is provided with the inputs x, but not with the paired correct output y for the input x. The machine itself, in this case, has to analyse the unlabeled input data and decide which features will be used to group them. This procedure is generally known as self-organization, or adaptation. The system is able to come up with its own classification for the input data, analysing them to discover features and properties, grouping together data with similar features and reflecting their different properties in the output. (Agatonovic-Kustrin & Beresford, 1999)

1.2.3 Reinforcement Learning

A third approach to train a machine is Reinforcement Learning. In this case, the information available in the training dataset can be considered at the crossroads between Supervised Learning and Unsupervised Learning: paired examples with the correct output y for a given input x are not available, but the training dataset provides the machine an indication about whether the action done is correct or not. (Jordan & Michell, 2020)

As a consequence, in Reinforcement Learning the agent operates through a sequence of observations, actions, and eventual rewards. (Mnih, et al., 2015)
In case the action is not correct, the machine still has to find which action is the correct one. If the action is correct, the machine receives a reward signal. More specifically, considering sequences of inputs, the reward signal is not provided after each action, but it refers to an entire sequence of actions. (Jordan & Michell, 2020)
Therefore, the goal of the machine can be considered as selecting a sequence of appropriate output actions to maximize its cumulative future reward over time, by learning how to interact with this initially unknown environment; the learning task consists in learning a control strategy, a policy, to choose the appropriate action for any given input . (Schmidhuber, 2014) (Mnih, et al., 2015) (Jordan & Michell, 2020)

The theory of Reinforcement Learning is closely related to other fields of research, such as psychology and neuroscience: for example, Reinforcement Learning is used to predict the response of monkeys' neurons in learning, to associate stimulus lights with subsequent sugar rewards. (Jordan & Michell, 2020)

Systems using this kind of approach are still in their infancy, but they have already achieved outstanding results in a variety of domains, such as classification tasks and computer games. (LeCun, Bengio, & Hinton, 2015) The most well-known Reinforcement Learning system is probably the world-class Reinforcement Learning backgammon player, able to achieve the human world championship level. (Schmidhuber, 2014) Anyway, the application of Reinforcement Learning is still mainly limited to situations in which useful features can be handcrafted, or with fully observed, low-dimensional state spaces. (Mnih, et al., 2015)

Despite these three different learning paradigms have been proposed as separated, it is fundamental to underline how boundaries between them may blur, so that much of the current research is focused on combining them.

For example, in Semi-Supervised Learning, Unsupervised Learning is used to facilitate Supervised Learning by using unlabeled data to augment labeled data. Or for example, Discriminative Training combines Unsupervised Learning systems with optimization algorithms making use of labeled data. (Jordan & Michell, 2020)

1.3 Deep Learning

Increasingly, Machine Learning applications are making use of a class of techniques called Deep Learning. (LeCun, Bengio, & Hinton, 2015)

For decades, Machine Learning methods have been limited in their capability to process raw data in input: building a Machine Learning system required a large domain expertise and a scrupulous design, to be able to develop a feature extractor to transform data in their raw form (for example the pixel of an image) into an internal representation, or feature vector, from which the learning subsystem, often a classifier, was able to detect patterns in inputs or to classify them. (LeCun, Bengio, & Hinton, 2015)
As opposed to these traditional Machine Learning techniques, Deep Learning techniques have generated increasing attention since the proposal of a first learning algorithm for a deep structured architecture in 2006. This is largely due to their capability to overcome the dependence on hand-engineered feature extractors, main limitation of traditional Machine Learning techniques, by automatically learning good features for patterns detection or classification from data. (Liu, et al., 2017)

Deep Learning is considered as belonging to the broader family of Representation Learning. This notion refers to a group of techniques allowing a system to be fed with raw data and automatically discover the representation of data needed for patterns identification or classification. Deep Learning methods, in particular, are Representation Learning methods characterized by multiple levels of representation. This means that using non-linear modules, they are able to transform the representation of data at each level (beginning with the input in the raw format) into a representation of data at a higher and more abstract level. (LeCun, Bengio, & Hinton, 2015)

Generally speaking, a Deep Learning algorithm consists of a hierarchical architecture composed of multiple layers, each of which constitutes a non-linear information processing unit. These several layers are used to construct a progressively more abstract representation of data, making it possible for deep architectures to learn concepts such as object categories from data in the raw form. (Liu, et al., 2017)
For example, an image in input is an array of pixel values. In the first layer of representation, the learned features are typically the presence or not of edges at particular orientations and locations of the image. In the second layer, motifs are typically detected, by identifying specific arrangements of edges. In the following layer motifs may be combined into combinations corresponding to parts, and subsequent layers may detect objects as a combination of these parts. (LeCun, Bengio, & Hinton, 2015)
These deep architectures, since composed of multiple layers between the input and the output layers, are generally referred as Deep Neural Networks. (Liu, et al., 2017)

Deep Neural Networks are a category of Artificial Neural Networks. Artificial Neural Networks are systems of fully interconnected neurons organized in layers: an input layer, an output layer and one or more hidden layers between them. As shown in Figure 4, two different types of Artificial Neural Networks exist: shallow or deep. Shallow Artificial Neural Networks have only one hidden layer (there is one layer between the input and the output layer). Deep Artificial Neural Networks have multiple hidden layers between the input and the output layers.

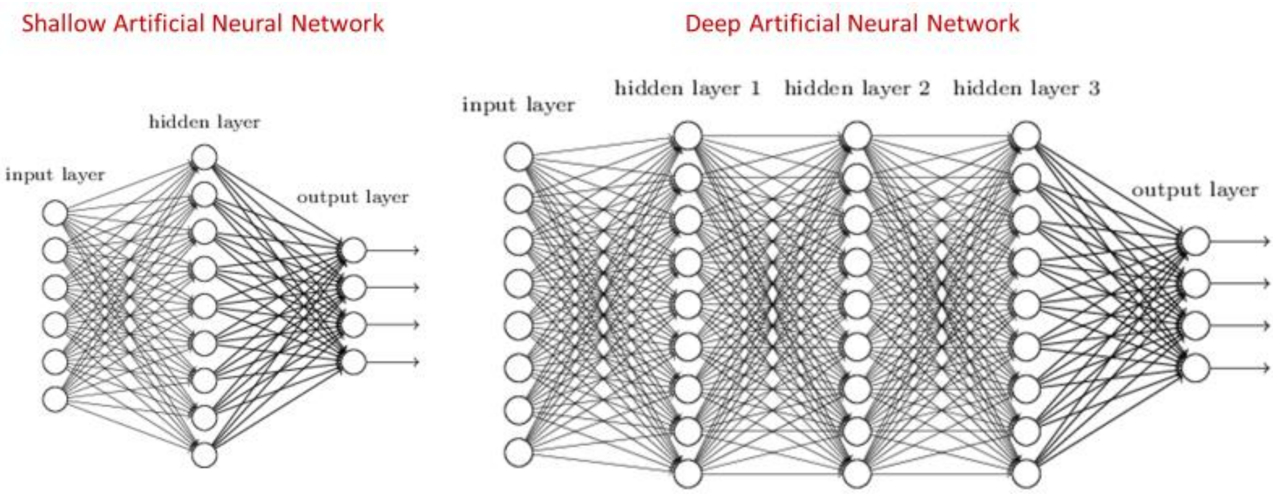

Figure 4: Representation of a Shallow and Deep Artificial Neural Network - http://neuralnetworksanddeeplearning.com/chap5.html

Artificial Neural Networks, and so Deep Neural Networks, are biologically inspired systems, designed to simulate the way the human brain works and processes information. Since Deep Learning is a class of techniques in the broader field of Machine Learning, these digitized models of the human brain are learning, or are trained, through experience, as humans do, and not from programming. (Agatonovic-Kustrin & Beresford, 1999)

Particularly, the human brain is excellent in pattern detection: if we look at a pen, we are able to associate that object with the object description "pen", since neurons in our brain have come across a similar input pattern in the past and have learned to link that specific input pattern with the object description "pen". Since an average brain comprises 100 billion neurons, each of which has between 1000 and 10000 connections with other neurons, an almost infinite range of input patterns can be learned and recognized. Neurons in the human brain are organized in a fully interconnected network and behave as messengers, receiving signals from neurons and sending them to other ones. The final result is a brain excellent in learning and recognizing patterns. (Agatonovic-Kustrin & Beresford, 1999)

Biologically inspired by the way through which the human brain processes information, an Artificial Neural Network consists of a few hundreds, or thousands, of single units, namely artificial neurons or processing elements, constituting the nodes of the neural structure. These artificial neurons are the building component of an Artificial Neural Network and are used to process information, but the real power of neural computation comes from connecting neurons in the network. As previously mentioned, neurons are organized in layers (with an input layer, an output layer and the hidden layers between them): neurons from the input layer receive input data, while neurons in the output layer provide the response of the network to the input data. Neurons are connected with neurons of other layers through connection coefficients (weights). These weights are adjustable parameters that makes an Artificial Neural Network a parametrized system. (Agatonovic-Kustrin & Beresford, 1999)

Each artificial neuron has many weighted inputs, a transfer function and a unique output. (Figure 5) The weighted sum of the inputs of a neuron activates the neuron itself. Then, the weighted sum of the inputs of the neuron is passed through the transfer function to produce the output of the neuron. (Agatonovic-Kustrin & Beresford, 1999)

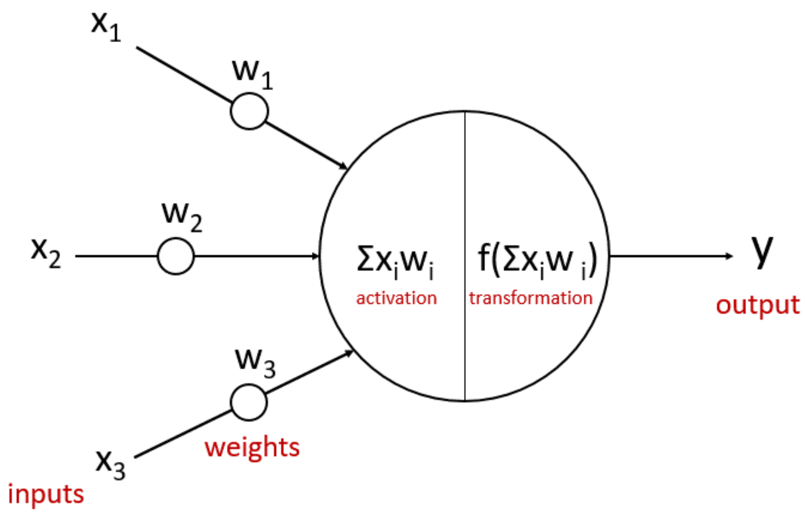

Figure 5: Representation of the functioning of an artificial neuron - (Agatonovic-Kustrin & Beresford, 1999)

Particularly, the functioning of an Artificial Neural Network is determined by: the transfer function of the artificial neurons; a learning rule; the architecture of the network itself. Different types of Artificial Neural Networks have been proposed, but all of them can be described in terms of transfer function, learning rule and connection formula. (Agatonovic-Kustrin & Beresford, 1999)

a. Transfer function of the neurons

An Artificial Neural Network consists of hundreds or thousands of single units, artificial neurons, constituting the nodes of the neural structure. These nodes are organized in layers, process information and are connected with nodes of other layers. Each artificial neuron has many weighted inputs, a transfer function and a unique output. The inputs, signals arriving in the neuron multiplied by the connection weights, are first of all summed. The weighted sum of the inputs activates the neuron itself, and the result is then passed through the transfer function to produce the output of the neuron. The transfer function is a non-linear function and introduces non-linearity in the Neural Network. (Agatonovic-Kustrin & Beresford, 1999) Modelling the output y of the neuron as a function of the weighted sum of the inputs x, the most commonly used transfer functions are (Krizhevsky, Sutskever, & Hinton, 2010):

- $f(x) = \tanh(x)$;
- $f(x) = (1 + e-x)-1$
- $f(x) = \max(0, x)$.

At present, the latter is the most popular one. (LeCun, Bengio, & Hinton, 2015) In particular, neurons using this non-linear function are generally referred as Rectified Linear Units (ReLUs). Use of ReLUs allows a faster training time in neural networks if compared with the first two transfer functions. (Krizhevsky, Sutskever, & Hinton, 2010)

b. Connection formula

The functioning of an Artificial Neural Network depends also on the way artificial neurons are connected in the network. Neurons can receive either excitatory or inhibitory inputs signals from another connected neuron, rather than being inhibited by other neurons in the same layer. Another type of possible connection is named feedback: in this case, the outputs of one layer are reused as inputs of the previous layer, or as inputs of the same layer.

Depending on the existence or not of feedback connections in a Neural Network, two different types of architecture can be defined: Feedforward Neural Networks (FNNs) and Feedback Neural Networks, or Recurrent Neural Networks (RNNs). (Figure 6) (Agatonovic-Kustrin & Beresford, 1999)

Feedforward Neural Networks are acyclic graphs, while Recurrent Neural Networks are cyclic graphs. (Schmidhuber, 2014) In the first case, the outputs of one layer are not reused as inputs for the same or previous layers, and so records of previous output values are not memorized. On the other side, in Feedback Neural Networks the outputs of one layer are used as inputs for the same or previous layers; in this case, records of previous output values are memorized, so that the next state of the network depends not only on inputs, but also on the previous states of the network. (Agatonovic-Kustrin & Beresford, 1999)

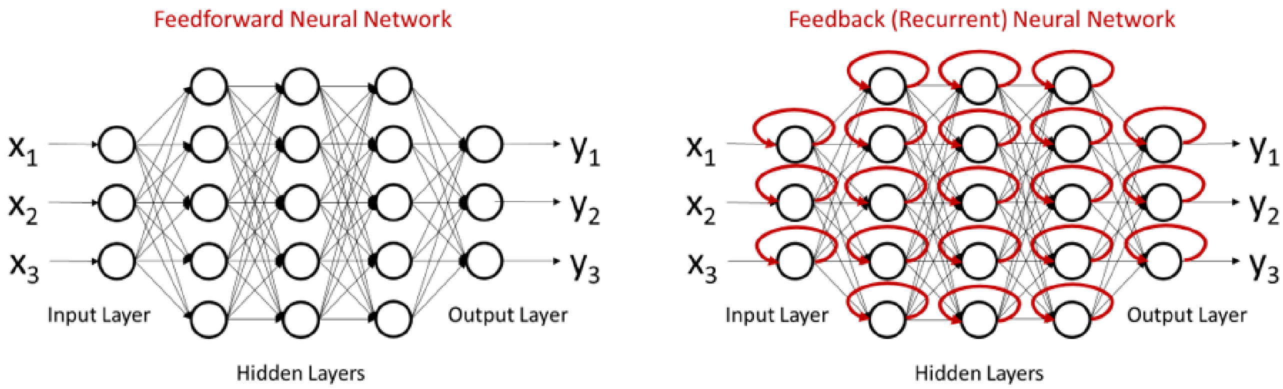

Figure 6: Representation of Feedforward and Recurrent Neural Network

c. Learning rule

Despite the existence of several learning rules to train an Artificial Neural Networks, the most commonly used method is the backpropagation one: the network is trained to map input data by iteratively adjusting weights, internal adjustable parameters used to compute the representation in each layer from the representation in the previous one.

Backpropagation is an efficient stochastic gradient descent algorithm to train an Artificial Neural Network: through backpropagation, internal weights between the neurons are optimized with the backward propagation of the error through the network during training, or learning, phase. (Liu, et al., 2017)

Basically, inputs and outputs values in a labeled training dataset are considered and weights progressively adjusted to minimize the error in prediction, difference between the predicted output values and the target output values, until the system reaches a satisfactory level of performance. (Agatonovic-Kustrin & Beresford, 1999)

For example, consider designing a system to classify images containing a person, a pet, a house or a car. The first step is the collection of a large number of images of people, pets, houses or cars, each one labeled with the correct category. This set of images constitutes the training dataset, and it is used to train the machine.

In the training phase, the system is fed with images from the training dataset and returns as output a vector of scores, one for each of the categories. Since we desire the correct category to have the highest score in this vector, but this is unlikely to happen at the beginning, we introduce an objective function measuring the prediction error, distance between actual scores in output and the desired ones. Through backpropagation, weights are then adjusted to reduce the prediction error, by computing a gradient vector that indicates for each weight how much the error would increase or decrease by slightly increasing the weight. Then, the weight vector is adjusted in the opposite direction with respect to the gradient vector, to reduce the output error and improve the system. This process is repeated until the system reaches a satisfactory performance level, and the weights values are saved. (Liu, et al., 2017)

After training, a separated test dataset is used: this is a set with different images, used to test the ability of the network to answer to new inputs not encountered in the previous training phase. Here the Artificial Neural Network works only by forward propagation of data, without any backpropagation of error. Once the network has successfully passed through the training and testing phases, it can be fed with new inputs to return an output. (Agatonovic-Kustrin & Beresford, 1999)

Being excellent in discovering intricate structures in data, Deep Learning systems have dramatically improved the state of the art in Computer Vision, Speech Recognition, Object Detection Natural Language Processing and many other fields, where they have yielded major improvements in performance over traditional Machine Learning techniques. In recent years, this class of methods have also won a multitude of contests in pattern recognition, image detection and Machine Learning. (Schmidhuber, 2014)
A great contribution to the booming of Deep Neural Networks has come from the use of low-cost high-performance graphics processing units (GPUs) originally developed for video gaming. The first GPU-based Deep Neural Network dates to 2006, and because of GPUs' outstanding performances if compared to traditional CPUs, they have become more and more relevant in the following years. (Schmidhuber, 2014) GPUs made it possible to build GPU-based Deep Neural Networks containing billions of parameters, allowing to train them 10 or 20 times faster than before with large datasets of images, videos, and

speech samples available on the Internet. (Jordan & Michell, 2020) (LeCun, Bengio, & Hinton, 2015)
To conclude, despite the great successes obtained with Deep Learning methods, these digitized models of the human brain are still far beyond the complexity, functions, and creative capacity of the human brain. Artificial Neural Networks used for o
Object Recognition or Natural Language Processing contain tens or hundreds of thousand units, compared with the 100 billion neurons of the human brain. (Agatonovic-Kustrin & Beresford, 1999)

In conclusion, in the following Sections two different types of Deep Neural Networks achieving outstanding results in the recent years will be reviewed: Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs).
In particular, Convolutional Neural Networks have been a major breakthrough in image and video processing and have been widely adopted by the Computer Vision community, while Recurrent Neural Networks have achieved many practical successes in processing sequential data such as text and speech. (LeCun, Bengio, & Hinton, 2015)

1.3.1 Convolutional Neural Network

Convolutional Neural Networks, or ConvNets, are a particular type of Deep Feedforward Neural Networks, designed to process input data coming in the form of multiple arrays (consider for instance a color image consisting of three 2D arrays containing pixel intensities in the three color channels). Many input data are in the form of multiple arrays: 1D for signals and sequences; 2D for images and audio spectrograms; 3D for videos and volumetric images. (LeCun, Bengio, & Hinton, 2015)

The typical architecture of a Convolutional Neural Network is biologically inspired by the animal visual cortex organization. As Deep Neural Networks, they are composed of an input layer, an output layer and several hidden layers. Moreover, a Convolutional Neural Network consists of three different types of hidden layers: convolutional layers, pooling (or sub-sampling) layers and fully connected layers. These different types of layers have different roles and perform different tasks. In particular, the first stages of the network consist of alternated convolutional layers and pooling layers. These two kinds of layers are connected alternately, and information is passed through them in a feedforward direction. (LeCun, Bengio, & Hinton, 2015) (Guo, et al., 2015)

On one side, convolutional layers are used to extract features from input data. When an image is received in input, this is combined with convolutional filters to create feature maps. Feature maps are the way units in a Convolutional Neural Network are organized. When an input image, or an intermediate feature map, is received in input, it is combined with convolution filters to produce transformed feature maps, achieving a higher level of abstraction. A filter, or filter bank, is a vector of connection weights used to extract the most useful information from the received inputs. A filter bank allows a unit in a feature map to be connected with local patches in the feature maps of previous layers. The result of the weighted sum is then passed through a non-linear transfer function, generally a ReLU. The ReLU non-linearity is applied to weighted input data of every convolutional layer and allows to produce additional feature maps used as inputs for the following pooling layers. (LeCun, Bengio, & Hinton, 2015)

On the other side, pooling layers are used to obtain invariance in image transformation: by computing the maximum of a local patch of units in one feature map and summarizing neighbourhoods in the same feature map, they are able to combine similar features. In this way, they both reduce the dimension of the representation for the next convolutional layer and achieve invariance to small image shifts and distortions in the input. (LeCun, Bengio, & Hinton, 2015) (Liu, et al., 2017)

Typically, in a Convolutional Neural Network two or three stages of convolution, non-linearity and pooling are typically stacked, with the same operations previously described repeated for each layer. These stages are followed by more convolutional layers and some fully-connected layers that convert feature maps into a 1D feature vector, fed forward into a number of categories for classification. (LeCun, Bengio, & Hinton, 2015) (Voulodimos, Doulamis, Doulamis, & Protopapadakis, 2018)

To conclude, as common Deep Learning Networks, Convolutional Neural Networks are trained by using standard Backpropagation algorithms to adjust the weights in all the filter banks for achieving a satisfactory level of performance. (LeCun, Bengio, & Hinton, 2015)

Since their introduction in the early 1990's, Convolutional Neural Networks have achieved outstanding performances in tasks such as Optical Character Recognition and Face Detection. (Ziler & Fergus, 2014) At present, these architectures have between 10 and 20 layers of ReLUs, billions of connections between the different units and hundreds of millions of connection weights. (LeCun, Bengio, & Hinton, 2015)

Training such a large network could represent a problem till a few years ago, requiring weeks to do so. Fortunately, recent progress in hardware (with the introduction of current GPUs), software and methodologies have reduced the training time to just a few hours. (LeCun, Bengio, & Hinton, 2015) Moreover, since objects in a real-world context exhibit relevant variability, training of Convolutional Neural Networks requires very large datasets to be able to recognize objects. At this regard, a great contribution has come from the emergence of new datasets with millions of labeled images, such as LabelMe and ImageNet. (Krizhevsky, Sutskever, & Hinton, 2010)

To conclude, at present Convolutional Neural Networks have been successfully applied in a variety of tasks, such as handwriting recognition, speech recognition, face detection, Natural Language Processing, image classification and much more. (Liu, et al., 2017) (Ziler & Fergus, 2014)
The excellent results of Convolutional Neural Networks have caused a revolution in the Computer Vision field, rapidly becoming the dominant architecture used for recognition and detection tasks. These promising results are the reason why all the relevant technology companies, as well as an increasing number of startups, have started research and development projects in this field of study and deployed image understanding products and services, with real-time vision applications in smartphones and cameras. (LeCun, Bengio, & Hinton, 2015)

1.3.2 Recurrent Neural Network

Alongside Convolutional Neural Networks, another type of Deep Neural Network has achieved outstanding results in recent years: Recurrent Neural Networks. When the backpropagation algorithm was firstly introduced, it was widely employed to train Recurrent Neural Networks. These are a type of Deep Neural Networks commonly used for tasks requiring the processing of sequential input data, such as text and speech. (LeCun, Bengio, & Hinton, 2015)

As previously introduced, Recurrent Neural Networks are cyclic graphs, in which the outputs of one layer are reused as inputs for the same or previous layers (Figure 7); in this way, by memorizing the records of previous output values, the next state of the network depends not only by the inputs, but also by the previous states of the architecture. (Agatonovic-Kustrin & Beresford, 1999)
This is done by processing a sequence of inputs one at a time, keeping in the hidden units of the Recurrent Neural Network a state vector. This allows to keep track of information about all the previous elements of the input sequence, to create and process memories of input sequences. (LeCun, Bengio, & Hinton, 2015)

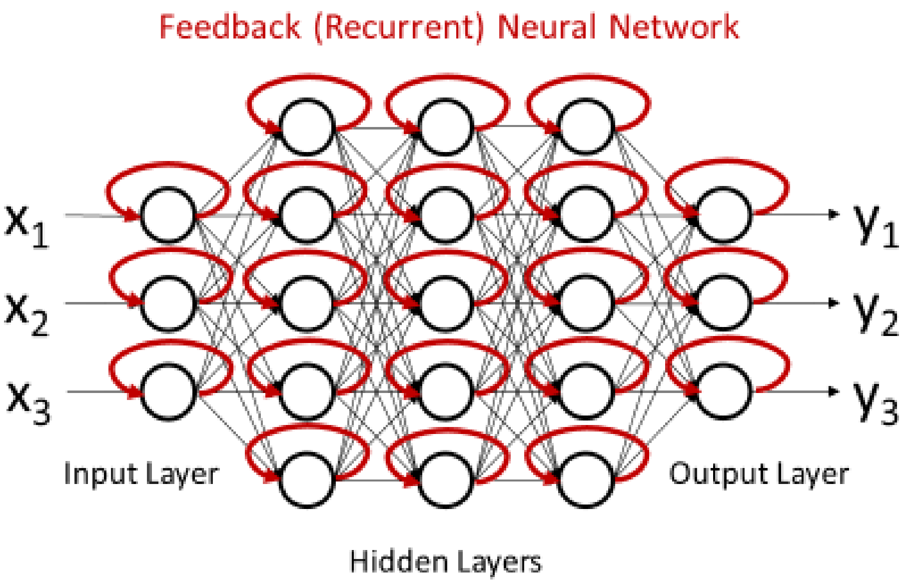

Figure 7: Representation of a Recurrent Neural Network

Particularly, one of the main drawbacks of conventional Recurrent Neural Networks is the difficulty in learning to store information for a long time. To overcome these difficulties, Recurrent Neural Networks have been improved with the introduction of the

Long Short-Term Memory method: Long Short-Term Memory Recurrent Neural Network introduces special hidden units that act as information accumulators and allows remembering inputs for a longer time. Long Short-Term Memory networks are much more effective than conventional ones, enabling for example a speech recognition system to process acoustics and to transcribe the entire sequence of related characters. (LeCun, Bengio, & Hinton, 2015) (Buchanan, 2006)

Recurrent Neural Networks have achieved excellent results in performing tasks such as predicting the next word in a sequence or the next character in a text, but they have also found application in more complex tasks such as language translation. For instance, after reading an English sentence word by word, the system can be trained so that the state vector in the hidden units represents the meaning of the sentence. This vector can then be used for a translation from English into French, providing as output a probability distribution for the first word of the translated sentence. The first word selected will then be used to generate a probability distribution for the subsequent word of the French sentence. This procedure allows to create a sequence of French words, on the basis of a probability distribution depending on the English phrase. (LeCun, Bengio, & Hinton, 2015)

1.4 History of Artificial Intelligence

After understanding the main theoretical concepts and methods behind Artificial Intelligence, this Section is intended to provide an historical contextualization of them. Therefore, a brief historical overview of the significant events and relevant progresses in the history of Artificial Intelligence and Machine Learning is here provided.

1940
The early Shallow Neural Network architectures are developed. These early systems are not able to learn yet. (Schmidhuber, 2014)

1950

Alan Turing's seminal paper "Computing Machinery and Intelligence" is considered the starting point in the history of Artificial Intelligence. The expression "Artificial Intelligence" has not been coined yet, but the paper introduces the idea of programming a computer to behave intelligently and includes a description of the Turing's Test, a test to understand if a machine is intelligent or not. (Buchanan, 2006)

1956

Allen Newell, Clifford Shaw, and Herbert Simon develop an earlier proof of concept, the Logic Theorist, considered the first Artificial Intelligence program. It shows intelligence and creativity by inventing proofs of logic theorems. It was presented at the Dartmouth conference on Artificial Intelligence, the meeting that coined the expression "Artificial Intelligence". (Buchanan, 2006)

1959

Artificial Intelligence laboratory is founded at the Massachusetts Institute of Technology (MIT), followed closely behind by Artificial Intelligence laboratories at Stanford and Edinburgh. (Buchanan, 2006)

1960

Research on the cat's visual cortex provides inspiration for Deep Learning, inspiring later Deep Neural Networks. A first Shallow Artificial Neural Network using Supervised Learning is developed. (Schmidhuber, 2014)

1965

A first learning Feedforward Deep Neural Network, trained with the Group Method of Data Handling (GMDH), is developed. (Schmidhuber, 2014)

1974-1980

Artificial Intelligence meets its first Winter, because of excessive expectations not met: interest and funding in Artificial Intelligence research declines, as a result of excessive enthusiasm and hype for the technology. In 1979 the Neocognitron, considered the first Convolutional Neural Network, was developed. Connection weights are not set through backpropagation yet, but by local WTA-based Unsupervised Learning rules. (Schmidhuber, 2014)

1981

Backpropagation, a gradient descent method for Supervised Learning, is for the first time applied to train a Neural Network. (Schmidhuber, 2014)

1987-1993

Second Artificial Intelligence Winter: again, interest and funding in Artificial Intelligence research declines, because of excessive expectations and hype for the technology.
In 1991 backpropagation training of Deep Neural Networks is also found to be difficult in practice, because of the problem of vanishing/exploding gradients: cumulative backpropagated error signals shrink rapidly or grow out of bounds. In the following years, several methodologies to alleviate this problem have been proposed.
In the same year, a very deep hierarchy of Recurrent Neural Networks using Unsupervised Learning is proposed, while the 1992 sees the introduction of the Neocognitron, a Convolutional Neural Network inspired by the Cresceptron, able to adapt its topology during training and substituting WTA methods with Max-Pooling layers. (Schmidhuber, 2014)

1995

The first Supervised Recurrent Neural Network using Long Short-Term Memory is proposed. (Schmidhuber, 2014)

1997

Major milestone in the history of AI: the chess world champion Gary Kasparov is defeated by an IBM computer. Despite this success, computers are still not able to compete with a child in tasks such as visual pattern recognition. (Schmidhuber, 2014) (Buchanan, 2006)

2003

Many applications are still dominated by non-neural Machine Learning techniques. Anyway, Neural Networks start to outperform conventional methods in certain domains, winning contests and setting records in features selection and speech recognition.

(Schmidhuber, 2014)

2006

Interest in feedforward Deep Neural Networks is strongly revived because of the integrated use of Unsupervised Learning and Supervised Learning to accelerate the training of Deep Neural Networks using backpropagation. This event marks the birth of Deep Learning: the expression Deep Learning is coined, despite the presence of the early learning Deep Neural Networks dating back to 1965. The same year saw an early GPU-based implementation of a Convolutional Neural Network, with the network trained 4 times faster than a CPU-based implementation. (Schmidhuber, 2014) (LeCun, Bengio, & Hinton, 2015)

2009

First official international competitions in pattern recognition won by Recurrent Neural Networks. Deep Learning is officially rehabilitated and attracts widespread attention. (Schmidhuber, 2014) (LeCun, Bengio, & Hinton, 2015)

2012

An ensemble of GPU-based Convolutional Neural Network, using Max-Pooling and Backpropagation wins the ImageNet Image Classification competition, popular in the Computer Vision community, achieving excellent results in the classification benchmark. Several other contests are won in image recognition and classification. (Schmidhuber, 2014) (Liu, et al., 2017) This success represents a revolution in the Computer Vision community: Convolutional Neural Networks are nowadays the commonly used method for recognition and detection purposes. In the same year, Long Short-Term memory Recurrent Neural Networks outperform statistical approaches on the IAM-DB benchmark in language modelling. (LeCun, Bengio, & Hinton, 2015)

2013

More contests are won and records set by GPU-based Convolutional Neural Networks and Long Short-Term Memory Recurrent Neural Networks in Optical Character Recognition (OCR), language identification and translation, speech recognition, text-to-speech synthesis, image and video classification, object detection. Convolutional Neural

Networks also support the detection of numbers in Google Street View pictures. (Schmidhuber, 2014)

2016

A Deep Neural Network developed by Google defeats in a Go Game match in South Korea one of the world's strongest players, Lee Se-dol (Liu, et al., 2017)

Present

At present, the majority of competition-winning or benchmark-record setting Deep Neural Networks are Long Short-Term Memory Recurrent Neural Networks and GPU-based Convolutional Neural Networks, using Supervised Learning and trained through the backpropagation algorithm. (Schmidhuber, 2014)

1.5 Classes of Solutions

Once presented the main theoretical concepts behind Artificial Intelligence and a historical contextualization of them, the role that the technology can have in a business context is considered in this Section: according to the Artificial Intelligence Observatory of the Politecnico of Milan, when researchers talk about Artificial Intelligence, they implicitly refer to a multiplicity of "Classes of Solutions". These Classes of Solutions have been explicitly defined by the Observatory to consider the possible applications of Artificial Intelligence in business, due to the rapid diffusion of the technology in a corporate setting.

Organizing Artificial Intelligence projects according to this classification helps to understand the different areas and scope that companies want to achieve using this type of technology. Some Classes of Solutions are created with Artificial Intelligence, while others come from different applications, like Analytics, IoT and Industry 4.0. Since Artificial Intelligence is widespread in this field, some categories of classification cover those classes. (Perego, et al., 2019) Each Class of Solutions is furthermore associated with different Specifications, defining the specific use of Artificial Intelligence solutions.

According to the Artificial Intelligence Observatory, 9 different Classes of Solutions can be identified:

a. Intelligent Data Processing: this category includes those solutions that use Artificial Intelligence on structured and unstructured data, with the aim to extract information present in them to support the decision making process. (Perego, et al., 2020)

b. Natural Language Processing: this category includes solutions for language processing. The objectives vary from understanding the content, to translation, up to production of text. (Perego, et al., 2020)

c. Virtual Assistant/Chatbot: this class includes software agents capable of performing action and providing services to a person, based on commands and requests received through written or spoken interaction. (Perego, et al., 2020)

d. Computer Vision: the category includes solutions for image and video analysis, oriented to the recognition of people, animals and things present within the image. It also includes biometric recognition and other solutions to generally extract information from an image or video. (Perego, et al., 2020)

e. iRPA: this class refers to the execution of recurring tasks by a software integrated with Artificial Intelligence capabilities. (Perego, et al., 2019)

f. Recommendation Systems: this category includes solutions with the objective of addressing the preferences, interests or more generally the decisions of a user based on information provided by him, directly or indirectly. The aim consists in providing personalized recommendations that can be placed at different points of the customer journey or the decision-making process. (Perego, et al., 2020)

g. Autonomous Robot: this class includes robots able to move themselves, or some of their parts, manipulate objects and perform various kinds of actions without human intervention, drawing information from the surrounding environment and adapting to unforeseen or coded events. (Perego, et al., 2020)

h. Autonomous Vehicle: this class includes self-driving vehicles used for the transportation of people, animals, or things, circulating on the road, intended for navigation or flight and capable of perceiving the external environment and identifying the correct manoeuvres to do to adapt to it. (Perego, et al., 2020)

i. Intelligent Object: this class of solutions includes objects and physical products capable of performing action and making decisions without requesting human intervention, interacting with the surrounding environment through the use of sensors (e.g. thermometers, video cameras, microphones…) and actuators (e.g. opening/closing doors, activating appliances…) and learning from habits or actions of people interacting with them. (Perego, et al., 2020)

To gain a better understanding of the topic, each Class of Solutions and the related Specifications have been review and discussed based on the existing literature in the following Sections, providing when possible additional information about practical applications in business, adoption and diffusion, benefits, challenges and concerns.

1.5.1 Intelligent Data Processing

This category includes those solutions that use Artificial Intelligence on structured and unstructured data, with the aim to extract information present in them to support the decision making process. (Perego, et al., 2020)

Companies are more and more overwhelmed by data, and this implies the necessity to use Artificial Intelligence to analyze them and extract relevant information. (Perego, et al., 2020)

While Artificial Intelligence was not developed in the past, also due to the existence of small datasets, nowadays there is a massive amount of data available, and this has contributed to both the development and wide use of Artificial Intelligence analytics algorithms. (Silwattananusarn & Tuamsuk, 2012)

Particularly, the expression “Big Data” has been coined to consider this data explosion trend. (Qiu, Wu, Ding, Xu, & Feng, 2016)
At present, Artificial Intelligence and Big Data have formed a very close relationship. We can easily say that there is no Artificial intelligence without Big Data, and vice versa. (Silwattananusarn & Tuamsuk, 2012)

1.5.1.1 Big Data: definition and challenges

The notion of Big Data refers to massive datasets that cannot be captured, processed, managed and analyzed by traditional software tools. These datasets are huge in size, heterogeneous, complex, and include operational, transactional, sales, marketing, and many other typologies of data. In addition, Big Data contains data that comes in a variety of forms including text, voice, video, image, and more.
While Structured Data are data with a predefined data model, such as Excel files or SQL databases, for which the application of computational analytic tools is therefore easier, Unstructured Data are data that do not have a predefined model or are not organized in a predefined manner, such as documents and PDF files, photos and images, videos, emails and web pages. These unstructured data are growing faster than structured ones, so that new types of processing capabilities are needed to get data insights for a better decision making. (Vassakis, Petrakis, & Kopanakis, 2018)

In the literature, Big Data are characterized by the so-called 5 “Challenges of Big Data” (Vassakis, Petrakis, & Kopanakis, 2018) or “Critical Issues” (Qiu, Wu, Ding, Xu, & Feng, 2016), identifying their main inherent characteristics and problems:

a. Volume: volume is referred to the large size of current datasets. The Internet of Things, in conjunction with the increasingly evolving Information and Communication Technologies including Artificial Intelligence, has led to an enormous generation of data, such as counting records, transactions, tables and files, through the creation and expansion of connected smartphones, sensors and other devices. The volumes of data generated are so impressive that data speed exceeds Moore’s law, and new measures for data storage, like exabytes and yottabytes, have been introduced. (Vassakis, Petrakis, & Kopanakis, 2018)

It is so evident how the primary attribute of Big Data is volume, which poses a greater challenge for Machine Learning to analyse them. For instance, every day Google alone needs to process about 24 petabytes of data, considering only digital data. (Qiu, Wu, Ding, Xu, & Feng, 2016) This characteristic of Big Data creates difficulty to train Machine Learning algorithms with central processor and storage. As a consequence, nowadays analysts use distributed frameworks with parallel computation. In addition, several parallel programming approaches have been proposed and applied to learning algorithms to deal with large scale data. (Qiu, Wu, Ding, Xu, & Feng, 2016)

b. Variety: variety reflects the increasing diversity regarding both data generation sources and data formats. The generation of different types of data is largely due to Web 3.0, which contributed to the growth of the web and social media networks, main providers of heterogeneous data. These ones range from messages, photos and videos posted in social networks to SMS, GPS signals from smartphones and smart objects, customer transactions in e-commerce platforms. Many of the main sources of Big Data are relatively new, including mobile devices that supply massive streams of data related to human activity, or web sources supplying data from records, click streams and actions on social media.
Meanwhile, Big Data often vary in the types of data generated, consisting of structured, unstructured and semi-structured data, as well as other difficult to classify data. (Vassakis, Petrakis, & Kopanakis, 2018)
A key solution to mitigate the problem of having heterogeneous data is data integration, referring to the integration of data from various sources, providing the user with a unified view of them. An effective approach to the problem of data integration is to learn good data representation from each individual data source and then incorporate the learned characteristics at different levels. Therefore, Representation Learning, a set of techniques allowing a system to autonomously identify the representations required for raw data detection, is therefore preferred in these situations. (Qiu, Wu, Ding, Xu, & Feng, 2016)

c. Velocity: velocity refers both to the high speed of data generation characterizing Big Data and to the importance of rapidly taking advantage of them. Since data generated by connected devices and web data arrive in companies in real-time, speed is extremely important for them to take actions resulting in competitive advantage. (Vassakis, Petrakis, & Kopanakis, 2018)

Companies must complete tasks within a certain amount of time in many real world applications, such as earthquake prediction, stock market prediction and autonomous buying/selling exchange systems; otherwise, the effects of processing data become less useful, or even useless. The potential value of data in these time-sensitive cases depends on the freshness of the data that need to be processed in real-time: through Big Data analytics and Artificial Intelligence, companies are now able to evaluate and understand data taking action in real time. (Qiu, Wu, Ding, Xu, & Feng, 2016) (Vassakis, Petrakis, & Kopanakis, 2018)

d. Veracity: data veracity refers to the uncertainty of data and degree of reliability, relevant since in collection of data information is not always clean and reliable. (Vassakis, Petrakis, & Kopanakis, 2018)

   In the past, Machine Learning algorithms have been typically fed with reasonably reliable data from well-known sources, with satisfactory outcomes of the learning algorithms. Therefore, veracity has never been considered a serious problem. At present, the accuracy and trust of the data source has quickly become a problem, since data sources often have several different origins, and the quality of data is often variable. (Qiu, Wu, Ding, Xu, & Feng, 2016)

   Moreover, data uncertainty is a reality when data reading and collection are no longer deterministic, but they are subject to random or probability distributions. Data uncertainty is frequent in many applications: for instance, in wireless networks, some spectrum data are inherently uncertain due to ubiquitous noise, fading and shadowing. One common approach to handle data uncertainty is to apply summary statistics to abstract sample distribution, such as means and variances. (Qiu, Wu, Ding, Xu, & Feng, 2016)

e. Value: the ultimate objective of companies with Big Data is to extract deep insights and useful information from vast quantities of data for commercial benefits, leveraging a range of learning methods to analyse big datasets. Knowledge discovery in datasets, data mining and Artificial Intelligence come into play to overcome this challenge, as these technologies offer potential solutions for finding the necessary information concealed in massive data. (Qiu, Wu, Ding, Xu, & Feng, 2016) Businesses solving obstacles and effectively leveraging Big Data, are able to provide more reliable information and to generate new knowledge from which they can enhance their strategy and business operations with well-defined targets such as productivity, financial results and market value, whereas big data plays a major role in digital transformation of innovation-driven companies. (Vassakis, Petrakis, & Kopanakis, 2018)

In summary, these key characteristics of Big Data represent different challenges for Artificial Intelligence and Machine Learning. Moreover, most Machine Learning techniques are not standardized for Big Data analysis, and so it is necessary to use specific learning methods according to different data. (Qiu, Wu, Ding, Xu, & Feng, 2016)

In several domains, Artificial Intelligence in data mining is commonly used to solve classification, planning, prediction, optimization problems, diagnosis, and computation, by collecting and analyzing customer information. Since Big Data are present almost everywhere, and because of the importance of not missing out any relevant information in them, the technology used to process Big Data is often Artificial Intelligence. Machine Learning and its multitude of techniques are all algorithm-based, and these algorithmic approaches are typically applied to Big Data to achieve the desired results and find trends, patterns and predictions. (Silwattananusarn & Tuamsuk, 2012)

1.5.1.2 Intelligent Data Processing: Specifications

The Intelligent Data Processing Class of Solutions is the broadest from the application point of view. This class includes all those solutions using Artificial Intelligence algorithms, on structured and unstructured data, for purposes related to the extraction of the information present in them. (Perego, et al., 2020)

It is important to underline how the use of Artificial Intelligence does not necessarily represent a discriminant for which a data analysis project is considered successful or not: indeed, it is the context that determines the most appropriate type of methodology to adopt. However, there are areas in which Artificial Intelligence is essential because, without it, it would not be possible to achieve the desired results. Consider, for instance, the identification of the relationships between variables in datasets of enormous size and with a high number of attributes for each record. (Perego, et al., 2020)

To provide a more precise picture and a better overview of Intelligent Data Processing, its main Specifications are briefly presented in this Section, according to the model proposed by the Observatory.

a. Forecasting

   This Specification includes solutions with the aim to identify the value that one or more variables may assume in the future. Application areas are numerous: from predictive maintenance, to demand forecast and churn prediction. (Perego, et al., 2020) Forecasting allows a company to assess its history and thereby be able to concentrate more clearly on its future, as well as to predict future market patterns. Therefore, companies can better prepare, use resources more efficiently, manage inventory, remain competitive and updated on trends, as well as many other things. Particularly, forecasting systems supported by Artificial Intelligence have proven to be much more accurate than the ones not implementing Artificial Intelligence, underlying the importance of the technology in Forecasting. (Annor-Antwi & Al-Dherasi, 2019)

   An example of practical application of Artificial Intelligence in Forecasting regards demand forecasting: companies can use Artificial Intelligence-based demand forecasts to produce and deliver the right amount of goods, with subsequent savings in inventory costs and improvements in capital performances. The combination of Artificial Intelligence and Machine Learning is essential, since it allows companies to get more accurate and reliable results. (Annor-Antwi & Al-Dherasi, 2019)

b. Classification & Clustering

   This Specification includes solutions aiming to assign a category to each record from a predefined set, or to divide input data into groups, or clusters. (Perego, et al., 2020)

   In particular, the aim of Classification is to map a data item in input into one of different predefined classes: for instance, a customer of a bank is assigned to a predefined class of customers based on personal and economic information. On the other side, Clustering is not working with predefined categories, but it groups data into homogeneous sets to define a finite number of classes, so that objects in a cluster are similar to each other and dissimilar to objects in other clusters: for instance, information about consumption of different users of an utility can be used to create groups of users based on similarities and dissimilarities in consumption and to target them with tailored offers. (Silwattananusarn & Tuamsuk, 2012)

For instance, examples of practical Classification & Clustering applications in business are the use of Classification algorithms in Healthcare, for the classification of patients from healthcare centres to specialists, or the use of Clustering algorithms in Retail, clustering clients into segments to better understand the target market for future product lines and brand extension. (Silwattananusarn & Tuamsuk, 2012)

c. Identification

This specification includes solutions aiming to identify anomalies or patterns within a large amount of data. (Perego, et al., 2020)

Identification is one of the broader Specifications, so that it is utilized by companies belonging to every sector. Examples of applications for this Specification can be found in Finance, where it is applied to economic transactions to identify potential illegal activities, or in Pharma, where it is used for the identification of recurring patterns or links between active ingredients, medicines and diseases (Perego, et al., 2020)

Or again, in cybersecurity, a malicious executable is a threat to the protection of the system, damaging it or obtaining confidential data without permission from the user: Artificial Intelligence algorithms based for Identification can be used to detect malicious executables before they are executed. (Padhy, Mishra, & Panigrahi, 2012)

d. Optimization

This Specification includes solutions aiming to achieve the optimal value for a target variable, by modifying the variables of the system to achieve the optimum.

Optimization Artificial Intelligence algorithms are very diffused in different industries, from Banking to Logistics. (Perego, et al., 2020)

For instance, optimization algorithms supported by Artificial Intelligence can be used for Mobile Network Optimization, enabling the design of mobile networks to ensure efficient service delivery. Or, for example, Artificial Intelligence optimization algorithms can be used to optimally distribute the budget to different advertising media. (Ajah & Nweke, 2019)

e. Content/Design Creation

This Specification aims to use data analysis to create new contents or plan new services or products. Companies can use Artificial Intelligence to have a better computational advantage and to create new things in a more efficient way, delivering more value to the customer. (Perego, et al., 2020)

1.5.2 Natural Language Processing

Natural Language Processing is a range of computational techniques used to analyse and represent naturally occurring text at different levels of linguistic analysis, with the objective to achieve human-like language processing for a variety of tasks or applications. (Liddy, 2001) Natural Language Processing mainly deals with text, or any sequence of words that express one or more messages, such as web pages, posts, tweets, company information. The pure text is not structured in computational terms, being not coupled with a further organization of the information like other texts such as HTML that has a well-defined structure. (Liddy, 2001)

Recently, Natural Language Processing has gained a great deal of interest in computational representation and study of human language, with applications ranging from machine translation to detection of spam email, to extraction of information. (Khurana, Koli, Khatter, & Singh, 2017) The discipline has made significant progress in recent years, mainly due to the enhanced computational capabilities of new Artificial Intelligence techniques, such as Deep Learning. (Perego, et al., 2020)

There are many aspects concerning the dialogue between a man and a machine, such as phonetics, phonology, morphology, syntax, semantics, pragmatics and speech as a whole. Therefore, there are many Natural Language Processing tasks to automate these areas, such as recognition of the language and decomposition of the sentence into elementary units, or semantic analysis and sentiment analysis. (Liddy, 2001)

As many other modern disciplines, the origins of Natural Language Processing are mixed: indeed, it received an important emphasis by different groups with backgrounds influenced by the most various disciplines. (Liddy, 2001) Primary contributors to Natural Language Processing's discipline and practice are:

- Linguistic, that focuses on formal, structural language models and the discovery of universal languages.
- Computer Science, dealing with the creation of internal data representations and the efficient processing of such structures.
- Cognitive Psychology, that looks at the use of language as a gateway into the human cognitive process and attempts to model the use of language in a psychologically realistic way. (Liddy, 2001)

Generally, Natural Language Processing can be divided into two main components: Natural Language Understanding, that helps to understand the human language, and Natural Language Generation, focused on the creation of a human response.

1.5.2.1 Natural Language Understanding

Natural Language Understanding involves Linguistics, the science studying the meaning of language, language context and the various forms of the language. Linguistics is composed of different levels, here reported from the lower to the higher to better understand how Natural Language Understanding operates.

a. Phonology: it is the part of Linguistics which refers to the systematic arrangement of sounds. (Khurana, Koli, Khatter, & Singh, 2017) Sound waves, in a Natural Language Processing system accepting spoken inputs, are analyzed and encoded for understanding by different rules, or by comparison with the specific language model used in a digitized signal. (Liddy, 2001)

b. Morphology: morphemes, the smallest units of a word, represent the smallest units of meaning. (Khurana, Koli, Khatter, & Singh, 2017) Since the meaning of each morpheme is the same across different words, to understand the meaning of an unknown word humans could break down it into its constituent morphemes. Similarly, to gain and represent meaning, a Natural Language Processing system should understand the meaning conveyed by each morpheme. (Liddy, 2001)

c. Lexical: as humans, Natural Language Processing systems interpret the meaning of individual words. There are various ways of processing the world-level understanding. The first, and most important one, is a part-of-speech tag to each word. (Khurana, Koli, Khatter, & Singh, 2017) Additionally, at this lexical level a semantic representation of a meaning may be substituted for those words that have only one possible sense or meaning. (Liddy, 2001)

d. Syntactic: the focus of syntactic is on the analysis of words in a sentence to reveal the grammatical structure of the phrase. At this level, both grammar and parser are needed. (Khurana, Koli, Khatter, & Singh, 2017) The output of this processing stage is a representation of the sentence that exposes the relationships between the words with structural dependence. Syntax conveys meaning in most languages because order and dependency too contribute to meaning. For instance, the two sentences “the dog chased the cat” and “the cat chased the dog” differ only in terms of syntax, although they convey quite different meanings. (Liddy, 2001)

e. Semantic: while most of the people assume that meaning is now decided, all the possible levels that offer meaning need to be considered. Semantic processing defines the possible meanings of a sentence by focusing on the interaction between word-level meanings in the sentence. (Khurana, Koli, Khatter, & Singh, 2017) Similarly to how syntactic disambiguation of words that can work as multiple sections of speech is done at the syntactic stage, semantic level can involve the semantic disambiguation of words with multiple senses. Semantic disambiguation enables one and only one meaning of polysemic terms to be selected and used in the sentence’s semantic representation. (Liddy, 2001)

f. Pragmatic: this level takes into account the purposeful use of language in situations, considering context over the content of text for understanding and explaining how additional meaning is read into text lines without being encoded in them. (Khurana, Koli, Khatter, & Singh, 2017) This task requires much world knowledge, including the understanding of intentions, plans and objectives, to be executed. (Liddy, 2001)

Natural Language Processing systems generally implement modules to accomplish just the lower levels of processing, for a variety of reasons. First, at the higher level the application does not require interpretation. Secondly, the lower levels have been studied and applied more extensively. To conclude, the lower levels are applied to smaller units of analysis such as morphemes, sentences and words as opposed to the higher levels of language processing dealing with text and knowledge of the words. The statistical approaches needed for text analysis have been validated on the lower levels of analysis, while there are still few working systems incorporating the higher levels. (Liddy, 2001)

1.5.2.2 Natural Language Generation

The other component of Natural Language Processing, Natural Language Generation, deals with the generation of words, sentences, and meaningful paragraphs from an internal representation. (Khurana, Koli, Khatter, & Singh, 2017)

Therefore, while the role of Natural Language Processing, referring to the analysis of language to produce a meaningful representation, is equivalent to the one of a reader/listener, the role of Natural Language Generation is equivalent to the one of a writer/speaker, since dealing with the production of language from a representation.

While these two components of Natural Language Processing share much of the theory and technology, in Natural Language Generation additional planning capabilities are required, to decide what the system should generate at each point. (Liddy, 2001)

1.5.2.3 History of Natural Language Processing

After having introduced the notion of Natural Language Processing and its two components, this Section is intended to provide a brief historical overview of the significant events and relevant progresses in its history, also to have a better understanding of the discipline.

1940s: research in Natural Language Processing started with initial work regarding Machine Translation; the expression Natural Language Processing had not been coined yet. In 1946, computer translation was applied for braking enemy codes during WWII, while in 1949 Weaver's memorandum introduced ideas from cryptography and information theory to be applied for Language Translation. This memorandum provided inspiration for many subsequent projects and contributed to diffuse the idea of Machine Translation. (Young, Hazarika, Poria, & Cambria, 2018)

1950s: in 1957 Chomsky published Syntactic Structures, introducing the idea of generative grammar and shedding light on how linguistics could help Machine Translation. Other Natural Language Processing fields started to emerge, such as speech recognition. The first Natural Language Generation and Natural Language Understanding algorithms were proposed. (Young, Hazarika, Poria, & Cambria, 2018)

1960s: several prototype systems like ELIZA, the first publicly known chatbot, simulating the conversation between a psychologist and his patient, were developed. (Young, Hazarika, Poria, & Cambria, 2018) In 1966, excessive expectations and inadequacies of existing systems led the ALPAC (Automatic Language Processing Advisory Committee of the National Academy of Science) to declare that Machine Translation was not achievable in the short term, suggesting to not fund it. As a result, interest in Machine Translation and Natural Language Processing evaporated. (Young, Hazarika, Poria, & Cambria, 2018)

1970s: more attention was given to semantic issues, discourse phenomena and communicative goals and plans, as well as a relevant work was done in Natural Language Generation. (Young, Hazarika, Poria, & Cambria, 2018)

1980s: increasing Natural Language Processing awareness and development of the first basic applications. (Young, Hazarika, Poria, & Cambria, 2018)

1990s: fast growth of Natural Language Processing due to increased availability of electronic text, computers, and the advent of the Internet. In the same period, researchers at IBM acquired a wide number of English sentences and the corresponding French translation, enabling the construction of a probabilistic model for Machine Translation. In 1992 the Linguistic Data Consortium (LDC) was established, with the aim to make linguistic data available to researchers in a digital format. This made the transformation of Natural Language Processing in a Big Data field possible. (Hirschberg & Manning, 2015) Machine Learning starts to be applied for Natural Language Processing purposes. (Perego, et al., 2020)

2000s to present: Deep Learning and Neural Networks are applied for Natural Language Processing tasks. The first relevant results were in the early 2000s, but significant progress has been registered only in recent years. (Perego, et al., 2020) (Liddy, 2001)

1.5.2.4 Specifications

In this Section, a brief description of the main Specifications for the Natural Language Processing Class of Solutions is provided, according to the framework proposed by the Artificial Intelligence Observatory of Politecnico of Milan. Each Specification is in turn divided into possible sub-Specifications, each one describing a different Natural Language Processing task.

a. Information Retrieval

   Natural Language Processing data is unstructured and comes from a multitude of sources, including emails, social media, polls, web posts and papers. Typically, these texts are complicated, time-consuming, and costly to study. Therefore, businesses are continuously searching for clever ways to handle them, to efficiently tackle their business challenges and to make decisions.

   The Information Retrieval Specification was developed to consider the emergence of a consistent cluster of Natural Language Processing applications with unique features to automatically extract relevant information, take insights from a multitude of unstructured data, assist data analysts to perform brand monitoring and product reputation, or market intelligence activities among the various Natural Language Processing applications. Therefore, the Natural Language Processing applications included in this Specification are primarily linked to the business dimension and offer relevant support to the day to day decision of firms, enabling to extract information relevant for business. (Osservatorio Artificial Intelligence - Politecnico di Milano)

To provide a better detail into this Specification, the Artificial Intelligence Observatory of Politecnico of Milan identified different sub-Specifications, each one with its characteristics and peculiarities.

- Brand/Product Monitoring

  The process of market analytics regards the "inspection" of the various web and media channels to gain insights about a brand and its products, as well as information directly linked to its reputation. Since the wave of social media and Big Data has had a major influence on the user attitudes, the ability to consider the point of view of different consumer clusters allows organizations to deliver the best level of purchasing experience and service. (Osservatorio Artificial Intelligence - Politecnico di Milano)

  Social media provides a great opportunity and a wide basis for these analyses. For instance, some monitoring tools may enable businesses to understand the number of mentions of a certain brand or product, or even to discover the related subject in the conversation of a customer, to incorporate all these data and create successful market research. (Osservatorio Artificial Intelligence - Politecnico di Milano)

- Sentiment Analysis

  In business, the automated method of classifying web or media information as positive, neutral, or negative is called Sentiment Analysis. Sentiment Analysis provides corporations the ability to understand how customers consider specific aspects of their business. (Osservatorio Artificial Intelligence - Politecnico di Milano)

  Particularly, Sentiment Analysis consists in the extraction of sentiments and associations through relationship analysis of a topic-specific feature phrase. For the analysis, two linguistic resources are typically utilized: the sentiment lexicon and the database of the sentiment pattern. (Khurana, Koli, Khatter, & Singh, 2017)

Companies use Sentiment Analysis in several applications, for instance to analyse survey responses or to find the competitors' unsatisfied customers, leading them to purchase their product. (Osservatorio Artificial Intelligence - Politecnico di Milano)

- Intent Monitoring

  This task is intended to process the words and terms expressed in a text or voice discourse to anticipate the actions or intent of a customer, like a future purchase or an unsubscription. (Osservatorio Artificial Intelligence - Politecnico di Milano)

  Intent Monitoring particularly helps businesses in areas such as customer service or commercial sales, enabling a higher customer centricity. With an intent classifier, a simple purchasing intent can be categorized, saving time and money in marketing and sales processes, or helping to produce effective product analytics. (Osservatorio Artificial Intelligence - Politecnico di Milano)

- Text Analysis

  Text Analysis is the automated method through which unstructured text is understood, sorted, and elaborated. All Natural Language Processing applications with the function of reviewing vast collections of documents to find new information, or helping to address specific research questions, belong to this category. Companies can better organise their data by analyzing text and extracting a variety of elements such as topics, people, dates, or locations, and use them to identify valuable trends and insights. (Osservatorio Artificial Intelligence - Politecnico di Milano)

  Natural Language Processing performs the linguistic analyses necessary to interpret and understand documents, while the information extracted is elaborated by Artificial Intelligence algorithms. (Osservatorio Artificial Intelligence - Politecnico di Milano)

b. Information Filtering

The second Specification for Natural Language Processing applications is Information Filtering. Many businesses are introducing filtering systems to reduce non value-adding time. Several applications for this specification exist, from the simple identification and classification of emails to the recruitment activity process. (Osservatorio Artificial Intelligence - Politecnico di Milano)
Again, Information Filtering can be divided into several sub-Specifications, helping to identify the several areas in which this Specification can be applied.

- E-mail Filtering

  This sub-specification consists in sorting and handling the emails that arrive in an inbox according to specified criteria. It is used to improve the protection of emails and preventing time wasted in filtering all the arriving emails. One example is the email filtering provided by Google, with a Natural Language Processing system recognizing whether emails belong, based on their content, to one of three predefined categories (primary, social or promotional). (Osservatorio Artificial Intelligence - Politecnico di Milano)

- Smart Screening

  This sub-Specification allows to filter specific data that companies are looking for from a huge quantity of data, allowing some consistent business practices to be accelerated and to make more specific search with less errors. Companies use Smart Screening to filter internal documents, web news or CVs received by candidates. In particular, it has been implemented by several businesses to improve the recruiting process. Recruiters take a long time to review thousands of resumes every day: therefore, Smart Screening can be used to automatically sort out the best candidates. (Osservatorio Artificial Intelligence - Politecnico di Milano)

c. Text Generation

As previously mentioned, Natural Language Processing can be divided into Natural Language Understanding and Natural Language Generation. Text Generation is a Specification strictly related to Natural Language Generation. In

recent years Natural Language Generation is improving its implementation rate and there are more and more applications based on it. Indeed, businesses are looking for software able to both understand text data and then write something about, just like a human analyst would do. (Osservatorio Artificial Intelligence - Politecnico di Milano)
Also here, sub-Specifications can be identified:

- Machine Translation
  The world is covered by a multitude of languages with different sentence structures and grammar. Machine Translation uses Artificial Intelligence algorithms to translate phrases from one language to another one. The real key point is not only directly translating word by word but translating while keeping the meaning of sentences intact along grammar and tenses. One of the most famous examples for Machine Translation is Google Translate, which announced a new machine translation system based on Artificial Intelligence in 2016. (Khurana, Koli, Khatter, & Singh, 2017)

- Content Creation
  Another Natural Language Generation task spreading in recent years is Content Creation. Content Creation refers to the autonomous generation of specific contents with the aim to create more efficient businesses, saving time and resources. More and more organizations use these solutions to generate complex reports from business data in input. Indeed, automated report generation offers business advantages such as increasing efficiency through the elimination of time-consuming and human measurement errors, and allows for deeper analysis. (Osservatorio Artificial Intelligence - Politecnico di Milano)

- Automatic Summarization
  In the digital era, a critical point is the overload of information. This creates problems to recognize and understand important information from large data sets. (Osservatorio Artificial Intelligence - Politecnico di Milano)

As a consequence, Natural Language Processing is relevant to automatically summarize data, while keeping intact the meaning. (Khurana, Koli, Khatter, & Singh, 2017)

d. Language Modelling

Dealing with text, errors regarding the use of syntax, lexicon or morphology can frequently be made. These errors make text less comprehensible and could create problems in other Natural Language Processing applications, such as Machine Translation. Language Modelling has the aim to avoid these mistakes, by correcting them when the person is writing the sentences. Some examples can be auto correctors in smartphones or the grammatical check in Word documents. (Osservatorio Artificial Intelligence - Politecnico di Milano)

Here sub-Specifications are:

- Spell and Grammar Check

  While the Spell Check finds the misspelled words in a text and provides the incorrect words with suitable suggestions, the Grammar Check deals with the most complicated and complex kind of language errors, such as incorrect word order or verb tense errors. An example of an application is WhatsApp or Google Search, using Language Modelling to correct grammar mistakes. (Osservatorio Artificial Intelligence - Politecnico di Milano)

- Words Auto Completes

  Another sub-specification in Language Modelling is the Words Auto Completes, consisting in a software which suggests the rest of the word a user is typing. This technology is used by a lot of search engine companies like Google, exploiting this feature since years. (Osservatorio Artificial Intelligence - Politecnico di Milano)

1.5.3 Virtual Assistant/Chatbot

The Virtual Assistant/Chatbot Class of Solutions includes software agents capable of performing actions and providing services to a human user, based on commands and requests received through written or spoken interaction. (Perego, et al., 2019)

Companies nowadays are more and more interested in the potentialities offered by Chatbots and Virtual Assistants, since they can help in the automation of activities, improve productivity, and increase customer engagement. This is the reason why this application of Artificial Intelligence has received a great deal of attention from media in recent years, especially when it comes to customer service. (Nimavat & Champaneria, 2017)

Particularly, Conversational Bots, Conversational Interfaces, Chatbots, Virtual Assistants and Digital Assistants are all expressions used to refer to the same concept: software applications conversing with a human using natural language text, or speech, to achieve a result. At present, the most visible and familiar achievements in this Class of Solutions are recent voice-driven Virtual Assistants such as Siri, Alexa, Cortana and Google Assistant, but they are just the tip of the iceberg: thousands of text-based chatbots, operating in a variety of channels, such as apps or websites, social media platforms and messenger platforms, and used to help with specific tasks in narrow domains, exist. (Dale, 2016) The above mentioned expressions are frequently confused in the literature, or approached in different ways. (Stieglitz, Brachten, & Kissmer, 2018) Generally, it can be said that the expression "Virtual Assistant", or "Digital Assistant", is used to consider voice controlled applications, such as Apple's Siri and Amazon's Alexa; as a consequence, the expression "Chatbot" is relegated to text based applications, even if part of the existing literature uses the term in a broader way to include both text-based and voice-based Conversational Bots. (Nimavat & Champaneria, 2017)

Chatbots have been around for a long time: after the Turing test in 1950, many researchers tried to create computer programs to create a conversation by using speech or text. The first and most famous chatbot was Eliza in 1966, able to simulate the conversation with a psychotherapist by using simple pattern matching and a rule-based response method. During these early years, several Rule-Based Chatbots have been invented, such as Parry in 1972, Colby in 1975 and Alice in 2000. (Heung-yeung & Xiao-dong, 2018) Therefore, in the past Chatbots were simple rule-based solutions. In contrast, today most Chatbots are endowed with Artificial Intelligence capabilities, allowing to provide more sophisticated and tailored responses to customers. In particular, they leverage Natural Language Understanding and Natural Language Generation to process input queries and generate appropriate sentences as a response. (Nimavat & Champaneria, 2017)

With levels of social media penetration and Internet accessibility expected to increase even more in the future, Chatbots are expected to dominate the market, supported by recent advancements in Natural Language Processing and Artificial Intelligence. (Nimavat & Champaneria, 2017)

1.5.3.1 Classification

To bring order in the wide context of Virtual Assistant/Chatbot solutions, it is required to understand the different typologies of Chatbots and the different functionalities they can provide. Chatbots and Virtual Assistants can be classified using a variety of perspectives, such as their level of interaction or the method of response generation. (Figure 8) (Nimavat & Champaneria, 2017)

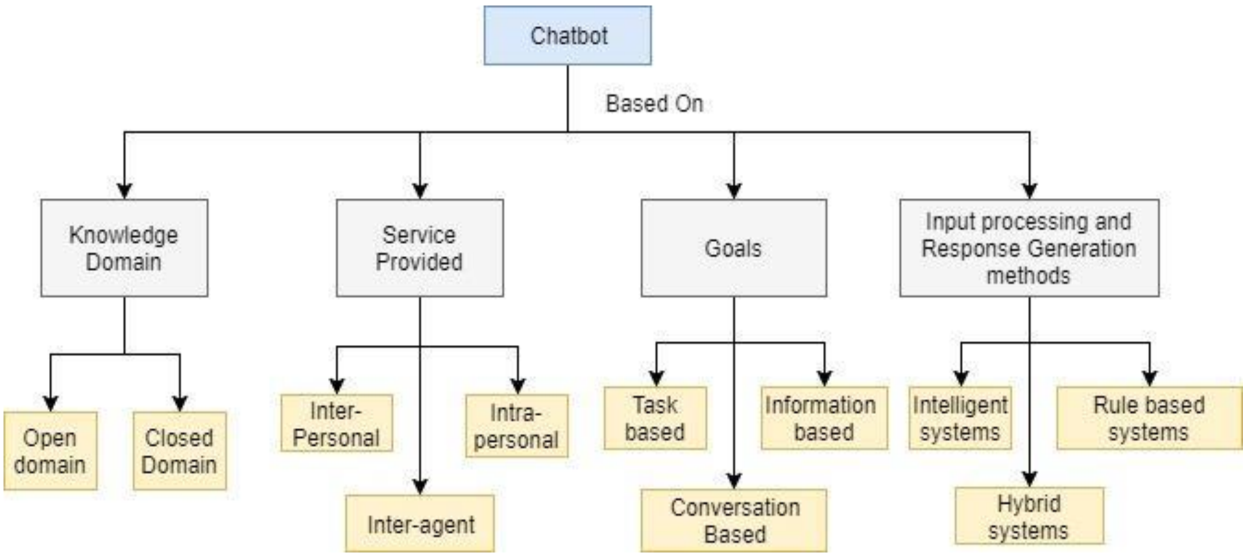

Figure 8: Classification of Chatbot – (Nimavat & Champaneria, 2017)

Based on the considered perspective, Chatbots and Virtual Assistants can be classified in the following categories (Nimavat & Champaneria, 2017):

a. Knowledge Domain

   With this classification, Chatbots are classified according to the knowledge they can access or the amount of data they process. Two main typologies of Chatbots and Virtual Assistants can be identified:

- Open Domain: these conversational bots can talk about general topics and provide appropriate responses to a variety of asked arguments. Examples of Open Domain Chatbots are Siri and Alexa.

- Closed Domain: these are more narrowly focused Chatbots, specialized in a specific knowledge domain and trained for that specific field only. For instance, a restaurant booking Chatbot is only capable of booking tables for the customers and providing them information related to the restaurant. (Nimavat & Champaneria, 2017) (Dale, 2016)

b. Service Provided

When adopting this perspective, Chatbots are classified based on the level of sentimental proximity with the user, that is the level of intimate interaction they are able to establish with him/her. According to the level of sentimental proximity, three different types of Chatbots can be identified:

- Interpersonal: these types of chatbot are not supposed to be companions of the user: they are just expected to act as enablers, getting the needed information and providing them to the users. This does not mean that they cannot have a friendly behaviour, but simply they are not required or expected to do so. Examples of these bots can be Restaurant booking bots, Flight booking bots, FAQ bots.

- Intrapersonal: these chatbots are companions to the user and they understand the user like a human. These types of Chatbots typically exist in the personal domain of the user, like chat apps, and perform tasks in the personal domain of the user, such as managing calendar or storing the user's opinion.

- Inter agent: this typology of bots is based on the idea of two systems communicating each other to achieve a desired goal, and will become dominant in the IOT field. As the diffusion of bots spread, certain Inter-

Agent communication capabilities will be more and more needed. An example of inter-agent communication is the integration of Alexa and Cortana.

c. Goals

With this Classification, Chatbot and Virtual Assistants are grouped based on their primary aim. Three different typologies can be identified:

- Information Based/Informative: this type of bots is typically used to provide information available in a static source. Consider for example a FAQ page of a website, or a warehouse database with inventory entry. Information Retrieval based algorithms are typically used to search for the requested information and provide them as a response.

- Chat Based/Conversational: the aim of these Chatbots is to talk with the user like a human being, correctly responding to the sentences given by the user. (Nimavat & Champaneria, 2017)

- Task Based: this class of bots is typically able to perform a specific task such as booking a flight or helping the client to navigate in an ecommerce website. In many cases, the possible actions are predefined, as well as the possible exceptions. (Nimavat & Champaneria, 2017)

d. Input Processing and Response Generation Methods

This perspective considers the method of processing inputs and generating responses by different Chatbots and Virtual Assistants. Three different typologies of systems can be identified:

- Intelligent Systems: these systems use Artificial Intelligence and Natural Language Processing to understand the phrases of the users and produce the most appropriate responses.

- Rule Based System: these systems use pattern matching and are consequently rigid. They are applicable when the potential outcomes are set, and all the possible scenarios can be considered.

- Hybrid Systems: these systems use a combination of rules and Machine Learning technology to find the right answers to give to the user. An example can be a system that uses a rule-based flowchart to guide the conversation but generates answers by using Natural Language Processing. (Nimavat & Champaneria, 2017)

Particularly, this last classification is especially relevant for the purpose of this thesis, since it considers the potential application of Artificial Intelligence and Natural Language Processing in the Chatbot domain and the shift from simple Rule-Based Chatbots to Hybrid and Intelligent Chatbots.
Then, it must be noticed how Chatbots and Virtual Assistants do not necessarily belong only to one class, but each Chatbot can belong to these categories in varying proportions. For Instance, a Chatbot used in a store may require different typologies of algorithms: it may need some chat capabilities, but also the capability to extract information from FAQs and to search the site to return some product results. To conclude, a Chatbot or Voicebot can be associated with multiple categories, based on the considered perspective: for instance, it can be Informative and Rule Based. (Nimavat & Champaneria, 2017)

1.5.3.2 Stages of Chatbots

Independently from the considered typology, typically Chatbots and Virtual Assistants follow a general architecture based on four sequential stages. (Figure 9)

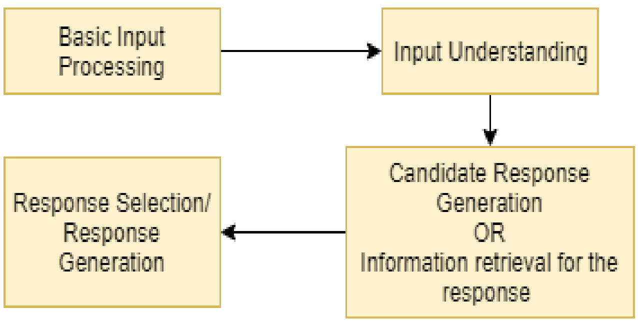

Figure 9: Chatbot Pipeline - (Nimavat & Champaneria, 2017)

The first stage is the Basic Input Processing, in which an input (in the form of voice or text) is simply recognized and processed into an appropriate format. In the following stage, Input Understanding, the input is elaborated, and the needed information is extracted. The next step is the Candidate Response Generation, in which the Chatbot elaborate the information and generate multiple responses to the user input. To conclude, in the Response Selection and Response Generation phase, the most appropriate response is selected and presented to the user. (Nimavat & Champaneria, 2017)

1.5.3.3 Benefits, Adoption and Challenges

Generally, when talking about Chatbots with Artificial Intelligence capabilities, the first thing that comes to mind is the customer care application. This declination of this Class of Solutions is typically the most widespread in companies, enabling them to manage in real time and 24/7 a multitude of customer requests, streamline the process, standardize and free up resources. (Perego, et al., 2019) (Accenture, 2018)

However, considering the application of Chatbots only for customer assistance purposes is reductive, since they have also started to be used for a variety of different purposes.
In particular, the notion of Enterprise Chatbots has progressively emerged: nowadays businesses use Chatbots and Virtual Assistants also internally to interact with their employees, solving problems or automating operations. Enterprise Chatbots can answer questions or perform tasks, supporting companies in the automation of processes and making user employees more efficient around the workplace. (Stieglitz, Brachten, & Kissmer, 2018)

To summarize, Chatbots and Virtual Assistance can have a positive impact on companies in a variety of ways, improving both the customer experience and the employee experience. (Accenture, 2018) In general, several application fields can be identified for Chatbots and Virtual Assistants, such as:

a. Customer Care, referring to the pre/post sale customer assistance.
b. Internal Help-Desk, providing answers to specific problems of the employees of a company.
c. Shop Assistant, providing information on product

d. Corporate Knowledge, providing information on the company, such as its organization and history.
e. Final Product, when the Chatbot is internally integrated in the product and it is one of its main functionalities.
f. Recruiting, enabling a first screening of candidates in a selection process.
g. HR Management, when providing information to HR managers such as the holidays of a worker or the percentage of absent workers.
(Perego, et al., 2019)

Despite the several benefits of Chatbots and Virtual Assistants and their applicability in a variety of fields, many companies are still reluctant to adopt this technology. According to a research published by Accenture in 2018, the most relevant motivations are two: on one side, many organizations believe that users will be reluctant to engage with a Chatbot; in this case, the adoption problem stands in the idea of Chatbot per se. On the other, they are critical about their expected performances, worrying about their ability to incorporate history and context to provide personalized experiences and their capability to adequately understand human input. In conclusion, despite difficulties remain, the future for Chatbots seems to be really bright, since the potential future advantages are too difficult to be ignored. (Accenture, 2018)

1.5.4 Computer Vision

Computer Vision is a branch of Artificial Intelligence that focuses on using algorithms and optical sensors to simulate human visualization, enabling machines to understand and process visual data, such as images and videos, in the same way that humans do. The Computer Vision field investigates how to extract valuable information from images or videos in input, by acquiring, processing and analyzing them to produce a numerical or symbolic information. (Liu, et al., 2017) (Wiley & Lucas, 2018) (Burger & Wheelock, 2015) Specifically, this Class of Solutions brings together solutions for analysing images, individual or in sequence (video), for recognizing people, animals and objects in the image, for biometric recognition and for the general extraction of information from the image. (Perego, et al., 2020)

The Computer Vision discipline can be considered at the crossroad of many different fields, combining techniques, concepts and ideas from digital image processing, pattern recognition, Machine Learning, and computer graphics. (Wiley & Lucas, 2018)
Computer Vision is not a new field, but scientists have been trying to find ways to endow machines with visual capabilities since the early 1960s.
One of the most influential papers in Computer Vision history was "Receptive fields of single neurons in the cat's striate cortex", published in 1959 by David Hubel and Torsten Wiesel, two neuroscientists from the Harvard Medical School. Their publication, result of a long series of experiments, described the responses of neurons in the visual cortex area of a cat while showing it several images, and how its visual experience shaped its cortical architectures. The results of the research described how in the visual cortex area there are both simple and complex neurons; the visual processing starts from detecting simple structures such as edges and corners with simple cells, and then the structures become more and more complex. These findings revolutionized the understanding of the human visual cortex and provided inspiration for Deep Learning, with the later implementation of Deep Neural Networks. In the 1960s Artificial Intelligence became an academic discipline and Papert, a professor at the AI Lab at MIT, launched the Summer Vision Project, with the aim to develop in just one summer a visual system able to solve the problem of machine vision. The project was a partial failure but marked the official birth of Computer Vision as a scientific discipline.

Later in 1982, inspired by the previous work of Hubel and Wiesel, David Marr, a neuroscientist at MIT, published "*Vision: A computational investigation into the human representation and processing of visual information",* a fundamental paper for research in Computer Vision. He introduced the idea of vision as a hierarchical procedure, inspiring the use of low-level algorithms to detect edges, curves and orientations in visual data and their use as inputs for a higher-level understanding of data with subsequent high-level algorithms. Relevant subsequent events mainly follow the footsteps of the Artificial Intelligence history described in Section 1.4, since Computer Vision is considered a branch of Artificial Intelligence and modern Computer Vision systems are mainly powered by Convolutional Neural Networks.

Just to mention the most significant events, the Neocognitron, generally considered the first Convolutional Neural Network, was built in 1979. Ten years later, backpropagation was applied on this architecture (Schmidhuber, 2014) In 2010 The ImageNet Large Scale Visual Recognition Competition was launched, rapidly becoming a benchmark in object detection. It leverages a dataset containing more than one million labeled images belonging to more than 1000 classes. In 2012 a Convolutional Neural Network joined the competition achieving an error rate of 16.4%. This event represented a significant breakthrough for Convolutional Neural Networks. In the following years, Convolutional Neural Networks have always been the winners of the competition, and the error rates fell to a few percent.

As a consequence, the great surge in popularity of Deep Learning over the last years can be considered as largely due to the great leaps it has enabled in Computer Vision: Deep Learning techniques have been shown to outperform traditional Machine Learning approaches in this field, with the use of Convolutional Neural Networks, in particular, very successful in visual understanding tasks such as object detection, facial recognition, image retrieval, semantic segmentation, action and activity recognition, human pose estimation, self-driving vehicles and autonomous robots. (Voulodimos, Doulamis, Doulamis, & Protopapadakis, 2018)

Indeed, despite a large number of Deep Learning methods having emerged, including Deep Belief Networks, Deep Boltzmann Machines and Stacked Autoencoders, Convolutional Neural Networks represents the current dominant approach in Computer Vision, with extensive use in recognition and detection tasks. The main reasons for their current widespread adoption are their ability to automatically learn and extract appropriate features from a set of data and the invariance in image transformation, being insensitive to small image shifts and distortions in the input. (Voulodimos, Doulamis, Doulamis, & Protopapadakis, 2018)

1.5.4.1 Computer Vision Tasks

In this Section the most common Computer Vision tasks, like Object Detection and Face Recognition, are introduced, providing a brief explanation for each one of them. Despite the focus on the most relevant ones, it is fundamental to be aware of the great variety of possible Computer Vision tasks. Just to mention, tasks not presented in the following Section include Event Detection, Instance Recognition, Scene Reconstruction, Object Posture Estimation, Image Restoration, Motion Tracking.

In this assortment, the choice of a specific approach to solve a Computer Vision problem strongly depends on the application domain, as well as the nature of the analyzed data. (Wiley & Lucas, 2018) At the same time, each application requires a specialized adaptation to the Computer Vision algorithm, although it is fundamental to underline how the most advanced Computer Vision applications, such as self-driving vehicles, may also rely on several techniques at the same time, in order to achieve the desired objective. (Burger & Wheelock, 2015)

a. Image Classification

   Image Classification, also known as Category-Level, or Generic, Object Recognition, consists of labeling images in input to categorize them into an appropriate class, such as “cat”, “car” or “bicycle”, assigning a probability to the presence of a particular general class. (Figure 10) (Guo, et al., 2015)

   Prior to adoption of Convolutional Neural Networks, the most adopted techniques were based on bags of visual words that associated a histogram of quantized visual words to an image and used it as input for a classifier. (Guo, et al., 2015) (Szeliski, 2010)

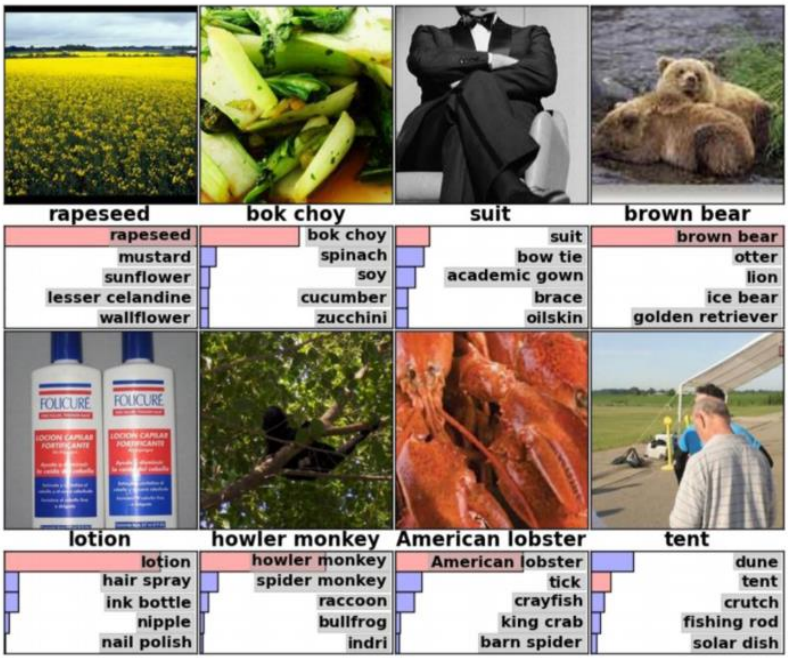

Figure 10: Example of Image Classification - (Guo, et al., 2015)

b. Object Detection

Object detection is a different task, but it is still closely related to Image Classification. In Image Classification the whole image is used as input to return a class label for the objects within the image. In Object Detection, not only a probability for the presence of a given class is returned, but also the position of the instance, or the instances, in the image is provided. The algorithm detects instances of objects across all locations of the image, by predicting the bounding boxes of the objects within the picture, and associate them to a particular class. (Figure 11) (Guo, et al., 2015)

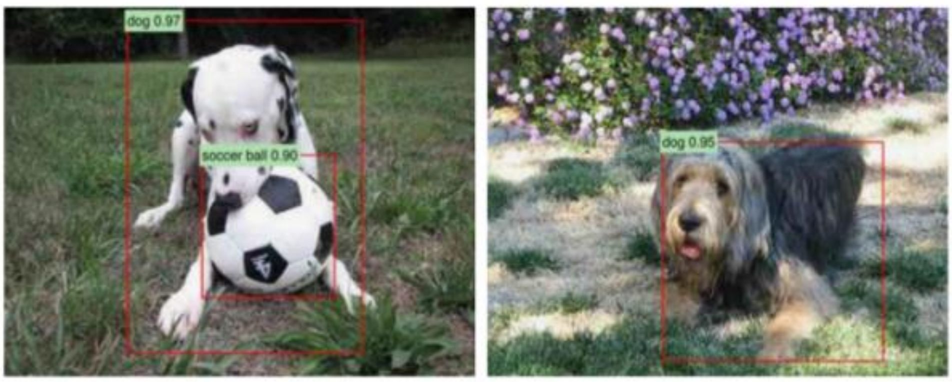

Figure 11: Example of Object Detection - (Guo, et al., 2015)

A detection window is considered as correct if the bounding box of the object has an overlap with the ground truth object, typically more than the 50%. Current common approaches in Object Detection use Deep Learning methods, especially Convolutional Neural Networks, to create a large set of candidate bounding boxes, extract relevant features for each proposal and then feed a classifier to decide if a detection window contains the object or not. Successful practical examples of Object Detection are for example Face Detection, integrated into the majority of current digital cameras to enhance auto-focus, and Pedestrian Detection used to identify pedestrians in automotive safety applications. (Guo, et al., 2015)

c. Image Retrieval

The aim of Image Retrieval, also known as Content-Based Image Retrieval, or Query by Image Content, is to identify images with a similar object or scene as in a query image given in input. (Figure 12) (Guo, et al., 2015)

Figure 12: Example of Image Retrieval (Guo, et al., 2015)

Convolutional Neural Networks are commonly used for extracting features in Content-Based Image Retrieval: for each query image, multiple sub-patches of different sizes and at different locations are extracted, and the distance between reference and query image is set as the average distance between each query sub-patch and the image of reference. (Guo, et al., 2015) Practical applications of this Computer Vision task can be found in Image Search, by using visual features and visual similarity for the retrieval of images from the Web. (Szeliski, 2010)

d. Face Recognition (Biometric Recognition)

Face recognition consists in identifying or verifying the identity of an individual by using images or videos of his face. This Computer Vision task has great commercial interest and a lot of possible practical applications. At present, Convolutional Neural Networks represent the state-of-the-art in Facial Recognition. For instance, both Google's FaceNet and Facebook's DeepFace are based on the use of Convolutional Neural Networks for Facial Recognition purposes. (Voulodimos, Doulamis, Doulamis, & Protopapadakis, 2018)

e. Action and Activity Recognition

This task regards the identification of actions in video sequences.

A series of video frames is received in input, and the system returns as output one or more predefined labels for actions. Both spatial information and temporal information are considered, to consider eventual information within the sequence of frames. Actions commonly recognized include for instance walking, running, dancing, jumping, picking up objects, standing up or sitting down.

Human Action and Activity Recognition has received a great interest in research, with the proposal of several Deep Learning techniques to perform this task. Typically, saliency maps are initially applied for event identification, then Convolutional Neural Networks are used to detect the most relevant frames corresponding to the specific event. (Voulodimos, Doulamis, Doulamis, & Protopapadakis, 2018)

f. Human Pose Estimation

The aim of Human Pose Estimation is to define the position of human joints from images or image sequences. Human Pose Estimation represents a very challenging task because of the variety of human appearances and silhouettes, cluttered background, illumination, and other nuisance factors.

Generally, Deep Learning techniques for Human Pose Estimation can be grouped into two categories based on how images, or image sequences, in input are processed. Holistic methods perform their task at a global level, without detection of individual human body parts and the definition of the spatial relationships among them. On the other side, part-based methods define an explicit model for each individual part and their spatial relationship. (Voulodimos, Doulamis, Doulamis, & Protopapadakis, 2018) (Guo, et al., 2015)

g. Semantic Segmentation

In Semantic Segmentation, algorithms are used to partition an input image into multiple regions, or groups of pixels, to be separately examined. (Guo, et al., 2015) Each identified area corresponds to a group of pixels with some characteristics in common, such as color, texture or gray level. Each region is then associated with a certain object class, with different regions estimated as belonging to different classes. The resulting pixel regions can vary based on the selected Image Segmentation approach. (Wiley & Lucas, 2018) Segmentation is the basis for Object Recognition and Identification, with Convolutional Neural Networks typically used to perform this Computer Vision task. (Guo, et al., 2015)

1.5.4.2 Specifications

To make more order in the Computer Vision area, characterized by a variety of possible tasks, a different perspective can be adopted, in which Computer Vision solutions are grouped into four main Specifications, as proposed by the Artificial Intelligence Observatory of Politecnico of Milan: Image Analysis, Video Analysis, Biometric Recognition and Image & Video Editing. These Specifications allows to provide an easier and more succinct overview of the Computer Vision field, representing a different perspective to include the Computer Vision tasks previously introduced.

To conclude, it should be noted that, despite Computer Vision is at the core of self-driving vehicles and vision-based autonomous robots, these two cases are considered different Classes of Solution, to offer a dedicated focus on these emerging applications.

Hereafter is a brief presentation of the four Specifications in Computer Vision, with examples of possible practical applications.

a. Image Analysis

This Specification refers to the extraction of information from images in input.

An example of a possible practical application is the use of Computer Vision in Healthcare, with systems for the detection of cancerous moles in skin images or the analysis of x ray scans to find symptoms. Image Analysis can be used to detect abnormal tissues or cells in medical images, with application of Deep Learning in the Magnetic Resonance Imaging (MRI) technique for brain tumour detection or in Computer-Aided Diagnosis (CAD) systems for the early detection of breast cancer. (Liu, et al., 2017)

Another possible application of Image Analysis is in the military field, with Image Change Detection. The goal of Image Change Detection is to identify differences between two images of the same scene taken at different moments. For example, this approach is widely used in maritime security monitoring to detect ships in spaceborn images. (Liu, et al., 2017)

b. Video Analysis

This Specification refers to the extraction of information from videos in input. For instance, Deep Learning has been successfully used for motion estimation in video tracking tasks, with possible applications in intelligent video surveillance systems. (Liu, et al., 2017)

c. Biometric Recognition

This specification refers to the notion previously introduced of Face Recognition, consisting in identifying or verifying the identity of an individual by using images of his face. Face Recognition is, among the several Computer Vision tasks, the one in which computers have had the most success. Systems based on Face Recognition are now starting to appear more widely, with practical applications that can be found in identity verification for mobile devices, such as smartphones and tablets. Another example is the use of Face Recognition by social media as Facebook, to detect faces in a photo and associate them with the correspondent identity, allowing users to tag friends. Face Recognition has potential applications also in many other contexts, including parental control, human-computer application, and patent monitoring. (Szeliski, 2010)

d. Image & Video Editing

This Specification refers to the creation or modification of images or videos, by receiving an image or sequence of images in input and returning a new version of it.

For this Specification, a wide variety of Image Editing applications exists. An example is the use of Scene Completion systems in Image Editing: given a photo in which the user wants to erase a large area and fill the gaps, the pixels to fill the gap are taken from similar images to replace a part of the original image. Another possible application is the use of systems to match a picture with a face to a large collection of Internet face images, replacing the face in the original photo supported by intelligent relighting algorithms. At the same time, Deep Neural Networks are used in Photo Restoration to remove noise, such as blurring, from the images. Deep Learning is also applied in Image-Based or Video-Based Rendering, with practical applications in immersive street-level navigation in online mapping systems. (Szeliski, 2010)

1.5.4.3 Adoption and Challenges

According to the "Global Computer Vision Market" report published by Verified Market Research in 2020, the Computer Vision market was valued at USD 13.75 billion in 2019, with projections to reach USD 24.03 billion by 2027, with a CAGR of 7.8% from 2020 to 2027. The stringency in quality assurance of industrial products and the increasing scope of vision-based robotics solutions are factors that have positively contributed to the rise of the global Computer Vision market, while the low adoption of Computer Vision solutions in some sectors might retain the market from growing at a high rate.[1]

In this scenario, several enterprises have applied Computer Vision solutions, often based on Convolutional Neural Networks, to real world applications: after being just a field of study in the past decades, Computer Vision is now starting to be commercialized in several markets, such as Automotive, Entertainment, Manufacturing, Security and Surveillance, Medical, Consumer and Mobile, with different degrees of adoption on the basis of the required application. (Mnih, et al., 2015)

While the technology is the same, the way information extracted from image and video analysis are used strongly depends on the specific application. For example, consider the Computer Vision task of tracking a person: it could be applied either in an airport to track a potential terrorist, or in a retail store to understand the movements of a buyer. Consequently, each application market is characterized by its own specific drivers and challenges. (Mnih, et al., 2015)

The Consumer and Mobile market includes, for instance, a wide range of applications based on consumers and mobile devices, leveraging onboard cameras of smartphones and tablets to deliver Computer Vision-based applications. Examples can be found in biometric recognition, gesture interface, virtual and augmented reality, character recognition. In the Manufacturing industry, Computer Vision systems are used for industrial automation, with cameras for defect detection systems. The Healthcare market too is exploiting the potentialities of Computer Vision, with systems to help the automation of tasks such as detecting cancerous moles in skin images or detecting symptoms in x-rays. Security and surveillance is another relevant industry, with solutions for government, retail chains and other interested subjects to detect humans, faces, animals or cars. (Mnih, et al., 2015)

To conclude, a brief overview about relevant challenges in the Computer Vision field is provided.

First, the notion of human-level vision can either be related to filling the performance gap between Computer Vision and human vision or integrating into Deep Learning systems new insights from human brain research.
Human vision has a remarkable ability in Computer Vision tasks, both in basic visual representations and in case of changes with occlusion, background variations and occlusion. In the last years, Convolutional Neural Networks have sometimes exceeded performance of human raters, but it is not enough to conclude that these computer architectures are able to compete with the human brain. For example, the difference between current Convolutional Neural Networks and human vision emerges when considering images that are completely not recognizable by humans, but computer architectures believe they contain a certain object class with a 99.99% confidence. Therefore, one major challenge at present is represented by the accuracy and correctness of the outcomes. (Schmidhuber, 2014)

Another relevant challenge could be represented by the limited availability of training data. This shortage could limit the ability of Convolutional Neural Networks to learn or their size, especially if obtaining already labeled information is difficult.
Consequently, two solutions are currently used to improve the available training data.

On one side, more training examples are produced from existing data by using data augmentation techniques, such as color casting, rotating, scaling or lens distortion. On the other, more training examples can be collected by applying weak learning algorithms on weakly labeled image sets. However, it is fundamental to identify new techniques for producing or collecting a wider range of training data, to make the network learn better features invariant to changes (Schmidhuber, 2014)

1.5.5 iRPA

This Class of Solutions refers to the execution of recurring tasks by a software integrated with Artificial Intelligence capabilities (Perego, et al., 2019): Intelligent Robotic Process Automation, or iRPA, is able to achieve flexible and intelligent automation by combining Robotic Process Automation (RPA) and Artificial Intelligence. (Zhang C. , 2018)
This second wave of RPA, in which the potentialities of RPA are amplified by the integration with Artificial Intelligence, is typically referred as “Intelligent” Robotic Process Automation, but many different names exist, such as Intelligent Automation, Cognitive Automation, Process Intelligence or Process Orchestration, all of them referring to the same phenomenon. (Perego, et al., 2019)

In particular, Artificial Intelligence and RPA are two separated technologies handling different tasks. (Zhang C. , 2018)

The notion of Artificial Intelligence has already been widely discussed in Section 1.1 of this thesis. Artificial Intelligence has been defined as “the branch of computer science that studies the development of hardware and software systems with specific capabilities typical of humans, able to autonomously pursue defined objectives, making decisions that previously were only made by humans”. (Perego, et al., 2019) In other words, Artificial Intelligence refers to the development of computer systems, methods and tools able to perform complicated and intelligent tasks commonly requiring human-like intelligence, such as visual perception, reading mails, chat and messaging, speech recognition, language translation or decision making. (Joseph, Dutta, Gillard, & Seewooruttun, 2018)

On the other side, Robotic Process Automation (or RPA) is a business process automation technology, allowing enterprises the automation of high volume routines. RPA technology is used to capture the execution of highly repetitive tasks previously carried out by human employees in their user interface and automate it by using software robots, or SW robots. These software robots are able to imitate and replicate the execution of routine tasks performed by human users in their computer system interface, replacing them in the execution of those processes involving routine work. (Agostinelli, Marrella, & Mecella, 2020) Routine tasks automated by software robots are typically office tasks such as extracting structured data from files, making calculations, copying and pasting information across the cells of a spreadsheet, opening mails and attachments. (Willcocks, Lacity, & Craig, 2017) These tasks are distributed along a wide variety of use-cases, such as customer onboarding, data migration and entry, payroll automation and accounting reconciliation. (Rughi & Ceriani, 2019)

RPA has two main differences if compared with traditional workflow automation tools, such as Enterprise Resource Planning (ERP) systems. Firstly, in RPA tools the software is developed through a user-friendly design interface, and therefore it does not require traditional programming skills. Then, the same software is "non-invasive", by means that it interacts with the existing IT systems through the same user interfaces used by human workers, allowing the RPA technology to run application software mimicking the way a person works with that software. (Zhang C. , 2018) (Willcocks, Lacity, & Craig, 2017)
This facilitates the integration of different systems and makes it non-invasive. (Zhang C. , 2018)

The advent of RPA opened eyes on how business processes can be automated, freeing people from highly repetitive routine tasks. Despite a growing attention and evident benefits, it is more and more recognized that traditional RPA tools need to be integrated with Artificial Intelligence to enlarge the scale and scope of process automation. (Burnett, 2017)
Indeed, conventional RPA creates numerous opportunities for automation, but also has its own limits.
A first difference between RPA and Artificial Intelligence stands in the typology of data they can process: RPA can process only structured and semi-structured data, while Artificial Intelligence can also work with unstructured data, such as scanned document, PDF, text or web contents. (Burnett, 2017) Then, traditional RPA is driven by predefined

rules defined by expert users, and so it can only automate simple and predictable rule-based processes and produce a single correct answer as output. On the other side, the learning ability of Artificial Intelligence systems allows the automation of tasks using inference-based processes, returning a set of likely outcomes. (Zhang C. , 2018) (Willcocks, Lacity, & Craig, 2017) (Agostinelli, Marrella, & Mecella, 2020)
The presence of these limitations imply that RPA can be used only for the automation of some processes, or more often sub-processes, limiting the potential for end-to-end process automation. In particular, there is a compelling case for using Artificial Intelligence in processes to automate unstructured parts of those processes and to learn, enabling the system to handle non-standard situations based on previous experience. (Burnett, 2017)

As a consequence, since mid-2017 organizations had progressively begun to combine RPA with capabilities provided by Artificial Intelligence systems to overcome these limits, enabling end-to-end process automation for any kind of process. (Joseph, Dutta, Gillard, & Seewooruttun, 2018) (Zhang C. , 2018)
The combination of these two technologies identify a new domain, namely Intelligent Robotic Process Automation, (or iRPA), in which RPA and Artificial Intelligence are combined to drive end-to-end process automation. (Zhang C. , 2018)
To avoid possible misconceptions, it should be noticed that RPA is not an old technology being replaced by Artificial Intelligence, but the two technologies should be considered as highly complementary: Artificial Intelligence can increase the value in RPA solutions, and vice versa. For instance, consider credit-risk modelling: this process could be supported by Artificial Intelligence only, but it is with the combined use of Artificial Intelligence and RPA, and so iRPA, that the insights delivered from Artificial Intelligence can be immediately actioned. Or vice versa, RPA can be considered as the oxygen feeding data into Artificial Intelligence. (Joseph, Dutta, Gillard, & Seewooruttun, 2018)

In particular, the different subfields and techniques of the Artificial Intelligence discipline, such as Natural Language Processing, Computer Vision, Optical Character Recognition or Voice Recognition, are fundamental for the extension of RPA capabilities. (Burnett, 2017)

The possible application fields of iRPA are a multitude, including both back-office activities (e.g. data entry of anagraphic information for new customers, crosscheck of financial information, mail delivery to clients when their orders have been processed) and front-end ones (e.g consultation of several systems when in line with a client, management of incoming customer emails) (Righetti & D'Aquino, 2018) (Burnett, 2017)

1.5.5.1 iRPA Interpretative Framework

To better describe the possible levels of integration between Artificial Intelligence and RPA, the AI Observatory of Politecnico of Milan has developed an interpretative framework to classify RPA applications and solutions. The framework defines three possible levels of integration between the two technologies. (Perego, et al., 2019)

a. Programmed RPA: RPA is not combined with Artificial Intelligence. The RPA is intended as traditional deterministic automation, working with structured data and commonly applied to repetitive high volume processes. (Perego, et al., 2019)

b. AI Assisted RPA: RPA is combined with Artificial Intelligence capabilities to improve efficiency in the execution of some tasks (for instance, by handling non programmable exceptions) or to add new capabilities to the solution. At this level, Artificial Intelligence is therefore included in some parts of the process, supporting the automation of tasks that cannot be handled by Programmed RPA. (Perego, et al., 2019)

c. AI Driven RPA: there is a strong integration between RPA and Artificial Intelligence. Artificial Intelligence guides the processes, optimizing routes and creating new connections between the tasks. Since the role of Artificial Intelligence moves from supporting individual robotized activities to the management and control of the process, at this level we can speak of "Intelligent Business Process Management". (Perego, et al., 2019)

1.5.5.2 Main Benefits

In this section, a brief overview of the main benefits coming from the implementation of iRPA is provided. These benefits are valid also for traditional Programmed RPA, but it is in the integration with Artificial Intelligence that RPA is able to realize its full potential. iRPA offers a multitude of advantages. Among them, the most relevant ones are:

a. Improved Time Efficiency: the execution speed of processes increases thanks to iRPA's implementation, reducing the process time of tasks. By substituting human employees in the execution of daily tasks, iRPA tools are able to run processes faster and, therefore, more effectively. (Rughi & Ceriani, 2019) (Reddy, Harichandana, Alekhya, & Rajesh, 2019)

b. Greater Productivity: the used resources are optimized and the volumes in output increases, due to the adoption of iRPA tools. (Rughi & Ceriani, 2019) (Reddy, Harichandana, Alekhya, & Rajesh, 2019) (Righetti & D'Aquino, 2018)

c. Larger Accuracy: by eliminating human errors from the process, the risk of errors in the execution of tasks is eliminated too. Considering that even the most careful employee can make occasional mistakes, the implementation of iRPA solutions allows to achieve performance levels unattainable by human beings. (Rughi & Ceriani, 2019) (Reddy, Harichandana, Alekhya, & Rajesh, 2019) (Righetti & D'Aquino, 2018)

d. Cost Savings: thanks to iRPA adoption, operative costs decrease. On one side, a software robot costs around ⅓ of a FTE off-shore and ⅕ of a FTE in-house. On the other, despite the upfront investment required, the increases in efficiency and productivity and the elimination of human errors largely justify the investment. (Righetti & D'Aquino, 2018) (Reddy, Harichandana, Alekhya, & Rajesh, 2019)

e. Greater Employee Satisfaction: since human workers have not boring, repetitive tasks to perform in their daily routine anymore, the satisfaction and retention levels typically increase. Moreover, employees can be reallocated on more stimulating added-value tasks, using their skills and knowledge in projects driving innovation and growth. (Berruti, Nixon, Taglioni, & Whiteman, 2017) (Rughi & Ceriani, 2019) (Righetti & D'Aquino, 2018) (Reddy, Harichandana, Alekhya, & Rajesh, 2019)

f. Better customer experience: because of the higher accuracy and efficiency in process execution, the customer experience improves as well. (Righetti & D'Aquino, 2018)

1.5.5.3 Example of Practical Applications

To provide a better understanding of the notion of iRPA and its functioning, here two different examples of possible practical applications for iRPA are provided.

a. E-commerce website audit refers to a typology of operational audit, in which an internal audit team is for instance in charge of verifying whether a certain product ranks or not at the top in a keyword search on a selected e-commerce platform. These ranking information are later used to evaluate the product performances and to provide recommendations for possible improvements. While the verification of the product ranking is typically done manually, iRPA can be used to automate this process.

   Here the primary workflow consists in providing an outcome "TRUE" or "FALSE" for each keyword searched on a digital platform, in which "TRUE" means that the product appears on the first page of results after searching the keyword, and "FALSE" that it does not appear. The outcome is then entered in a dedicated column of a spreadsheet.

   The secondary workflow consists in opening the web browser, entering the desired e-commerce platform, type one by one keywords from a list of useful keywords, and for each one of them understand if the product is present in the front page of the website or not. In particular, recognizing the eventual presence of a product requires human vision capabilities, since detection of the product logo is required.

   As a consequence, iRPA can be used to streamline this boring and repetitive process, leveraging Computer Vision to recognize the eventual logo of the product. The software robot mimics the actions of a human employee to open the web browser, search the desired e-commerce platform and enter in the website, type each keyword from a keywords list, press enter and understand if the logo is present on the first page of results or not. To conclude, it accordingly enters the expression “TRUE” or “FALSE” in a dedicated column of a spreadsheet. (Zhang C. , 2018)

b. Another possible example is represented by the application of iRPA in pension audit, in which a team of auditors typically analyze pension plans and extract key terms from them. These ones are then used to feed the Employee Benefit Plan (EBP) audit. In case of iRPA implementation, RPA can be used to collect the pension plans, organize them and deliver them to a Natural Language Processing module through which key terms are identified. RPA is then used to enter the extracted key terms in Excel spreadsheets and execute test queries for substantive testing. Notable items are then analyzed by Machine Learning modules to select suspicious items to be further analyzed by auditors. After suspicious items are investigated, auditors can complete the audit and prepare its conclusion. Lastly, a report of the audit can be created using Natural Language Generation. (Zhang C. , 2018)

1.5.5.4 Adoption & Challenges

According to a research published by Deloitte in 2019, the global market of Intelligent Automation is valued around USD 8,2 billions in 2019, and is expected to grow with a CAGR of more than 10% in the next five years, reaching around USD 14,39 billion in 2025. (Rughi & Ceriani, 2019)

iRPA has allowed the realization of many business cases within several industries. Use cases can be found in accounting, billing management, data and document validation, inventory list updating and much more. (Reddy, Harichandana, Alekhya, & Rajesh, 2019) Despite the use of this technology is expanding into a multitude of industries, the BFSI sector (Banking, Financial Services and Insurance) is the one with the faster rate of adoption. (Righetti & D’Aquino, 2018) Companies in these industries have been embracing RPA for a while, but now its combination with Artificial Intelligence is

attracting a great interest because of the potential achievable benefits. (Joseph, Dutta, Gillard, & Seewooruttun, 2018)
This dominance can be reconducted to the high volume of back office tasks, the necessity to execute processes by respecting norms, rules and standard, and a great availability of structured data, that makes this sector particularly suited to benefit from iRPA implementation. (Righetti & D'Aquino, 2018) (Perego, et al., 2019)

If we focus on the Italian market, according to a report published by the AI Observatory of Politecnico of Milan in 2020, the adoption of solutions combining Artificial Intelligence and RPA is still limited. Indeed, 41% of adopted RPA solutions are still traditional Programmed RPA. The 54% refers to AI assisted RPA, while AI Driven RPA is just the 5% of cases. (Perego, et al., 2020) Similar data can be found in a research about the Italian market published by Deloitte in 2019: 4 companies out of 10 have already started an automation journey, and among them the 55% is including Artificial Intelligence into its strategy of Intelligent Automation. In addition, another 25% plans its utilization within the end of 2022. (Rughi & Ceriani, 2019)

To conclude, despite the remarkable benefits of iRPA, a common concern is that robots will substitute human employees and will result in layoff. In reality, the purpose of this technology is to support humans in doing their work and to allow the reallocation of human workers on more stimulating and relevant tasks. (Reddy, Harichandana, Alekhya, & Rajesh, 2019)

1.5.6 Recommendation System

This Class of Solutions includes all those solutions aiming to address the preferences, interests or more generally the decisions of a user based on the information that he provides, directly or indirectly. The objective is to deliver personalized recommendations that can be placed at different points of the customer journey or the decision-making process. (Perego, et al., 2020)
Recommendation Systems can be defined as programs aiming to recommend the most suitable items, products or services, to certain users, individuals or businesses, by predicting their interest in items based on information about the items themselves, the users, and the interactions between the items and the users. (Jannach & Jugovac, 2019)

Particularly, Recommendation Systems can be considered one of the most visible and successful practical applications of Artificial Intelligence. (Jannach & Jugovac, 2019)
Nowadays they represent an area of huge interest, because of the wide variety of practical applications helping users to deal with information overload by suggesting products, services or contents tailored to their interests. Some well-known applications for these systems can be found in products recommendations of Amazon.com and film recommendations for Netflix (Adomavicius & Tuzhilin, 2005) (Burke, Felfernig, & Göker, 2011) A multitude of Recommendation Systems for e-commerce websites has been developed, providing tailored purchasing recommendations to online individual customers. (Lu, Wu, Mao, Wang, & Zhang, 2015)

The use of Artificial Intelligence and Machine Learning techniques have greatly contributed to advance the state-of-the-art in Recommendation technologies, in contrast with the early systems based on heuristics. (Adomavicius & Tuzhilin, 2005) At present, Bayesian techniques, Artificial Neural Networks and Markov Decision Processes are often used to build effective Recommendation Systems. (Lu, Wu, Mao, Wang, & Zhang, 2015)

As previously introduced, the most common application for Recommendation Systems is in an e-commerce scenario. For instance, a person interested in literature can visit his favourite online bookstore and find in the homepage, among the others, a tailored list of recommended books. This list may include books published by his favourite author, related to his personal interests or other books that for a certain reason he could be interested to read. (Burke, Felfernig, & Göker, 2011)
The capability to offer a personalized view of data to the user, the bookstore's inventory in this case, is a key characteristic of recommendation systems. Their purpose is to reduce the search effort of the user and to point out products that he might be most likely to buy, with relevant benefits for both the user and the owner of the e-commerce shop. (Burke, Felfernig, & Göker, 2011)

1.5.6.1 Classification of Recommendation Systems

Based on the used recommendation method, Recommendation Systems are generally classified in three main categories. Content-Based Recommendation Systems make recommendations based on the previous choices of the user, recommending items similar to the ones he chose in the past. In Collaborative Recommendation Systems, the user is recommended with items that people with similar interests and preferences chose in the past. To conclude, Hybrid Recommendation Systems are based on the combination of these two methodologies. (Adomavicius & Tuzhilin, 2005) To provide a better understanding of Recommendation Systems, these three typologies of Recommendation Systems are briefly explained in this Section.

a. Content Based Method

   In Content Based Recommendations, the utility of an item for a user is estimated based on the utilities of similar items that the user chose in the past. For instance, if considering a movie recommendation application, a Content-Based Recommendation System recommends movies to the user by identifying similarities, such as the actors, the director, the genre, with the movies that the same user has rated highly in the past. Therefore, he will be recommended with similar films coherent with his preferences. (Adomavicius & Tuzhilin, 2005) (Lu, Wu, Mao, Wang, & Zhang, 2015)

   Particularly, improvements over traditional approaches come from building user profiles collecting information about their tastes, interests and needs. The information necessary to build user profiles can be collected both explicitly and implicitly, using questionnaires or learning from the user behaviours over time. (Adomavicius & Tuzhilin, 2005)

   Opposed to conventional approaches, Machine Learning techniques such as Bayesian classifiers, decision trees and Artificial Neural Networks are frequently applied in Content Based Recommendation Systems. (Adomavicius & Tuzhilin, 2005)

   Despite the advantages, this recommendation methodology is also characterized by some inherent limitations. For instance, over specialization refers to the fact that the system is not able to present a variety of alternatives to the user, but just a homogenous set of options similar to what the user has previously chosen: by defining the utility of an item only based on the previous choices of the user, the system will only recommend items similar to the ones rated highly by him. Therefore, a user with no experience in Chinese cuisine would never be recommended with a Chinese restaurant. A typical solution to this problem is the introduction of some randomness. (Adomavicius & Tuzhilin, 2005)

Moreover, this methodology requires a new user to rate a certain number of items, before the system can adequately understand his preferences and generate accurate recommendations. Therefore, the Recommendation System will not be immediately effective with new users. (Adomavicius & Tuzhilin, 2005)

b. Collaborative Method

Collaborative Recommendation Systems estimate the utility of an item for a user based on the utilities assigned to the same item by similar users. For instance, in the movies recommendation application, an user will be recommended with a film that other users with similar preferences and interests in movies have ranked highly. Also in this case, in contrast with traditional heuristic-based approaches, model-based algorithms have emerged, leveraging Machine Learning techniques, especially Artificial Neural Networks, to build a model for making ratings prediction. (Adomavicius & Tuzhilin, 2005)

This typology of Recommendation Systems does not have the problem of Over Specialization of Content Based Recommendation Systems: formulating recommendations based on the ratings of different users, they enable suggestions different from what a user has seen in the past. Anyway, also in this case shortcomings can be identified, such as the new user problem previously mentioned with Content-Based Recommendations: to generate effective recommendations, the system must first learn the interests of the new user based on his choices. One possible solution is the use of Hybrid Recommendation Systems, later described. Another problem with Collaborative Systems is the one of sparsity. For instance, in movie recommendations, typically there is a small number of movies rated by a lot of people, while the majority of them are rated by few people. Therefore, these films are recommended very rarely, even if their ratings are very high. Moreover, for users with unusual interests, it is more difficult to find similar users with common tastes. Also for the sparsity problem a variety of solutions have been proposed, such as the use of profile information for evaluating similarity: for example, a film can be recommended to a user also considering if it has been rated highly by users with the same age, geographical origin, gender or education. (Adomavicius & Tuzhilin, 2005)

c. Hybrid Methods

Hybrid Recommendation Systems combine the best features of the two recommendation methods previously introduced, allowing to overcome some of their limitations. (Lu, Wu, Mao, Wang, & Zhang, 2015)

A wide variety of Hybrid Recommendation Systems are used, classified on the basis of how they combine Collaborative and Content Based methods. For instance, they can implement Collaborative and Content Based methods separately and then combine their predictions, or create a unique model integrating Content Based and Collaborative characteristics. Other possible options are the integration of some Content Based characteristics into a Collaborative approach, or the integration of Collaborative characteristics into a Content Based approach (Adomavicius & Tuzhilin, 2005)

1.5.6.2 Specifications

According to the framework proposed by the Artificial Intelligence Observatory of Politecnico of Milan, also for this Class of Solutions four main Specifications can be identified:

a. Purchasing Recommendation

This Specification considers the use of Recommendation Systems to suggest a user the purchase of products or services aligned with his interests. For instance, in the online catalogue of e-commerce websites, products viewed or searched are used to suggest complementary or related items.

b. Content Recommendation

This Specification considers the use of Recommendation Systems to suggest a user contents aligned with his interests. For instance, the posts that a user visualizes on Facebook, movies recommendations by Netflix or video suggestions by YouTube are based on Recommendation Systems based on Artificial Intelligence.

c. Online Advertising

    This Specification has been introduced to consider the broader variety of ways through which a company can use the Internet to provide promotional marketing messages to consumers. For instance, Recommendation Systems can be used to send clients emails with other recommended articles following a purchase.

d. Dynamic Pricing

    This Specification considers the internal use of Recommendation Systems to suggest the selling price for products or services, based on the current market demand.

1.5.6.3 Benefits and Adoption

Due to advancements in recommendation methodologies and techniques, more and more practical applications for Recommendation Systems have been developed. (Lu, Wu, Mao, Wang, & Zhang, 2015) At present, they have found widespread application across a multitude of sectors and contexts. (Thompson, 2015)

Particularly, they have been largely applied in the areas of e-commerce, enabling commerce websites to provide personalized product recommendations. Also in TV, music, and movie services, they allow to analyze the choices of the customers and use them to provide customized recommendations for television programs, songs and films. (Thompson, 2015)

To conclude, Recommendation Systems represent a source of value for both consumers and the businesses. (Jannach & Jugovac, 2019)
From a consumer perspective, Recommendation Systems allow to reduce the information overload by narrowing the set of possible choices, but also to find interesting articles and discover new options. (Jannach & Jugovac, 2019) (Thompson, 2015)
From a business perspective, companies are not used to publicly share detailed information about their benefit from the implementation of Recommendation Systems. Anyway, the use of Recommendation Systems to support Cross Selling and Up Selling can lead to notable increases of Sales and Revenues. Just think that in 2006, for instance, Amazon declared that the 35% of their sales are due to cross-sales recommendations.

(Jannach & Jugovac, 2019) Other benefits can be found in a stronger customer relationship, an increased client trust and loyalty, higher click-through and conversation rates, new possibilities for promotion and a better understanding of the customer. (Thompson, 2015)

1.5.7 Autonomous Robot

This Class of Solutions includes robots able to move themselves or some of their parts, manipulate objects and perform a variety of different actions without requiring human intervention, collecting information from the surrounding environment, and adapting to unforeseen or coded events. (Perego, et al., 2020)

Solutions belonging to this class are finding application in a surprising multitude of contexts and are extremely different in terms of characteristics, functionalities, and aspects, ranging from humanoid robots, to collaborative robots, to in-store assistants. Autonomous Robots can be used both by corporations, to automate and speed up the processes creating competitive advantage, and by people, to receive support in their daily routine. Artificial Intelligence is at the core of Autonomous Robot, allowing to make them more automated and able to interact with the external environment. (Hinds, Roberts, & Jones, 2011)

1.5.7.1 Examples of Practical Applications

Due to the wide variety of possible application fields and the diversity of solutions, it is impossible to provide a single unifying perspective about this Class of Solutions. Therefore, some relevant applications for Autonomous Robot are briefly described in this Section, to provide a better understanding of this Class of Solutions.

For instance, in the retail area the technology of social service robots has matured considerably in recent years, with the consequent introduction of commercial applications in many locations such as shopping malls. Humanoid robots, robots with the physical appearance of a person, can support the customer decision making process, helping them to make better use of their time during the shopping experience. (Niemelä, Heikkilä, & Lammi, 2017) Autonomous Robots can also work as shopping assistants, carrying the shopping with robotic arms and using Natural Language Processing

capabilities to interact with the users through speakers. (Iwamura, Shiomi, Kanda, Ishiguro, & Hagita, 2011)
Interesting use cases can also be found in the medical field: Autonomous Robots endowed with Artificial Intelligence capabilities are very popular in this context, and several applications are possible. For instance, a promising use of Autonomous Robots is in care assistance, with robot companions for elderly patients affected by cognitive decline or limited mobility. Another interesting application is the use of Autonomous Robots in surgery, either as assistants or as solo performers. For instance, Da Vinci is widely used for complex urologic and gynaecologic procedures. (Hamet & Tremblay, 2017)

Also in an industrial context, Autonomous Robots ranging from Collaborative Robots to robotics arms can help organizations to increase automation, improve production efficiency and reduce errors. While full robot autonomy factories are at the moment out of reach for the foreseeable future, the adoption of collaborative robots for coworking between people and robots is becoming a spread reality in industrial businesses: human beings and robots can share the same workspace in factory setups, home and office environments, and perform various object manipulation tasks in a collaborative manner. (Kragic, Gustafson, Karaoguz, Jensfelt, & Krug, 2018) Collaborative robots are expected to work side by side with human beings, supporting and collaborating with employees in the execution of different tasks. (Hinds, Roberts, & Jones, 2011)
An example of Autonomous Robots in this industry is the one developed by ABB, using Artificial Intelligence and Computer Vision to recognise and grasp targeted objects. (Kragic, Gustafson, Karaoguz, Jensfelt, & Krug, 2018)

1.5.7.2 Challenges

Autonomous Robots have also raised some concerns that is essential to consider for the future development of this technology, often depending on their specific capabilities and application domain. For instance, a typical concern about the use of robots for the elderly is that such use goes against their preferences, since they prefer to receive services from humans rather than cold machines. Another problem is that their use might increase their isolation from other people. (Winfield, McDermid, Müller, Porter, & Pipe, 2019)

Some general concerns about the adoption of Autonomous Robots are:

- Bias: Autonomous Robots using Artificial Intelligence can have bias in their decision making.
- Deception: humanoid robots present the risk of emotional attachment or dependency especially to vulnerable users.
- Employment: the introduction of Autonomous Robots might displace certain classes of workers.
- Oversight: the ability to oversee or govern is an issue, since operators are typically not able to understand and manage the behaviour of the system for which they are responsible.

  (Winfield, McDermid, Müller, Porter, & Pipe, 2019)

1.5.8 Autonomous Vehicle

One of the most relevant applications of Artificial Intelligence and Machine Learning are intelligent Autonomous Vehicles. (Zhang, et al., 2016)

The last decade brought significant advancements in the field of Autonomous Vehicles, or AVs, and self-driving cars, largely because of the progresses in Artificial Intelligence, and especially Deep Learning approaches. (Grigorescu, Trasnea, & Macesanu, 2019)
Artificial Intelligence algorithms are indeed widely applied in Autonomous Vehicles for a variety of purposes, including perception of a driving scene, identification of the appropriate navigation path, detection and classification of objects, behaviour arbitration and motion control. (Grigorescu, Trasnea, & Macesanu, 2019) (Bimbraw, 2015)

In particular, according to the definition proposed by the Artificial Intelligence Observatory of Politecnico of Milan, Autonomous Vehicles are "autonomous means of transport used to transport people, animals or things, either driving on the roads (vehicle) or navigating in the sea, lakes or rivers (vessel), or even flying in our atmosphere or in space (aircraft), able to perceive the external environment and identify the correct manoeuvres required to adapt to that environment. (Perego, et al., 2019)
Therefore, these intelligent unmanned vehicles are systems capable of autonomously taking decisions without any kind of human intervention, with the aim to improve the efficiency of transportation systems and road safety. (Zhang, et al., 2016)

1.5.8.1 Self-Driving cars

Among this Class of Solutions, the most evident application of Autonomous Vehicles can be found in autonomous, or self-driving, cars.

Autonomous cars are autonomous decision-making systems, able to process streams of environmental data collected by a variety of on-board sensors, such as radars, LIDARs, cameras, and GPSs units, with the aim to make driving decisions. (Grigorescu, Trasnea, & Macesanu, 2019)

Most of research and development in the autonomous driving field is done through research programs based on the collaboration between car manufacturers, research centres and equipment manufacturers. (Gruyer, et al., 2017) Car manufacturers are making enormous efforts trying to develop a new generation of automobiles, characterized by a high autonomy level, and not requiring human intervention anymore. (Zanchin, Silva, Santos, & Linares, 2017)

In particular, the Society of Automotive Engineers (SAE) introduced a classification model to describe the degree of autonomy for self-driving cars, with six possible levels of autonomous driving. The higher the level, the higher the autonomy of the vehicle, the functions it is able to perform and its capabilities, and the lower the requirement for driver involvement. (Zanchin, Silva, Santos, & Linares, 2017)
These six levels are organized in a scale from 0 to 5. The lower levels 0,1 and 2 of driving automation are characterized by the mere presence of basic ADAS (Advanced Driving Assistance Systems), eventually warning the driver or directly intervening if needed. The driver at these automation levels is still responsible for the driving task. The next levels exhibit increasing degrees of automation, till arriving at the full automation of level 5: at this level, the driving system is completely autonomous and human intervention is not required at all. (Grigorescu, Trasnea, & Macesanu, 2019) (Gruyer, et al., 2017)

In particular, at Level 0 there is no automation; the driver performs the entire Driving task. (Gruyer, et al., 2017)

At Level 1 system, or driver assistance, the driving system performs only the longitudinal or the lateral motion control under certain conditions. Anyway, the two tasks are never simultaneously performed by the system, and the driver is responsible for the remaining task. (Zanchin, Silva, Santos, & Linares, 2017) At this level, the main functionalities of the vehicle are typically the Active Cruise Control (ACC), allowing to maintain a selected speed while keeping distance from the vehicles ahead, and the Lane Departure Warning System (LDWS), able to recognize the line marking on the road and alerting the driver if he unintentionally changes the lane.
These functionalities are enabled by the use of an Ultrasonic sensor, exploiting sound propagation for object detection, a Long-Range Radar sensor, using electromagnetic waves, and a camera for surround, combining four-to-six fish lenses for object detection. (Zanchin, Silva, Santos, & Linares, 2017)

At Level 2, or partial driving automation, under certain conditions the driving system performs both the longitudinal and the lateral motion control, with the driver just supervising the execution of these tasks or performing other ones. In this case, one of the main functionalities added to Level 1 is Lane Keep Assist (LKA), moving from just warning the driver to autonomously intervene in cases of unintentional lane change and driver not responding to alerts. In other cases, these systems are used to support the driver in staying at the centre of the lane, by acting on the steering system. Another functionality typically introduced at level 2 is Park Assist (PA), with a multitude of different systems currently available.
Again, the technologies commonly used for implementation of these functions are Ultrasonic sensors, Long-Range Radar sensors, Short-Range Radar sensors and the camera for surround. (Zanchin, Silva, Santos, & Linares, 2017)

At Level 3, or conditional driving automation, under certain conditions the driving system is able to perform the driving task in low-speed, stop-and-go freeway traffic; however, the driver is still expected to keep attention and be ready to take over in case of failure. (Gruyer, et al., 2017) (Zanchin, Silva, Santos, & Linares, 2017)
Here the main added functions are the Automatic Emergency Braking (AEB), the Driver Monitoring (DM) and the Traffic Jam Assist (TJA).

The Automatic Emergency Braking supports the prevention of accidents, by warning the driver in case of a dangerous situation or autonomously braking if the situation gets worse. The Driver Monitoring uses infrared sensors or video cameras to monitor the biological and psychological state of the driver, eventually warning him if needed or directly acting on the vehicle to reach a safe condition. To conclude, the Traffic Jam Assist supports driving in traffic conditions, by enabling the driving system to autonomously brake, start, keep a safe distance from other vehicles and so on.
At this level the complexity of technology increases, with the addition of a Long Distance Camera, a Stereo Camera based on stereo vision, LIDAR (Light Detection and Ranging) sensors and eventually thermal sensors. (Zanchin, Silva, Santos, & Linares, 2017)

Level 4, or high driving automation, is much more challenging. At this level, under certain conditions the driving system is able to perform the entire driving task in motorways or freeways, but differently from Level 3 it must also be able to autonomously make emergency manoeuvres in critical scenarios, without any human intervention. (Gruyer, et al., 2017) (Zanchin, Silva, Santos, & Linares, 2017)
At this level there are no major changes in the technology; the main difference stands in the algorithm controlling the driving system and in the capability to combine the different sensors to have a more precise understanding of the environment surrounding the car. (Zanchin, Silva, Santos, & Linares, 2017)
To conclude, at Level 5, or full driving automation, the driving system is able to perform the entire driving task under every possible condition and without any human involvement, making emergency manoeuvres if needed and taking the best decisions to guarantee the highest possible safety level. (Gruyer, et al., 2017) At this level, the driving system is capable of driving the car throughout an entire distance, independently from starting or arriving point, road status, weather and traffic conditions. (Zanchin, Silva, Santos, & Linares, 2017)

Once introduced these six possible levels of driving automation, it must be remembered how Artificial Intelligence and Machine Learning are a central component for the development of both Advanced Driver Assistance Systems (ADAS) and fully automated vehicles.

When talking about ADAS, we are referring to the functionalities previously introduced, such as Adaptive Cruise Control, Parking Assist, Lane Keeping Assist, Lane Departure Alert, Emergency Brake Automation. A modern ADAS architecture is typically a combination of four different components: longitudinal control, lateral control, parking assistance and, driving vigilance monitoring system, each including a variety of possible functions. (Moujahid, et al., 2020)

Machine Learning algorithms have a major role in ADAS: for instance, consider the use of cameras to enable lane detection. Another example is the application of Computer Vision techniques and Convolutional Neural Networks to control the state of the driver: monitoring facial features and early signals of stress, anger or fatigue are detected and potential traffic accidents avoided. Deep Learning, and especially Computer Vision techniques, are also used for detection of obstacles, road lanes, pedestrians. (Moujahid, et al., 2020)

Artificial Intelligence in self-driving vehicles is strongly relying on Computer Vision methods, but not only: human vision systems constituted by front and back cameras are always combined with several sensors like Radars, LIDARs, GPS units, accelerometer, steering angle and wheel speed. The data gathered through these sensors cannot be processed without an intelligent system using Machine Learning algorithms to analyze them and make the most appropriate decisions. (Moujahid, et al., 2020)

The role of Artificial Intelligence and Deep Learning becomes even more relevant when moving to basic ADAS to full driving automation. For example, perception and understanding of the surrounding environment are fundamental requirements for a fully autonomous vehicle, especially in urban areas where a multitude of objects and visual occlusions may be present. Deep Learning, and particularly Convolutional Neural Networks, are relevant for the analysis of 2D images collected by cameras and 3D point clouds collected by LIDAR sensors, enabling to perform several Computer Vision tasks such as Object Detection and Recognition, Semantic Segmentation and Semantic Localization. (Grigorescu, Trasnea, & Macesanu, 2019)

Another key component of self-driving vehicles is localization, aiming at calculating the pose of the vehicle, and so its position and orientation, while on the road. Deep Learning methods are used to improve the accuracy of vehicle localization, through vision-based analyses of consecutive video frames or LIDAR intensity maps. (Grigorescu, Trasnea, &

Macesanu, 2019) Another example is the use of Deep Learning to support understanding of the environment surrounding the car: typically, occupancy maps are used to map the environment by dividing it into cells and computing an occupancy probability for each one of them. Also in this context, Deep Learning plays a key role in making probabilistic predictions for occupancy maps (Grigorescu, Trasnea, & Macesanu, 2019)

1.5.8.2 Benefits of Self-Driving Cars

A successful implementation and deployment of autonomous vehicles is expected to generate several benefits in the future. In this Section, the most notable ones are briefly summarized.

The most evident advantage resulting from self-driving cars' implementation would be a decrease in road accidents: road safety represents a major concern for the automotive industry, with the number of people dying in road accidents increasing year after year. Most of car collisions is due to human mistakes, including distractions and long reaction time. (Moujahid, et al., 2020) The number of traffic injuries and collisions is expected to decrease with the advent of self-driving cars, because of a higher reliability and shorter reaction compared to human drivers. (Bimbraw, 2015)
At the same time, Autonomous Vehicles are expected to reduce traffic congestions, by enabling a better management of traffic flows and improving mobility in overcrowded urban areas. (Bimbraw, 2015) The transit time is expected to decrease also because self-driving cars are supposed to navigate at higher speeds while keeping the safety level high. At the same time, the ride experience of the passengers is expected to improve, using the time previously devoted to the driving task for doing other activities or relaxing. (Bimbraw, 2015) (Zanchin, Silva, Santos, & Linares, 2017) In the future, also the problem of scarce parking spaces could become a thing of the past, with Autonomous Vehicles able to drop off the car's occupants at destination, autonomously start looking for a free park or going home, and return back later to pick up passengers. (Grigorescu, Trasnea, & Macesanu, 2019) Furthermore, also the necessity of road signs is expected to decrease, since self-driving cars will receive all the needed information via network. (Bimbraw, 2015)

1.5.8.3 Specifications

As the other Classes of Solutions, also the Autonomous Vehicle class can be organized in different Specifications, according to the framework proposed by the Artificial Intelligence Observatory of Politecnico of Milan, to provide a better overview of the category and easily classify the possible applications. The proposed Specifications are:

a. Advanced Driving Assistance Systems, or ADAS. This Specification allows the consideration of autonomous driving systems corresponding to the lower levels of the SAE classification, including ADAS making use of Artificial Intelligence techniques.

b. Autonomous system along a defined path. To better understand this Specification, the broader notion of Autonomous Vehicles, going beyond the narrow idea of self-driving cars, must be considered. As previously mentioned, this Class of Solutions includes all the “autonomous means of transport used to transport people, animals or things, able to perceive the external environment and identify the correct manoeuvres required to adapt to it”. (Perego, et al., 2019) Consequently, autonomous systems along a defined path are, for instance, all those autonomous transportation systems used to handle inventories in a well-defined area like the warehouse of a company.

c. Autonomous system along a non-defined path. As opposed to the previous Specification, this one includes autonomous driving systems able to move along non predefined paths, enabling the consideration of the higher levels of the SAE classification. Autonomous systems along a non-defined path are, for instance, self-driving cars or trucks with a high level of autonomous driving automation, rather than aerial or underwater autonomous transportation systems.

1.5.8.4 Adoption & Challenges

Since 2010, research and development in the field of Autonomous Vehicles entered a new phase, because of a growing interest by car manufacturers, startups, equipment manufacturers and IT organizations. (Zhang, et al., 2016)
Mercedes-Benz, BMW, Volkswagen, Ford, and several other companies launched research & development programs for self-driving cars. Chinese companies too, ranging from car manufacturers such as BYD, Yutong and SAIC, to the search engine giant Baidu, heavily invested in this field. Even Google started to test its own autonomous car on the road in 2015. (Zhang, et al., 2016) In this dynamic context, several organizations have developed their own autonomous driving prototype, to enable tests relevant for further research and developments. (Gruyer, et al., 2017)

Despite the recent advancements in autonomous vehicles, relevant obstacles still exist and need to be overcome to achieve fully autonomous vehicles. (Zhang, et al., 2016) This is the reason why, at present, only a limited number of automobile manufacturers, such as Hyundai, Mercedes, Tesla, Volkswagen and Volvo is able to offer partial driving automation, selling consumer vehicles typically classifiable as SAE Level 2 or 3. Therefore, fully Autonomous Vehicles still require notable efforts by research laboratories, and it is reasonable to estimate that the first driving systems enabling full automation will be not available on the market before 2025. (Gruyer, et al., 2017)
On the other side, many of the ADAS previously introduced are already sold on the mass market, with a recent shift from only high-end vehicles to mainstream vehicles too, even if some mid-range vehicles are still waiting for implementation of these functionalities. (Gruyer, et al., 2017) (Moujahid, et al., 2020)

Furthermore, despite the previously mentioned benefits of self-driving cars, many disadvantages and obstacles also exist, determining challenges in the autonomous driving field. Some of these challenges are related to the technology or technical obstacles in implementation, hindering the development of fully Autonomous Vehicles. Other ones are general concerns about Autonomous Vehicles that could limit the adoption of this technology.
Here some of these challenges are briefly summarized:

a. Perception ability: Autonomous Vehicles require the ability to understand the surrounding environment, enabling a safe navigation in the driving scene. Deep Learning has allowed accuracy progresses in detection and recognition of objects, but currently available systems are typically able to perform these Computer Vision tasks for just a limited number of trained object classes. Future methods should allow us to instantaneously detect and track a wider variety of objects. (Grigorescu, Trasnea, & Macesanu, 2019)

b. Real-time computing and communication: constraints of real-time computing and communication are present, to immediately process the wide variety of data collected from sensors and take the most appropriate decisions. These requirements can be met thanks to the progress in specific computational hardware developed for Autonomous Vehicles and 5G communication. (Grigorescu, Trasnea, & Macesanu, 2019)

c. Availability of training data sources: despite simulation environments have been used to fill the gap between the scarcity of training examples and the needed data, the gap between simulated and real world driving still exists. In particular, training data are scarce for corner cases, such as hazardous situations leading to car accidents, implying the necessity to design methods able to generalize and learn from limited training data (Grigorescu, Trasnea, & Macesanu, 2019)

d. Cybersecurity: another threat is represented by possible cyber-attacks at the computer or communication systems of self-driving vehicles. This may pose a risk for use in criminal activities and terrorist attacks. (Bimbraw, 2015) (Zanchin, Silva, Santos, & Linares, 2017)

e. Reliability of the system: additionally, the reliability of autonomous driving systems, with the subsequent implications on safety, represents a point of debate. The implementation of self-driving vehicles is affected by a barrier of uncertainty between the autonomous vehicles themselves and the safety of pedestrians and car’s occupants. (Zanchin, Silva, Santos, & Linares, 2017)

f. Regulation: another challenge is represented by the need to develop a regulatory framework and establish laws for self-driving cars, to determine for instance who is to be considered liable in case of car accident: if the owner of the car, the automobile manufacturer or the government. (Zanchin, Silva, Santos, & Linares, 2017) (Bimbraw, 2015)

g. Impact on employment: a common concern is that the adoption of self-driving vehicles will cause job losses in driving-related jobs (Zanchin, Silva, Santos, & Linares, 2017) (Bimbraw, 2015)

1.5.9 Intelligent Object

This Class of Solutions refers to objects able to perform actions and take decisions without requesting human intervention, interacting with the surrounding environment through sensors (e.g. thermometers, cameras, microphones) and actuators (e.g. open/close door/windows, switching on of domestic appliances or installation) and learning from the habits and the actions of the people interacting with them. (Perego, et al., 2019)

Specifically, this Class of Solutions is intended to make room for all those applications in which Artificial Intelligence algorithms, and particularly Deep Learning ones, run directly on the hardware of smart objects and IoT devices, avoiding the elaboration of data on dedicated cloud platforms. Conversely, all those solutions in which the sensory data collected or generated by Internet connected devices are elaborated through Artificial Intelligence algorithms on cloud platforms are classified following the Classes of Solutions previously introduced, based on the utilized techniques.

The main reason for having a dedicated Class of Solution is that nowadays the Internet of Things, or IoT, represents a very hot topic: this is due to the huge number of devices with sensing capabilities gathering or generating sensory data for a variety of aims and applications. (Mohammadi, Al-Fuqaha, Sorour, & Guizani, 2018)

In particular, the notion of IoT is related to having a worldwide network of uniquely addressable and interconnected physical objects, with all the things in the world

connected to the internet to provide a seamless integrated system. The IoT can be defined as an "interconnection of sensing and actuating devices providing the ability to share information across platforms through a unified framework, developing a common operating picture for enabling innovative applications. This is achieved by seamless ubiquitous sensing, data analytics and information representation with Cloud computing as the unifying framework." (Gubbi, Buyya, Marusic, & Planiswami, 2012)
Previously, the expression Internet of Things was coined by Peter Lewis at the Congressional Black Caucus 15th Legislative Weekend Conference in 1985, describing the IoT as "the integration of people, processes and technology with connectable devices and sensors to enable remote monitoring, status, manipulation and evaluation of trends of such devices". (Saha, Mandal, & Sinha, 2017)

The expression IoT was initially used to refer to RFID (Radio-Frequency Identification) devices, a now mainstream technology consisting in a circuit for computation and an antenna for communication, used for the automatic identification of objects. (IOTPROP) The notion of IoT strongly evolved over time, with progresses in cloud computing, Machine Learning, wireless networking, control systems and embedded systems, till arriving at the current concept of IoT. At present, IoT applications have expanded into different fields of domestic life and businesses, with smart wearables, smart home devices, smart security solutions, smart health monitoring devices, and the notion of Smart City. (Saha, Mandal, & Sinha, 2017)

Particularly, one of the most relevant characteristics of most of these applications is the use of intelligent learning algorithms for prediction, pattern recognition, classification, data mining or generally data analytics. (Mohammadi, Al-Fuqaha, Sorour, & Guizani, 2018) Intelligence has become a key point for IoT applications, and for this reason researchers have focused on making things in the IoT as intelligent as possible, instead of simply following predefined rules to make decisions. (Tsai, Lai, & Vasilakos, 2014)

The necessity to use intelligent learning algorithms is mainly due to the complex and challenging task of analysing the huge amount of data collected and produced by IoT systems, going beyond the capabilities of traditional learning methods. Progresses and advancements in Machine Learning enabled the complex Big Data analytics and knowledge extraction required by IoT systems. (Mohammadi, Al-Fuqaha, Sorour, & Guizani, 2018)

In particular, Deep Learning is the Machine Learning approach most commonly used to overcome this challenge, with Deep Learning techniques as the basis of many current applications in the IoT domain. (Lane, Bhattacharya, Georgiev, Forlivesi, & Kawsar, 2015)

When considering IoT Data, they have characteristics and properties that make them different from general Big Data, implying different requirements for their analytics. IoT Data are characterized by the following properties (Mohammadi, Al-Fuqaha, Sorour, & Guizani, 2018):

a. Large-Scale Streaming Data: a great number of IoT Data acquisition devices are utilized in IoT applications, generating continuous streams of huge volume data.
b. Heterogeneity of Data: IoT devices capturing data typically collect a wide variety of different information, implying data heterogeneity.
c. Time and space correlation of Data: IoT devices are typically placed in a certain location. Consequently, each sensory data item will have a corresponding location and time-stamp.
d. High noise Data: being so small pieces of data, in IoT applications data may be subject to errors during collection or transmission.

In particular, according to the desired application, IoT Data can either be accumulated as a source of Big Data or streamed continuously. (Mohammadi, Al-Fuqaha, Sorour, & Guizani, 2018)

IoT is one of the main sources of Big Data, expression through which we refer to large datasets that conventional hardware and software systems are typically not able to store, handle and analyze. The identification and extraction of meaningful insights from large amounts of input data is the main purpose of Big Data Analytics.

In particular, insights coming from IoT Big Data can be used to guide decision making even many days after the same data are generated. This is the reason why Cloud computing is typically used for IoT Big Data Analytics. (Mohammadi, Al-Fuqaha, Sorour, & Guizani, 2018)

Beyond Big Data Analytics, IoT Data also requires a different new class of analytics, generally referred to as Fast and Streaming Data Analytics. Streaming data are data collected or generated in a very short time interval and typically require to be quickly analysed, to immediately extract insights and enable real-time or near real-time decisions. As a consequence, insights coming from Streaming Data Analytics should be provided in a very short time, ranging from hundreds of milliseconds to few seconds.
Consider, for instance, applications like self-driving cars, advanced driver assistance systems, fire detection or checking the consciousness or health of drivers and older people. In all these situations, the collected data require to be rapidly processed since strict real-time actions are needed. Many other IoT applications leveraging Artificial Intelligence, such as human pose detection for smart home applications, require treating data with a fast analytic approach, instead of accumulating them for later analysis. (Mohammadi, Al-Fuqaha, Sorour, & Guizani, 2018)

Also in the case of Fast and Streaming Data Analytics, approaches based on cloud infrastructures have been proposed by researchers. Anyway, the transfer of sensory data to the cloud for analysis and the return of an answer are characterized by a certain time delay that makes this kind of approach usually not suitable for the required applications. (Mohammadi, Al-Fuqaha, Sorour, & Guizani, 2018)

This is the reason why several time-stringent applications, like the ones previously mentioned, require Fast Data Analytics close to the source of data, and so at the system edge, or even at the source of data, and so on the IoT device itself, allowing the elimination of excessive delays. (Mohammadi, Al-Fuqaha, Sorour, & Guizani, 2018)
As a consequence, Fast and Streaming Data Analytics is brought closer to the source of data to enable rapid data processing. For this reason a new research area has emerged, focused on bringing analytics to resource-constrained hardware and enabling real-time analytics. In addition, the “Intelligent Object” Class of Solution has been introduced, to take into account an emerging number of smart objects that directly support a form of light-weight intelligence. (Mohammadi, Al-Fuqaha, Sorour, & Guizani, 2018)

1.5.9.1 Specifications

As for other Classes of Solutions, also in this case several Specifications can be identified, enabling a better overview and classification for Intelligent Objects applications.
Following the framework proposed by the Artificial Intelligence Observatory of Politecnico of Milan, the main Specifications are:

a. Wearables, such as bracelets, rings, belts, shoes, and glasses.

b. Home Devices, such as smart speakers, lightbulbs, televisions.

c. Smart Cameras, for both public and private use, and domestic or business application.

d. Viewers

e. Other Smart Objects, including all the objects not covered in the previous Specifications

1.5.9.2 Adoption & Challenges

According to a report published by Verified Market Research in 2020, the global IoT Market Size was valued at USD 212.1 Billion in 2018, and it is expected to grow 25.68% in the period 2019-2026, reaching USD 1,319.08 Billion by 2026[2].
Recently, IoT applications have found a great expansion into several industries and different fields of domestic life and businesses, with practical applications in smart wearables, smart home devices, smart security solutions, smart health monitoring devices, smart city, and smart agriculture. (Saha, Mandal, & Sinha, 2017)

However, it should be remembered that the "Intelligent Object" Class of Solutions has not been introduced to describe the general rise of IoT technology and the related use of Artificial Intelligence algorithms to make sense of IoT data, but to provide a focus on the emerging number of applications in which intelligent algorithms are hosted on the same IoT devices. In this sense, the adoption of these solutions is still limited.

The main reason is that analytics performed at the IoT device level is still characterized by many challenges and limitations: at present, Deep Learning algorithms are rarely used in IoT systems' hardware because of constraints in terms of computing, memory and storage capacity, power resources. (Mohammadi, Al-Fuqaha, Sorour, & Guizani, 2018) (Lane, Bhattacharya, Georgiev, Forlivesi, & Kawsar, 2015) These are the main reasons why the adoption of Deep Learning within IoT hardware is still facing relevant barriers, making it impossible for a variety of models to be directly executed on the hardware and leading to a subsequent use of cloud platforms. (Lane, Bhattacharya, Georgiev, Forlivesi, & Kawsar, 2015)

In this context, different techniques have been proposed to overcome existing barriers. On one side, Deep Neural Networks are made more lightweight to be hosted at the device level, using Deep Neural Networks compression, approximate computing for Deep Learning algorithms, Tinymotes with Deep Learning and accelerators. On the other side, these progresses have been accompanied by the emergence of advanced hardware solutions. Anyway, despite getting closer and closer to real-time analytics on IoT devices, this still remains a clear open challenge. (Mohammadi, Al-Fuqaha, Sorour, & Guizani, 2018)

Because of these challenges and barriers, the number of well-developed applications in this Class of Solutions is still limited. Anyway, it is expected that once these challenges are overcome, the amount of IoT applications with Deep Learning algorithms directly hosted on the IoT device will rise.

1.6 The Future of Artificial Intelligence: Research Directions and Challenges

To conclude this Literature Review, after having introduced the main theoretical concepts behind Artificial Intelligence and the nine Classes of Solutions proposed by the Observatory, this concluding Section is intended to provide information about future research directions for Artificial Intelligence and challenges that the technology is facing.

Indeed, despite the successes of Machine Learning in the recent years, it must be noticed that this remains a young field with several opportunities to explore.
Some of them can be easily identified by comparing the differences between the way an Artificial Neural Network and the human brain process information, rather than differences between the learning procedure of Machine Learning systems and the one of natural learning systems such as humans and animals. (Jordan & Michell, 2020)

Just to mention few of these opportunities: consider, for instance, how Machine Learning systems typically learn one specific function or data model from a unique data source, while people learn a wide variety of skills and knowledge with different training experiences along their life, using both Supervised and Unsupervised Learning and performing tasks in an increasingly difficult sequence (for example, humans are first of all learning to crawl, then to walk and only later to run). This is the reason why a current area of research regards the construction of computers able to continuously learn skills and functions for years without interruption, improving their ability to learn a new skill because of the previously learned skills. (Jordan & Michell, 2020)

Another opportunity under investigation is the concept of team-based mixed-initiative learning. Consider how people with different backgrounds and skills often operate in teams to collect and analyze relevant data, while Machine Learning systems typically analyze information in isolation. New Machine Learning systems may be developed to tightly collaborate with human beings in data analysis, combining their ability to detect patterns in massive amounts of data with the ability of human experts to provide plausible explanations and generate new hypotheses. (Jordan & Michell, 2020)

To conclude, despite the relevant improvements in performance that GPU-based implementations allowed, current GPUs require much more energy than biological brains; these are characterized by neurons efficiently communicating by using brief spikes and remaining often quiet. Consequently, energy-efficient hardware is under development, to allow on top construction of Deep Neural Networks with similar biologically inspired spiking neurons. However, these Deep Neural Networks with spiking neurons are at present still not able to compete with conventional Deep Neural

Networks. (Schmidhuber, 2014)

On the other side, it is relevant to consider how Machine Learning, as any powerful technology, is not only opening up a lot of potential opportunities, but also a considerable amount of challenges and critical questions.
For instance, Artificial Intelligence plays a key role in the development of autonomous intelligent robots and intelligent software robots. If on one side these technologies are perceived as an opportunity, part of the popular opinion also considers them as potential threats to human labour. Despite the fear of robots substituting human workers and creating an employment problem, these technologies can lead to rethinking human-machine collaboration or to the reallocation of human workers on more stimulating value-added tasks. (Rughi & Ceriani, 2019) (Reddy, Harichandana, Alekhya, & Rajesh, 2019) Or, just to mention another example, the value-driven push in collecting more and more personal data has generated several privacy issues. (Jordan & Michell, 2020)

Moreover, questions are emerging about who should have access to data, ownership of them and take advantage of them. At present, many data are collected by companies and internally used to raise profits, with little or no reasons for sharing them outside of the organizational boundaries. Despite these data are privately owned, the profits and advantages that the whole society could realize if they were made available for public good are inestimable. (Jordan & Michell, 2020)

These reflections suggest that Artificial Intelligence will probably be one of the most transformative technologies of this century. Even if it is impossible to understand what will happen in the future, it is essential for society to start reflecting on how to maximize the benefits coming from this disruptive technology. (Jordan & Michell, 2020)

2. Research Questions and Methodologies

The aim of this Chapter 2 is to introduce the research questions this thesis work is intended to answer, as well as to provide information about the research methodologies followed to answer them. In Section 2.1 the research questions are introduced, while section 2.2 provides information about the followed methodologies.

2.1 Research Questions

The aim of this Section 2.1 is to introduce the research questions at the basis of this thesis. To better understand the scope of this research, it must be noticed that this thesis work is supported and supervised by the Artificial Intelligence Observatory of Politecnico of Milan, born in 2017 with the aim of answering to the growing interest of public and private companies in the opportunities offered by Artificial Intelligence. Each year a research aiming to analyse Artificial Intelligence applications at a global level is commissioned, to provide information about the level of adoption of Artificial Intelligence solutions by companies operating all over the world.

Indeed, although the presence of an extensive body of literature related to Artificial Intelligence techniques, information about the adoption of Artificial Intelligence in enterprises is limited, and mainly referred to past years. Moreover, Artificial Intelligence is a rapidly growing technology, and its enterprise adoption is rapidly evolving, implying the necessity to provide a constantly updated picture of its evolution. Therefore, the primary contribution of this thesis consists in investigating the current state of adoption of Artificial Intelligence solutions at international level, answering the following research question:

RQ1: What is the present state about the international adoption of Artificial Intelligence solutions?

To answer this research question, several perspectives need to be considered, such as the level of diffusion and maturity of the nine Classes of Solutions proposed by the Observatory, each one with its own Specifications, or the geographical adoption of Artificial Intelligence solutions. Furthermore, information should be provided about the involved sectors and the application fields of these solutions, for instance if they are internally used by organizations or if they are enablers of the products and services sold by companies. To conclude, it is necessary to understand if technological partners are involved in Artificial Intelligence projects, as well as their typology, and eventual benefits or critical issues related to the implementation of Artificial Intelligence solutions. The main results for this research question are presented in Chapter 3.

Meanwhile, the broad applicability of Artificial Intelligence is currently enabling companies in different sectors to adopt it at advantage of their specific business needs. In this regard, the secondary aim of this thesis work is to provide a focus on how Artificial Intelligence is currently being adopted in some specific industries, with a sectoral analysis to answer the following research question:

RQ2: How Artificial Intelligence adoption changes from industry to industry and what is the potential contribution for the sector?

Particularly, the Food & Beverage, Manufacturing, Retail and Banking & Finance industries have been analysed. The main results for these research questions are shown in Chapter 4.

2.2 Research Methodologies

To provide an answer to the research questions, a precise methodology has been adopted.

a. Existing Literature Review: firstly, the existing literature has been reviewed to achieve an appropriate level of comprehension of the theoretical concepts behind the technology, as well as to gain knowledge about the nine Classes of Solutions introduced by the Artificial Intelligence Observatory of Politecnico of Milan, at the basis of the adoption of Artificial Intelligence in a business context.

b. Analysis of Secondary Sources: an analysis of secondary sources on the web has been carried out, to identify active Artificial Intelligence initiatives of companies. The search for information has been performed with the Google Chrome research engine. To be rigorous and effective in the search for information, a well-established methodological scheme has been adopted, allowing to identify for each analysed company the corresponding Artificial Intelligence projects. At the end, 1089 Artificial Intelligence initiatives of 165 companies have been identified by applying this methodology. The methodology adopted for the analysis of secondary sources is presented into detail in Section 2.2.2.

c. Census of the Projects: the identified initiatives have been used to feed a database, specifically created for the purposes of the research. Each of the initiatives has been categorized according to specific criteria presented in Section 2.2.3. The database has been later used as basis for analysis, to provide answers about the identified research questions.

d. Development of a Qualitative Framework to Support Sectoral Analysis: finally, to help answer *RQ2*, a Qualitative Framework has been specifically developed. The model is organized into three different Levels (AI-enabled Products & Services; AI in Customer-Oriented Processes; AI in Enterprise-Oriented Processes), each of them including possible application fields of Artificial Intelligence solutions based on proximity with the specific customer of the industry. Therefore, the model has been applied to the Food & Beverage, Manufacturing, Retail and Banking & Finance industries to analyse the way companies operating in these sectors are applying Artificial Intelligence solutions. At this regard, the model has been used to support both the analysis of how Artificial Intelligence is being applied within a certain industry and a structured and orderly presentation of results.

These Research Methodologies have been applied to a sample including the first 235 public companies rated in the 2020 Forbes Global 2000 list, with the aim to identify any active Artificial Intelligence project of these firms. Information on the 2020 Forbes Global list and how it has been drafted are provided in Section 2.2.1. In the sample, a cap of maximum 20 companies analysed per industry has been placed, to obtain a homogeneous overview of the sectors in which projects are developed and avoid biases in the database. By excluding vendors of Artificial Intelligence solutions and companies exceeding the maximum number of 20 firms per industry, 169 companies have been analyzed.

2.2.1 Forbes Global 2000

The Forbes Global 2020 is an annual ranking list of the top 2000 public companies in the world. The ranking list is done by taking into account 4 different parameters for each company: sales, profit, assets and market value.

In a 2018 article, Forbes discloses the general methodology used to rank companies in the list[3].

Firstly, they create 4 separate lists of the 2000 biggest companies for each parameter: 2000 companies ranked by sales, 2000 companies ranked by profit, 2000 companies ranked by assets and 2000 companies ranked by market value. Each list has a minimum value in order to be qualified, and a company is eligible to enter in the Global 2000 if it respects the minimum value for at least one parameter. Each company receives then a score for each parameter based on the ranking in the related list; if the value for a parameter is under the minimum, it receives a zero score for the specific list.

After that each score is added up for all the four parameters, and a new list with the aggregate score is created: this is the final Forbes Global 2000 list.

2.2.2 Analysis of Secondary Sources

Secondary Sources have been used and analysed to find information about any active Artificial Intelligence initiative of companies under analysis. The information needed to map the adoption of Artificial Intelligence solutions have been so collected online, paying attention in using secondary sources of proven reliability, such as the website of the company and its other information channels, news websites, interviews, verified online articles. Particularly, to provide a reliable picture of the current adoption of Artificial Intelligence solutions, only information from 2018 onwards for Operative, Pilot and implementation projects has been considered, and from 2019 onwards for Project Proposals (Project Proposals of 2018 or before, without any available information in most recent years, have been considered as being discarded).

To be rigorous in the search for information a methodological scheme has been adopted, to identify for each company the corresponding Artificial Intelligence projects. The research has been performed with the Google Chrome research engine, is partially organized according to the Classes of Solutions, and imply for each company the typing of the following words:

a. Name of the company + Artificial Intelligence (search results are consulted up to page 5)
b. Name of the company + Machine Learning (search results are consulted up to page 5)
c. Intelligent Data Processing
    - Name of the company + Artificial intelligence + “Forecasting” (up to page 1)
    - Name of the company + Artificial intelligence + “Classification” (up to page 1)
    - Name of the company + Artificial intelligence + “Pattern” (up to page 1)
    - Name of the company + Artificial intelligence + “Optimization” (up to page 1)
d. Chatbot
    - Name of the company + Artificial intelligence + “Chatbot” (up to page 1)
    - Name of the company + Artificial intelligence + “Assistant” (up to page 1)
e. Natural Language Processing
    - Name of the company + Artificial intelligence + “Natural language processing” (up to page 1)
f. Computer Vision
    - Name of the company + Artificial intelligence + “Computer Vision” (up to page 1)
g. iRPA
    - Name of the company + Artificial intelligence + “RPA” (up to page 1)
    - Name of the company + Artificial intelligence + “Intelligent Automation” (up to page 1)
h. Intelligent Object

- Name of the company + Artificial intelligence + "Intelligent Object" (up to page 1)

i. Recommendation System
   - Name of the company + Artificial intelligence + "Recommendation" (up to page 1)

j. Autonomous Robot
   - Name of the company + Artificial intelligence + "Robot" (up to page 1)

l. Autonomous Vehicles
   - Name of the company + Artificial intelligence + "Vehicle" (up to page 1)

Following this methodology, Artificial Intelligence initiatives have been identified for each of the analysed companies and used to feed the database for later analysis of results.

2.2.3 Census of the Projects

The identified initiatives have been used to feed a database, specifically created for the purposes of the research, categorizing each of the identified initiatives based on specific criteria. Once completed, the database has been used as basis for analysis, providing relevant indications about the identified research questions. Consequently, the database can be intended as an instrument to represent the state-of-the-art adoption of Artificial Intelligence solutions: inside it, all the relevant information regarding Artificial Intelligence projects developed by the analysed public companies can be found.

Below, the different columns (the criteria used to categorize projects) constituting the database are presented, accompanied by a brief explanation when necessary.

a. Source

   It indicates the source of reference for mapping a specific Artificial Intelligence project.

b. Ranking

   It indicates the position of the considered company in the Forbes Global 2000.

c. Sales

   It indicates the sales of the considered company (in USD millions)

d. Company

Name of the company

e. Industry

It specifies the industry of the company. The possible industries considered for the database are the following:

- Aerospace & Defence
- Agriculture
- Automotive
- Banking and Finance
- Chemical
- Construction & Real Estate
- Energy, Resources & Utility
- Food & Beverage
- Healthcare & Assistance
- Hospitality
- Insurance
- Transportation & Logistics
- Manufacturing
- Media & Entertainment
- Pharma
- Retail
- Technology
- Telco

f. Country

Here the country in which the project is carried out is selected (Worldwide if the project is adopted all over the world).

g. Continent

Corresponding continent (Africa, Asia, North America, South America, Europa, Oceania; Worldwide if the project is adopted all over the world)

h. Description

Brief description of the project

i. Status of the Project

In this section the level of progress in the project is specified:

- Project Proposal: the project is still at an early stage; there is an idea and an allocated budget, but the project has not been implemented yet. In this phase there are no concrete results yet, since the project is still in an embryonic stage.
- Pilot: the idea is developed on a small scale. The first concrete results are shown, or successful tests are carried out.
- Implementation: the project is currently being released on a large scale
- Operative: the project is completely operative and a continuous improvement process for it is in place.

j. Reference Year:

It indicates the year associated with the corresponding Status of the Project.

k. Technological Partner

It indicates the name of any technological partner eventually involved in the realization of the project. Only technological partners have been considered.

l. Category of Technological Partner

It indicates the category the technological partner belongs to:

- Cloud Platform Provider
- Consulting Company
- Hardware Provider
- Mixed (if more than one Technological Partner not belonging to the same category is present)
- Robotics Company
- Startup
- System Integrator
- Universities & Research Institutes
- Vertical Solution Provider

m. Class of Solution

The Class of Solution for the project is specified:

- Autonomous Robot
- Autonomous Vehicle

- Computer Vision
- Intelligent Data Processing
- Intelligent Object
- iRPA
- Natural Language Processing
- Recommendation
- Virtual Assistant/Chatbot

n. Specification

Based on the considered Class of Solution, the Specification is selected (if the Class of Solutions is further organized into Specifications).

- Autonomous Vehicle
    - Advanced Driving Assistance Systems
    - Autonomous System along a not defined path
    - Autonomous System along a defined path
- Computer Vision
    - Biometric Recognition
    - Image Analysis
    - Image & Video Editing
    - Video Analysis
- Intelligent Data Processing
    - Classification & Clustering
    - Content/Design Creation
    - Forecasting
    - Identification
    - Optimization
- Intelligent Object
    - Home devices
    - Sensor
    - Smart Camera
    - Viewer
    - Wearable
    - Other Smart Objects
- Natural Language Processing

        - Information Filtering
        - Information Retrieval
        - Language Modelling
        - Text Generator
    - Recommendation
        - Contents Recommendation
        - Dynamic Pricing
        - Online Advertising
        - Purchasing recommendation
    - Virtual Assistant/Chatbot
        - Chatbot
        - Voicebot

o. Application

Here the application scope of each project is specified:

- Customer Interaction
- Customer Profiling
- Customer Service
- Cybersecurity/Fraud detection
- Demand Forecasting
- Employees Support
- Energy Management
- Engineering Asset Management
- Ethics & Awareness
- External Logistics
- HR Management
- Internal Logistics
- Maintenance
- Marketing Optimization
- New Product Development
- Operations Automation
- Price Recommendation
- Product
- Product Planning

- Quality Control
- Recruiting
- Regulatory Compliance
- Risk Management
- Safety
- Security
- Sentiment analysis
- Service
- Supply Chain Management

p. B2B/B2C/B2G

This column specifies to whom the project is addressed.

- B2B Internal (the solution is internally used by the company)
- B2B
- B2C
- B2E
- B2G

q. Target Market

In case of solution sold or used outside the boundaries of the company (B2B/B2C/B2G), this column is used to indicate the market that the project addresses. The same sectors used for the column Industry have been considered, with the addition of new elements:

- Fintech
- Insurtech
- Smart City
- Smart Home
- General (when a solution is not attributable to a unique target market)

r. Benefits

In this column, eventual benefits related to the implementation of the Artificial Intelligence solution are provided

s. Critical Issues

In this column, eventual critical issues related to the implementation of the Artificial Intelligence solution are provided

To conclude, some general rules followed in the mapping process are provided, alongside some clarifications:

- In the database, solutions sold by Artificial Intelligence vendors have been excluded. These are for instance Amazon, Google, IBM, Oracle or Microsoft, selling IT solutions leveraging Artificial Intelligence as their core business. This has been done since we are interested in mapping the adoption of Artificial Intelligence solutions, and not its offer. As a consequence, solutions of these companies have been considered only when internally applied, or intended for a B2C market.
- For each project, a line has been dedicated in the database: if a company has multiple projects under exploration, each project has its dedicated line.
- If an analyzed company is involved in a project just as a technological partner of another company, the solution is mapped in the database under the perspective of the other company, while the technological partner of the project is reported in the corresponding column.
- In case of companies investing in startups or other companies developing Artificial Intelligence solutions, eventual Artificial Intelligence solutions of the invested firms have not been considered, except if the investing company declares its intention to adopt their solution.
- Initiatives of a company with the same Class of Solution, Specification and Application have been aggregated in a unique line.
- If a company is clearly diversified in different business lines and sectors, the column Industry for a solution is filled out according to the specific sector in which the project is used.

3. Applicative Scenario of Artificial Intelligence

The aim of this Chapter 3 is to show and comment the results obtained through the adoption of the research methodologies described in Chapter 2; the objective is to provide an overview about the current enterprise adoption of Artificial Intelligence at an international level.

Following the research methodology 235 companies have been considered, and 169 of them actually analyzed. Of the remaining 66, 61 were excluded since exceeding the maximum number of 20 companies analyzed per industry. This threshold was introduced to avoid the creation of a database unbalanced towards sectors more present in the first 235 positions of the Forbes Global 2000 list. The excluded companies are all operating in the sectors “Banking & Finance”, “Energy, Resources & Utility” and “Insurance”, for which the maximum number of 20 analyzed companies was reached. The other 5 companies were excluded, since vendors of Artificial Intelligence solutions; indeed, the aim of this thesis work is to analyze the adoption of Artificial Intelligence, and not its offer.
Of the 169 analyzed companies, no Artificial Intelligence initiatives have been found for 13 of them. 2 of these companies operate in the Manufacturing sector, 6 in the Construction & Real Estate one, 1 in the Insurance one. The remaining 4 companies have not been associated with an industry, being holding companies simultaneously operating in a variety of different sectors.
This is the reason why the created database contains 156 companies, corresponding to 1089 projects mapped and an average of around 7 Artificial Intelligence initiatives per mapped company (or 6,4 per analyzed company).

As follows, Section 3.1 provides an overview of the main results at a very high level, while section 3.2 focuses on each Class of Solutions to analyze in detail the results and describe how the different solutions are being adopted by companies.

3.1 General Results

Starting with a general overview of the sectors, Figure 13 summarizes the number of companies analyzed per industry, projects mapped per industry and the average projects per company in each of the industries, taking into account also the companies with no projects found. To conclude, aggregated values are provided. Despite the actual number of analyzed companies is 169, an aggregated value of 187 companies is shown. This is because it must be considered that a company can operate in different industries at a same time, when adopting a sectoral perspective.

Industry	Mapped Companies	Mapped Projects	AVG Projects per Company
Aerospace & Defense	4	32	8,0
Agriculture	3	12	4,0
Automotive	10	99	9,9
Banking & Finance	20	115	5,8
Chemical	1	5	5,0
Construction & Real Estate	12	15	1,3
Energy, Resources & Utility	20	119	6,0
Food & Beverage	7	69	9,9
Healthcare & Assistance	9	41	4,6
Hospitality	3	17	5,7
Insurance	20	103	5,2
Manufacturing	15	75	5,0
Media & Entertainment	7	50	7,1
Pharma	14	78	5,6
Retail	12	62	5,2
Technology	11	96	8,7
Telco	14	74	5,3
Transportation & Logistics	5	27	5,4
Total	187	1089	6,0

Figure 13: Number of projects and companies per Industry

Although these data are probably not enough to make conclusions about the investment level of different industries, they are useful to show how Artificial Intelligence is a very pervasive technology, finding adoption in a variety of different sectors. Among the most active sectors we can find the Automotive, Banking and Finance, Insurance, Energy, Resources & Utility, and the Aerospace and Defense, with several active Artificial Intelligence initiatives. Technology companies as well are leading the adoption of Artificial Intelligence, as one would expect. One sector in which the number of Artificial Intelligence projects is still limited seems to be the Construction & Real Estate industry, in which several companies with no active projects have been found. Anyway, we can generally say that many very different industries have currently adopted, or are experimenting, solutions based on Artificial Intelligence, confirming the idea of a technology applicable in a variety of contexts.

Other meaningful information emerge from the analysis of data about the Status of the Project and the corresponding Reference Year.

Figure 14 shows how practical applications for Artificial Intelligence technology have emerged only in very recent years, despite Artificial Intelligence has been a field of research for a long time.
Particularly, the increase of Artificial Intelligence projects along the recent years is impressive, showing how this technology is rapidly spreading in a business context.
The reduced number of Artificial Intelligence projects in 2020 must not be interpreted as a setback for this technology: simply, the number of initiatives is inevitably lower since the mapping has been carried out between May and October 2020. What is relevant, is the impressive rise in Artificial Intelligence projects in the last few years. This information is also confirmed by the high average number of projects per analyzed company (6,4), showing how firms are simultaneously adopting Artificial Intelligence in different contexts, or exploring different opportunities.

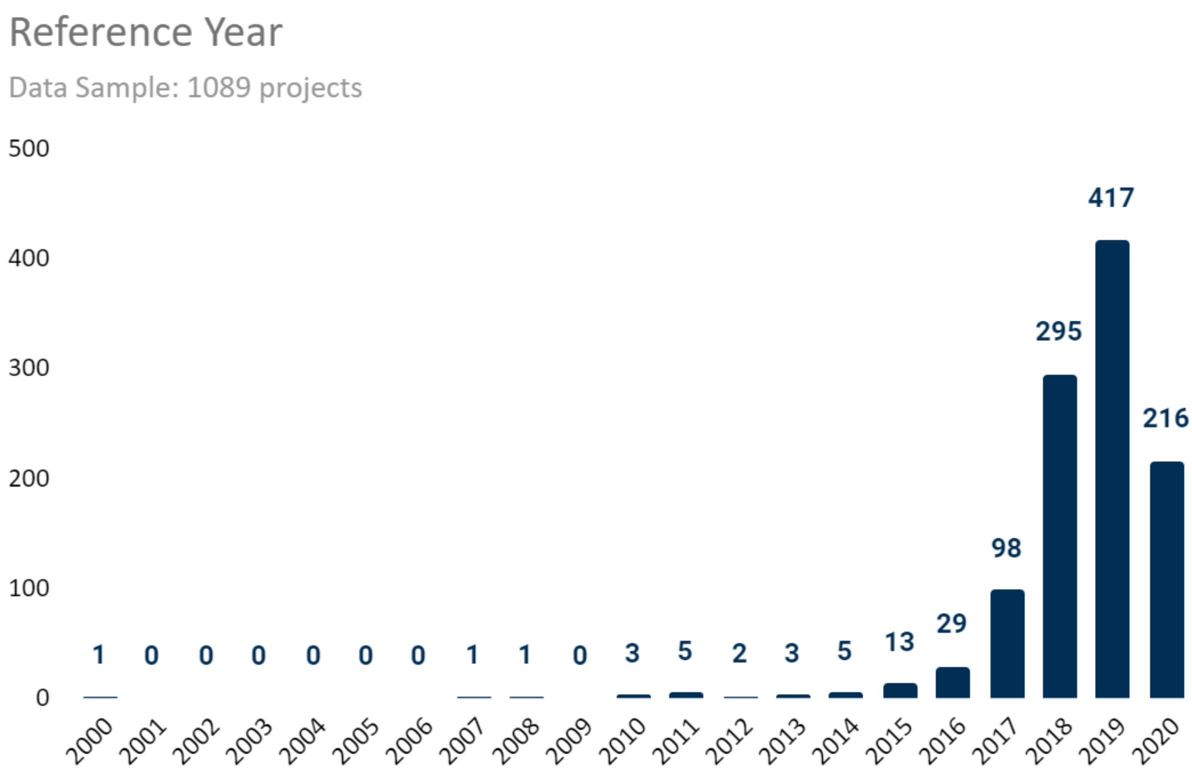

Figure 14: Number of Projects per year

Additional information about the adoption of Artificial Intelligence can be provided by complementing the analysis with data about the status of Artificial Intelligence projects.

As shown in Figure 15, the 57,9% of the projects are already Operative, while the remaining 42,1% are Project Proposals (11,8%), Pilots (21,9%) or projects in the Implementation stage (8,4%).

On one side, this indicates that a large number of projects launched in recent years has been already implemented at a large scale within organizations, signalling a successful and beneficial adoption of the solutions. On the other, the substantial percentage of non Operative projects suggests that many companies are starting only now to consider this technology for business applications and to explore the opportunities offered by Artificial Intelligence. This confirms the idea of the Artificial Intelligence area as a highly dynamic and fervent environment.

However, it must be noticed that the analysis is influenced by the fact that only large firms have been considered and by the use of secondary sources: on one side large companies have been the first to start investing in Artificial Intelligence, on the other side the intercepted projects are just the ones made publicly available, and these typically are the ones at the most advanced stages. Anyway, the results suggest that Artificial Intelligence is not only limited anymore to research laboratories, but has become a technology with practical relevance too.

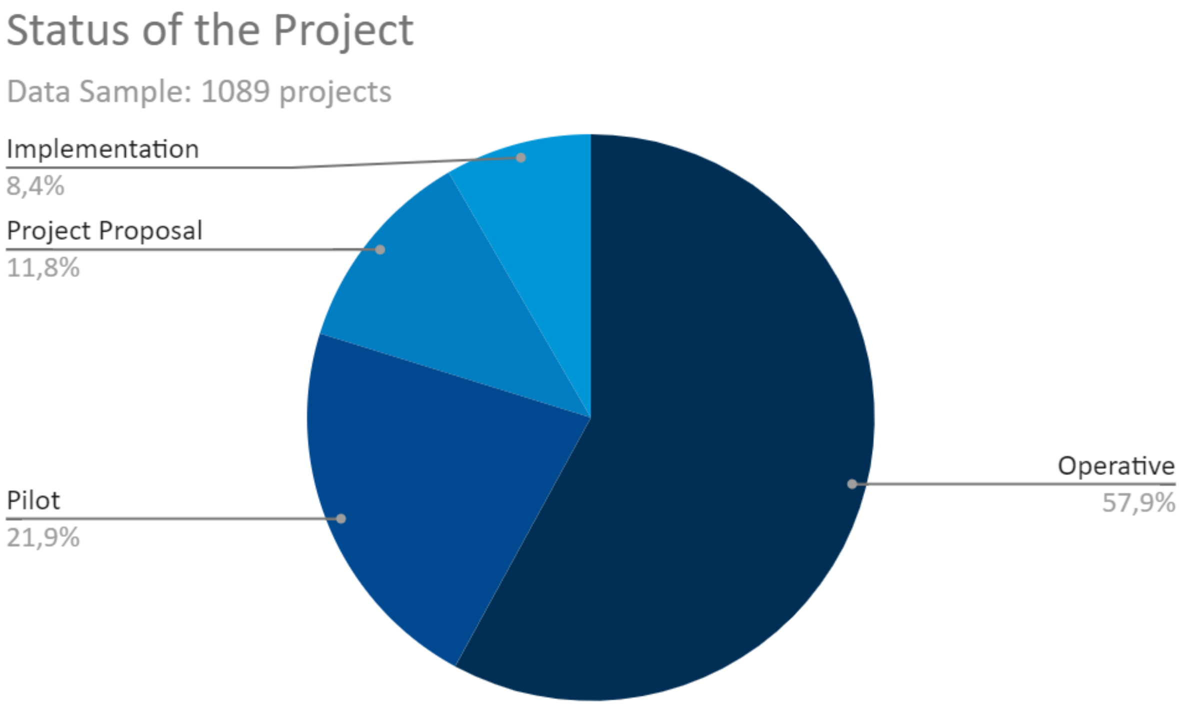

Figure 15: Status of the Projects

If shifting the focus to the Countries and the Continents in which projects are carried out, Figure 16 shows how a huge percentage of them (36,8%) is carried out Worldwide. However, this result is strongly influenced by the large presence of multinational companies in the analyzed sample. Immediately after, Asia and North America are the geographical areas with the majority of Artificial Intelligence projects (respectively, 25,2% and 22,6%), pulled by several initiatives in the United States and China. Europe follows with 13,1% of the projects. Anyway, data are not enough to evaluate the effectiveness of national strategies for Artificial Intelligence development or to provide information about which of the world powers is winning the Artificial Intelligence race.

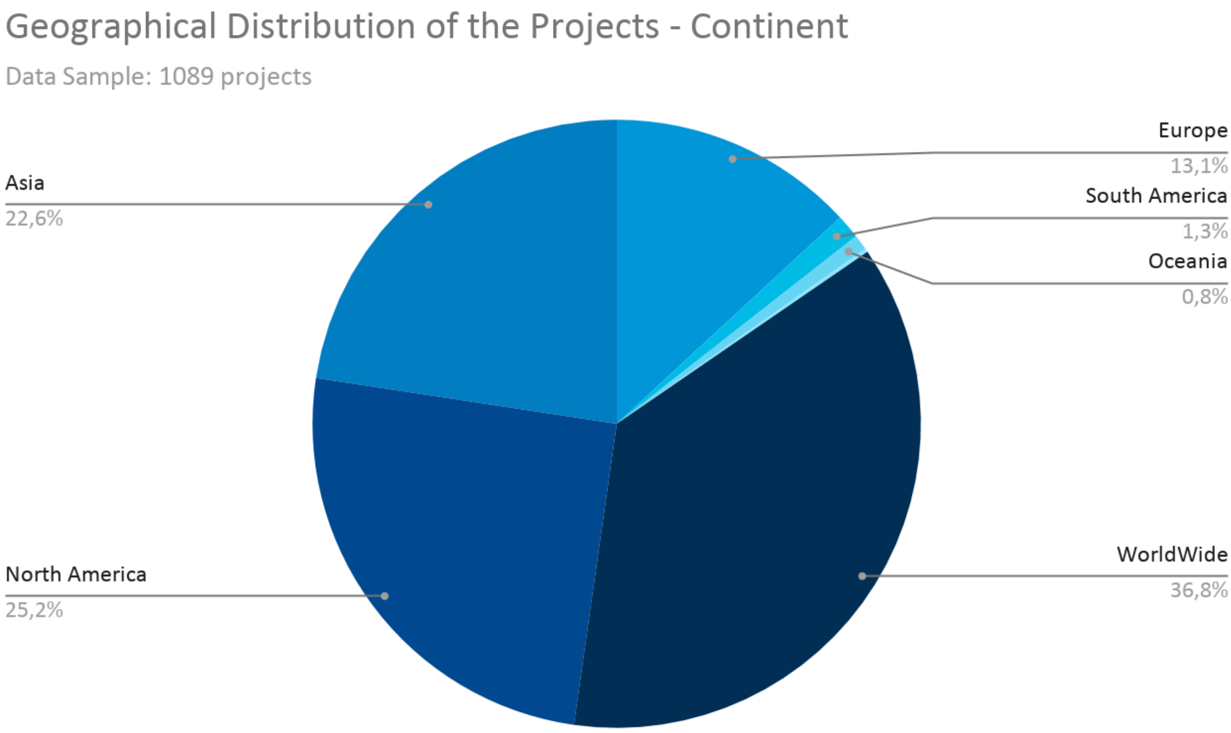

Figure 16: Geographical Distribution of the Projects per Continent

Moving on to analyze the diffusion of the different Classes of Solutions, the main results are summarized in Figure 17. Intelligent Data Processing is the class with the highest diffusion (37,2%), underlying the key importance of Artificial Intelligence analytics algorithms to make sense of the amount of data companies are overwhelmed with and to extract relevant information from them.

Computer Vision, Virtual Assistant/Chatbot and Natural Language Processing solutions are diffused in companies too, constituting respectively the 17,7%, 14% and 11,3% of the mapped initiatives.

Recommendation Systems represent just the 5,2% of the projects, but it should be considered how their implementation is generally relevant only in some industries, such as the Retail and Media & Entertainment ones. Autonomous Vehicles and Autonomous Robots are respectively the 5,7% and the 5,2% of the solutions, underlying how these two Classes of Solutions are still at an embryonic stage and further advancements are needed to reach a wider diffusion: indeed, complex Artificial Intelligence capabilities are required for their purposes.

Similarly, iRPA solutions constitute just the 3,8% of the initiatives, suggesting that the shift from Programmed RPA to Artificial Intelligence Supported or Driven RPA is still far. To conclude, the low percentage of Intelligent Objects (1,3%) suggests that Artificial

Intelligence and Machine Learning algorithms are typically still running on the cloud, while the number of devices able to directly elaborate information on their hardware is limited.

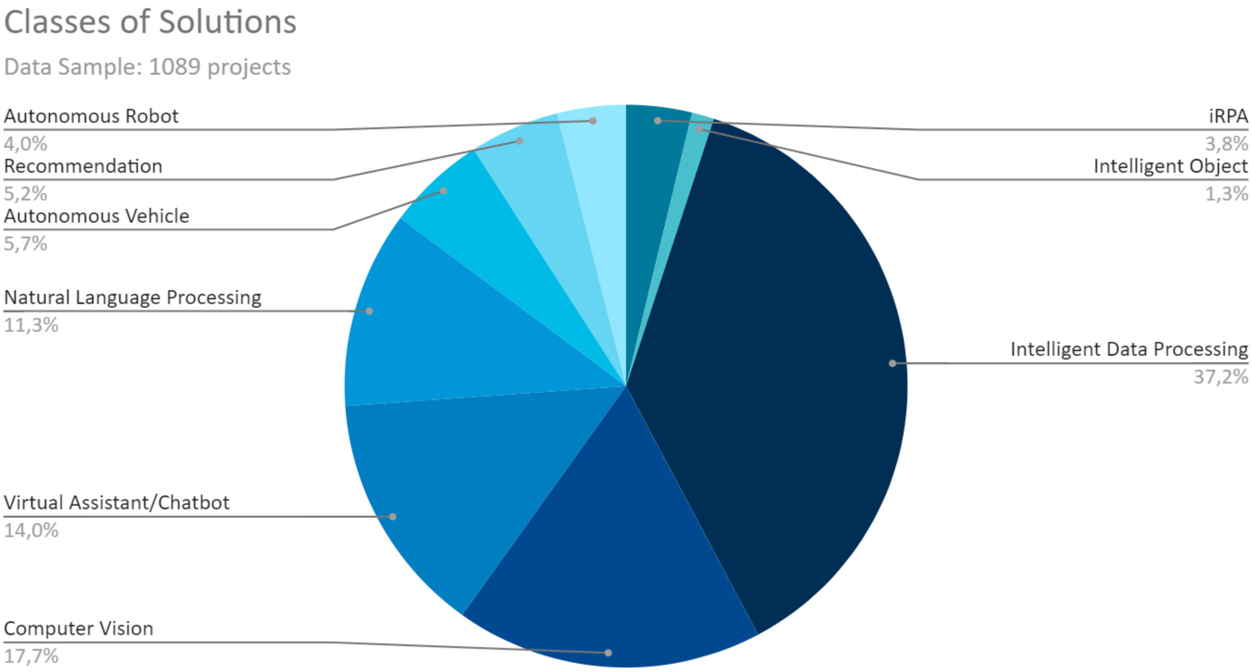

Figure 17: Distribution of Classes of Solutions

The analysis of these results can be supported by considering the Status of the Project for each Class of Solutions (Figure 18): Autonomous Robots and Autonomous Vehicles have the lowest percentages of Operative and in Implementation projects, showing how these solutions are still far from being implemented at scale within companies. Consequently, the majority of projects are at a Pilot Stage or just Project Proposals (in total, the 86,1% of the cases for Autonomous Robots and the 79% for Autonomous Vehicles). This justifies the reduced number of initiatives for these technologies: several companies are investing in these fields, but the fact that they are still under development hinder a widespread adoption of these classes. On the contrary, Recommendation Systems and Virtual Assistant/Chatbot are the classes with the highest percentage of Operative and in Implementation projects, pointing out how these solutions are the ones with the higher level of maturity.

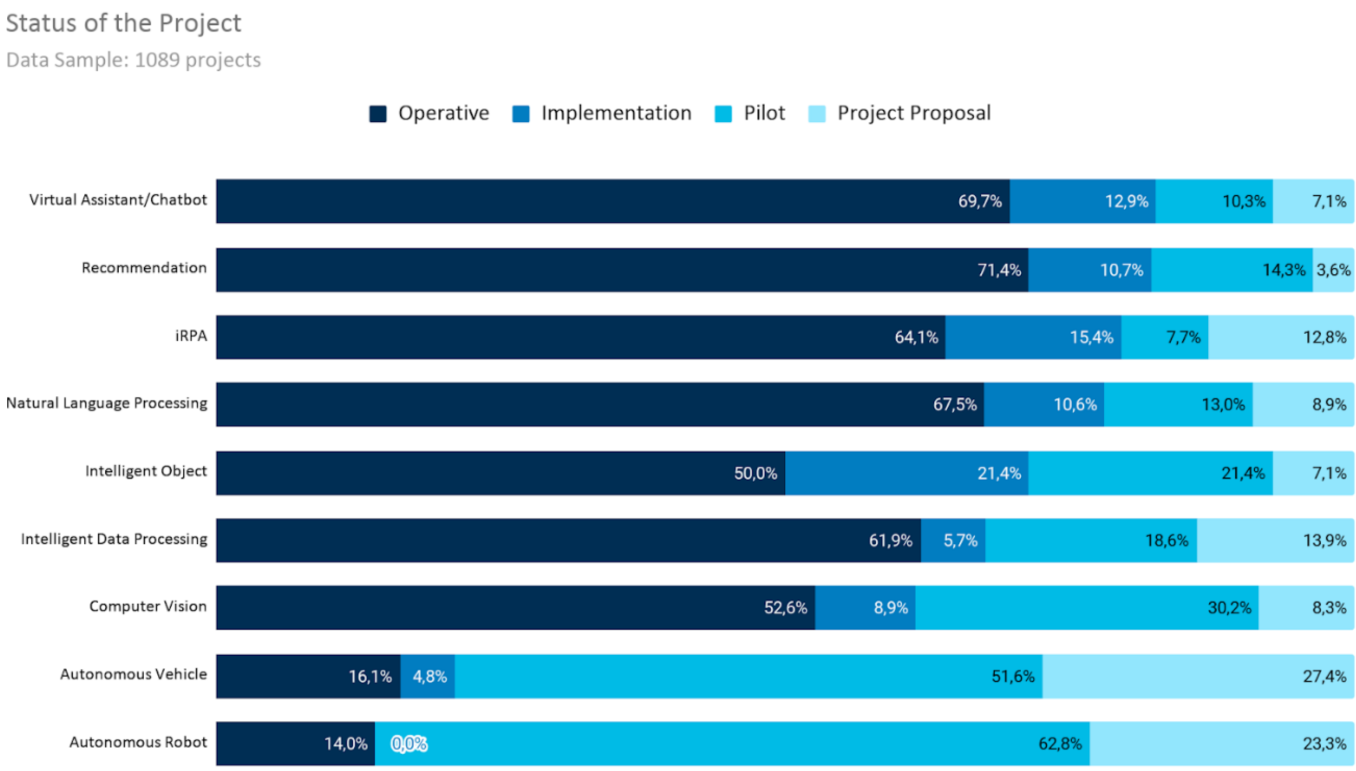

Figure 18: State of the Projects per Class of Solutions

Then, moving to consider where Artificial Intelligence solutions are applied (Figure 19), it emerges that they are mainly internally used by organizations (B2B Internal: 57,2% of the cases). Figure 20 describes the way these solutions are internally used, showing a plurality of possible applications that ranges from New Product Development to Marketing Optimization and Maintenance. Particularly, Operations Automation constitutes the 29,8% of the projects, showing the key role of Artificial Intelligence in the automation of tasks previously performed by employees.

B2C solutions follow, with the 24,8% of the initiatives devoted to final consumers. In this case, as shown in Figure 21, Artificial Intelligence can be the enabler of the functionalities of Services (50,9% of the projects) and Products (15,7%) offered to consumers, rather than being used to improve the interaction between consumers and the company itself (33,3%).

B2B projects represent the 9,6% of the cases, with solutions sold to companies operating in the same or different markets. B2G initiatives are the 5,3% of the projects and can be traced to three Target Markets: Healthcare & Assistance, with a variety of projects to support the public healthcare, Aerospace & Defense, with complex solutions for military purposes, and Smart City. Particularly, the latter represents a scenario in which several practical applications are starting to emerge, ranging from traffic management and optimization to waste collection. To conclude, B2E comprises the remaining 3,1% of solutions, with projects aiming to support employees in their daily work routine at work or to improve their employee experience. Figure 22 summarizes the diffusion of applications independently from internal use or destination to a market.

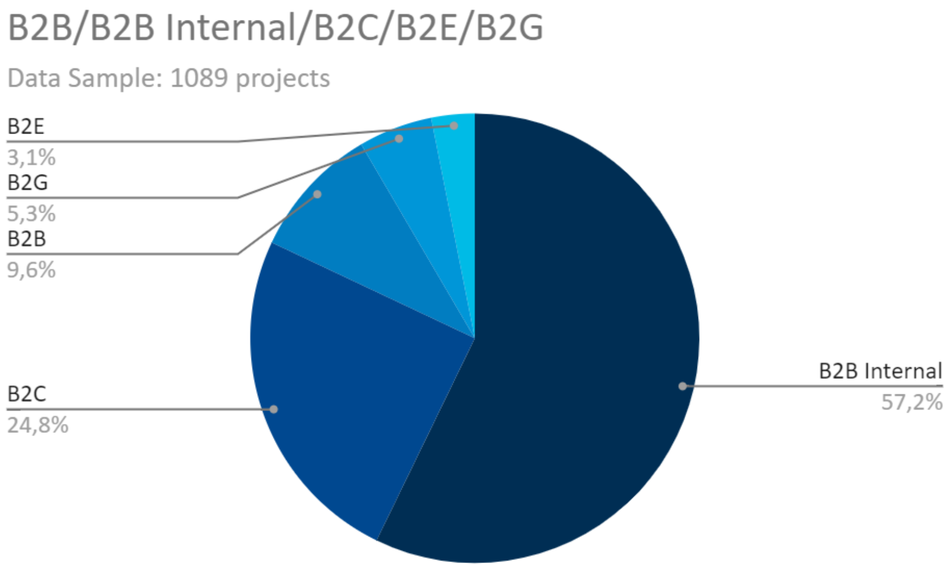

Figure 19: Destination of Artificial Intelligence solutions

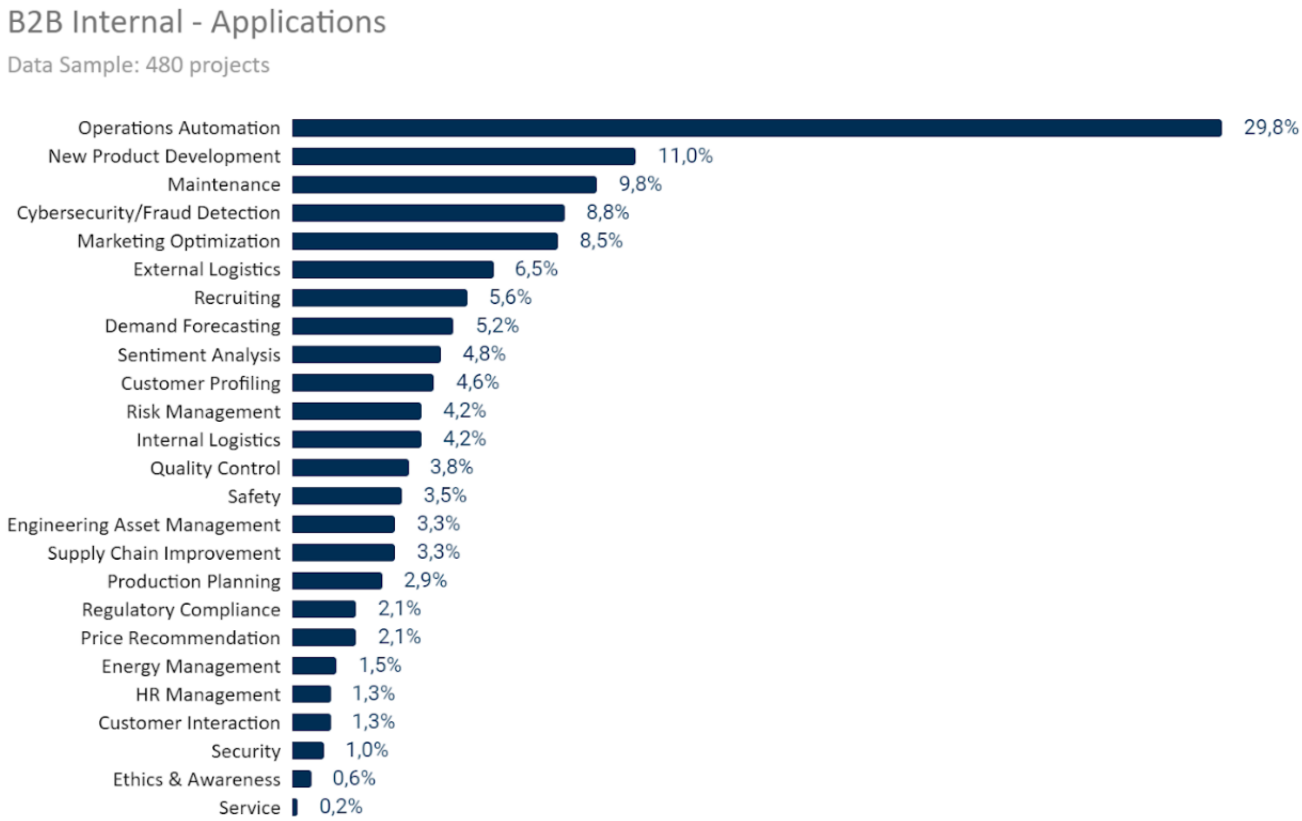

Figure 20: Distribution of Internal Artificial Intelligence Applications

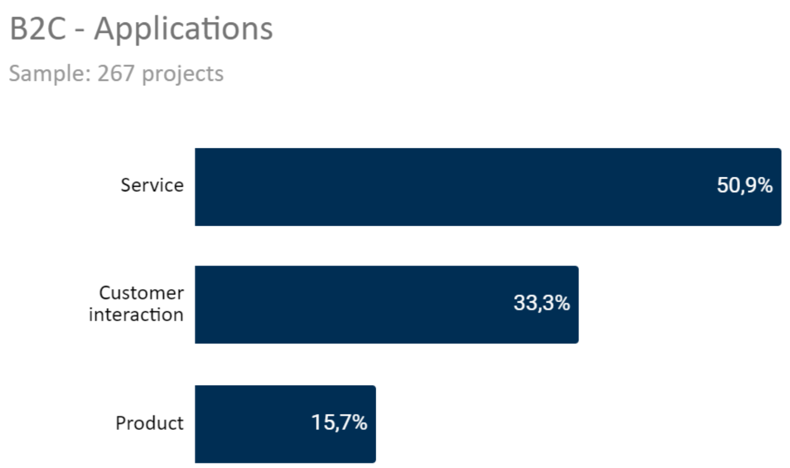

Figure 21: Distribution of B2C Applications

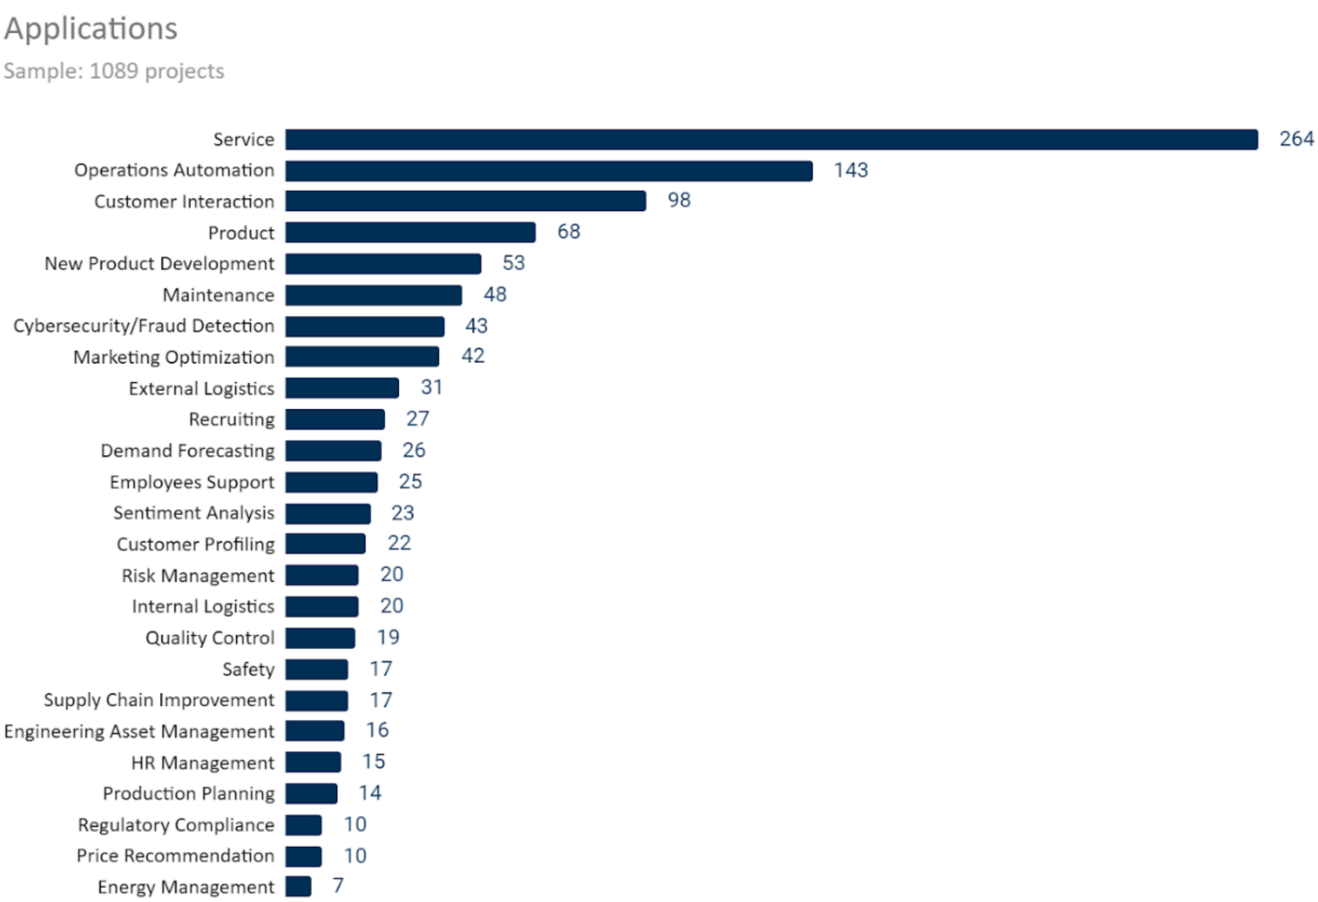

Figure 22: Distribution of Artificial Intelligence Applications

To conclude, external Technological Partners have been involved in the 38,4% of the 1089 projects mapped (Figure 23), suggesting a certain propensity for companies to collaborate, asking for external support when developing their own Artificial Intelligence solution rather than using the Artificial Intelligence solutions sold by other firms.

Considering only the projects with partners, Cloud Platform Providers (27,5%), Startups (23%) and Vertical Solution Providers (20,6%) are the most involved partners.

Consequently, companies are strongly relying on the cloud computing services offered by Cloud Platform Providers; the most relevant ones are the ones of Big Tech companies, with Microsoft's Azure, Google's Cloud AI or Amazon's AWS.

Universities and Research Institutes collaborate in the 7,2% of the projects, while Hardware Providers like Intel and Nvidia have been explicitly involved in the 5,3% of the cases and are starting to provide not only the hardware for storage and elaboration of data, but also the software. To conclude, Systems Integrators, allowing the integration of hardware or software systems, have been involved in the 3,6% of the cases, while Consulting Companies in the 4,3%. Projects in which more partners of different categories have been involved (Mixed) are the 6% of the cases.

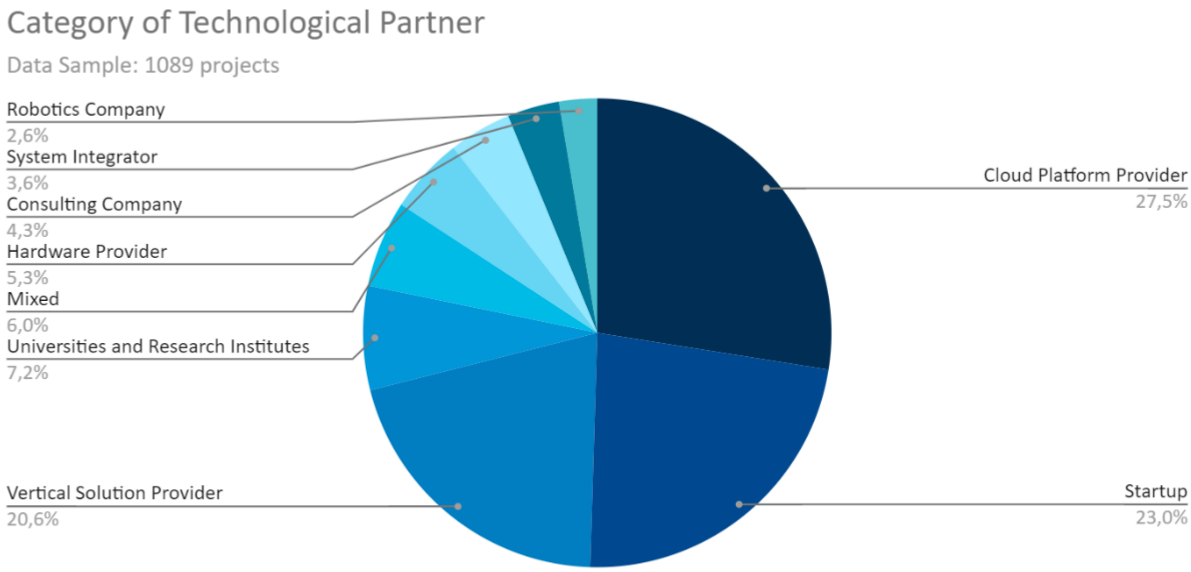

Figure 23: Category of Technological Partner Distribution

3.2 Classes of Solutions

The aim of Chapter 3.2 is to provide a focus on each Class of Solutions and its corresponding Specifications, showing how Artificial Intelligence solutions are applied in a variety of sectors and contexts. When available, additional information about Benefits and Critical Issues of Artificial Intelligence applications are provided.

3.2.1 Intelligent Data Processing

Intelligent Data Processing is the Class of Solutions with the highest diffusion, including 405 out of 1089 analyzed projects (37,2%). This widespread adoption of Intelligent Data Processing solutions is reasonable, because of the impressive amount of data businesses are overwhelmed with: Artificial Intelligence analytic algorithms are needed to make sense of these data, extracting relevant information from them, and to support the decision making process.

When considering the diffusion of Intelligent Data Processing solutions among different industries, results are shown in Figure 24.

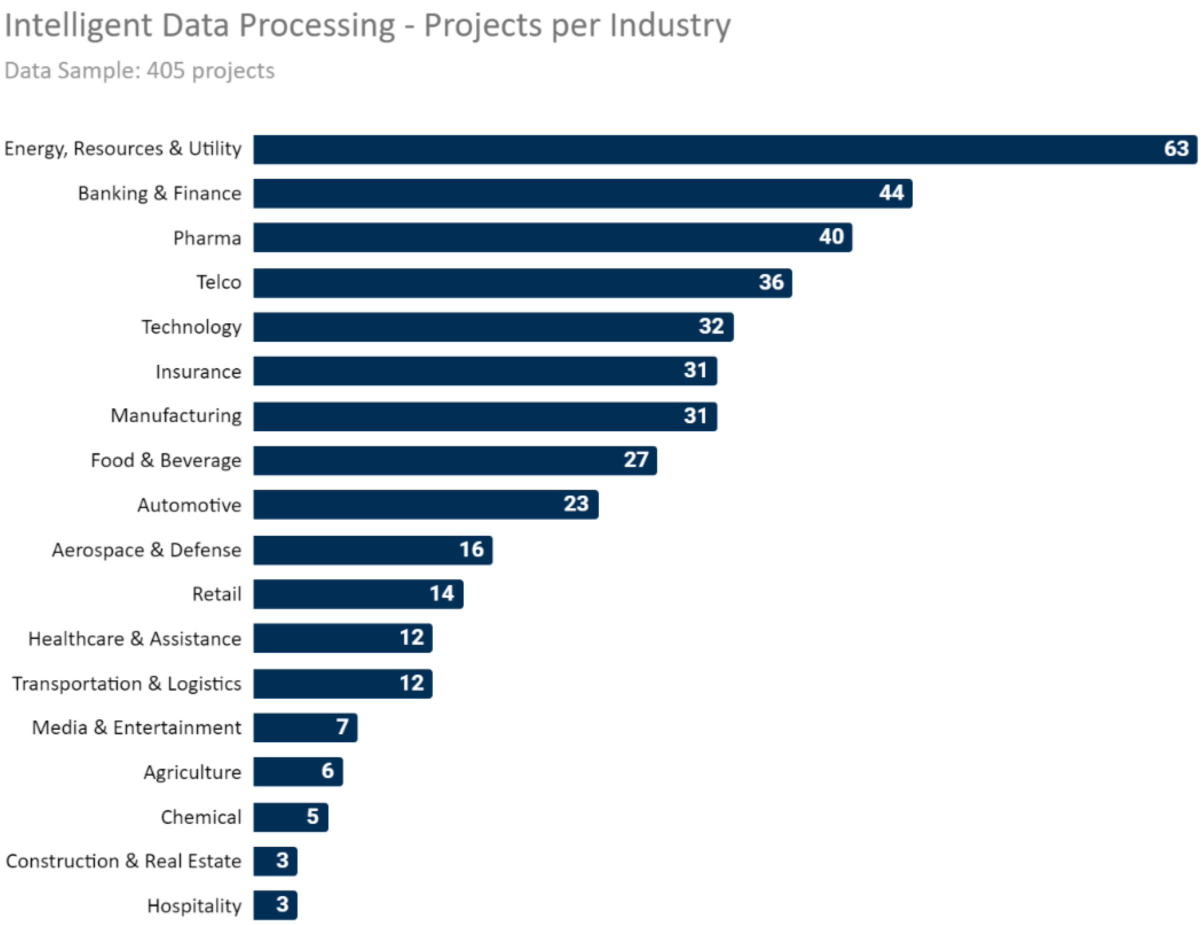

Figure 24: Distribution of Intelligent Data Processing Projects per Industry

As shown, Artificial Intelligence solutions based on Intelligent Data Processing are applied in all the sectors, pointing out how all the industries can exploit, and are actually exploiting, the potentiality of Artificial Intelligence to make sense of data. Particularly, the Energy, Resources & Utility, Banking & Finance, and Pharma industries are the ones with the highest number of mapped initiatives, with different practical applications later explained into detail. For instance, banks are widely using Artificial Intelligence algorithms in Cybersecurity and Fraud Detection and in Services offered to their customers, while companies in the Energy, Resources & Utility sector are applying analytics algorithms to Maintenance, for the identification of possible failures in assets and equipment. To conclude, pharmaceutical companies leverage Artificial Intelligence to analyze data and identify relevant patterns in them for the development of new drugs. Figure 25 shows how the majority of Intelligent Data Processing solutions are typically internally applied by enterprises, with B2B Internal including the 74,1% of the solutions.

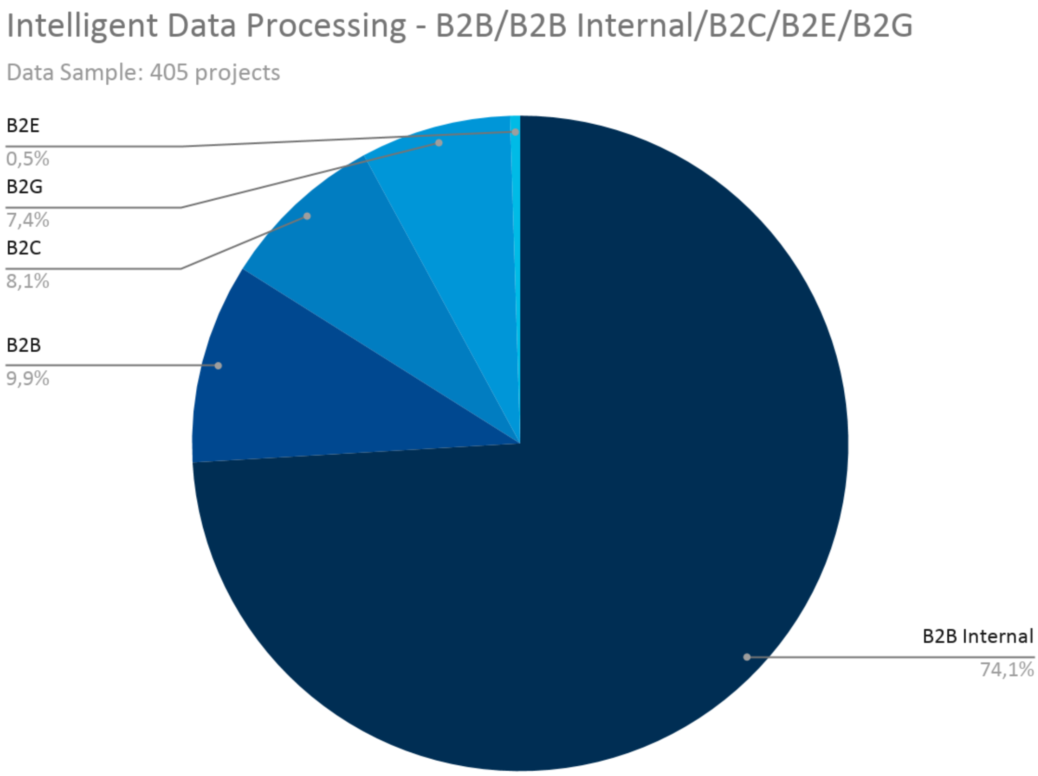

Figure 25: Destination of Artificial Intelligence solutions

To conclude, Figure 26 shows where Intelligent Data Processing projects are applied. Again, it is possible to obtain some general indications about the most diffused Applications, including Services (22,2% of the projects within this Class of Solutions), New Product Development (10,1%), Operations Automation (9,9%) and Cybersecurity/Fraud Detection (9,4%). However, also because of the huge amount of initiatives for Intelligent Data Processing, it is more reasonable to analyze these data at a Specification-level, to have a better understanding of this Class of Solutions.

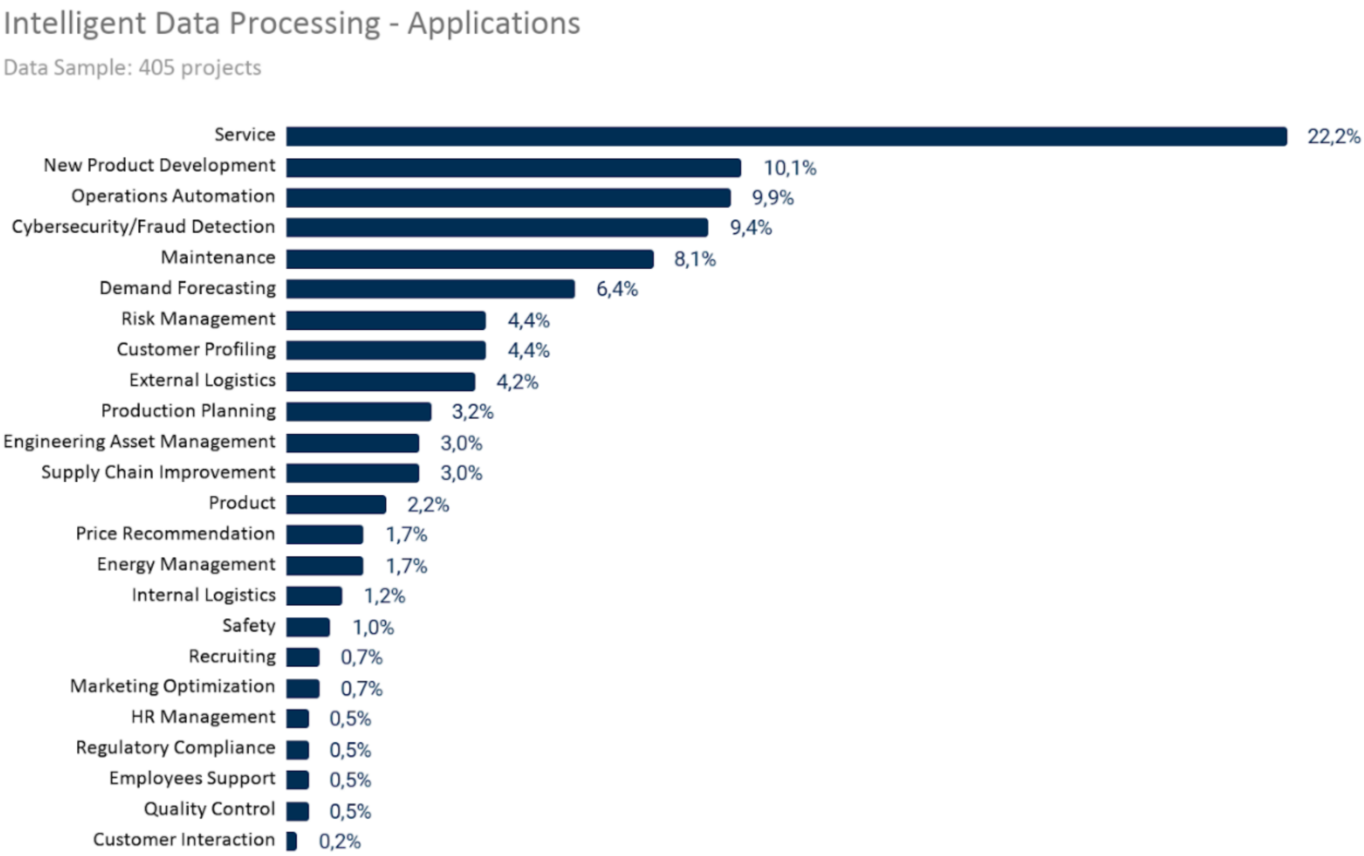

Figure 26: Distribution of Intelligent Data Processing Applications

When considering the distribution of projects in Specifications, results are shown in Figure 27. Identification is the Specification with the highest number of projects (47,4%) in this class, underlying the relevance of Artificial Intelligence to identify patterns or anomalies into Big Data. Forecasting and Optimization follows, respectively with the 23,2% and 21,2% of initiatives. To conclude, Classification & Clustering (4,9%) and Content/Design Creation (3,2%) are the Specifications with the lowest number of initiatives. Each Specification is further analyzed hereafter, to provide a better understanding of the Intelligent Data Processing Class of Solutions.

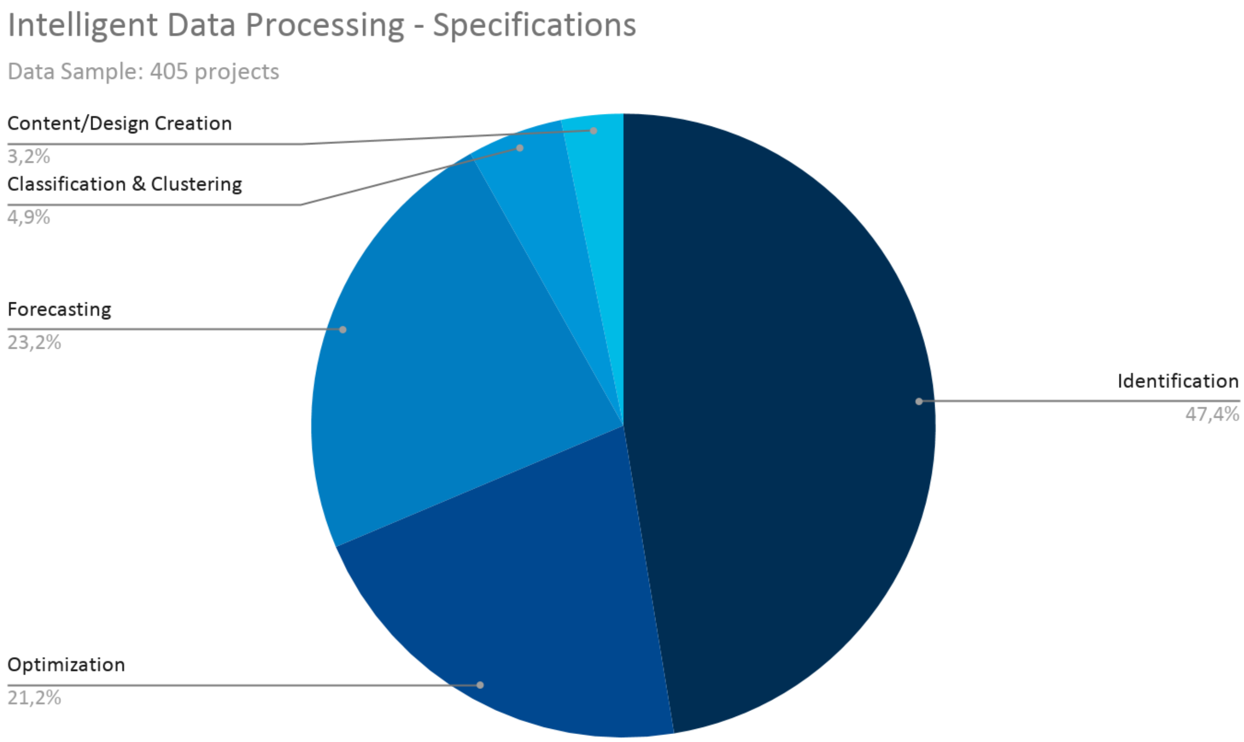

Figure 27: Distribution of Intelligent Data Processing Specifications

Identification

Identification is the most diffused Specification in Intelligent Data Processing, representing the 47,4% of the cases in the class. This Specification includes all the projects aiming to identify anomalies or patterns within large amounts of data. By analyzing which are the most diffused Applications (Figure 28), they include Service (19,8%), Cybersecurity/Fraud Detection (19,3% of projects in this Specification), Maintenance (15,1%) and New Product Development (14,6%).

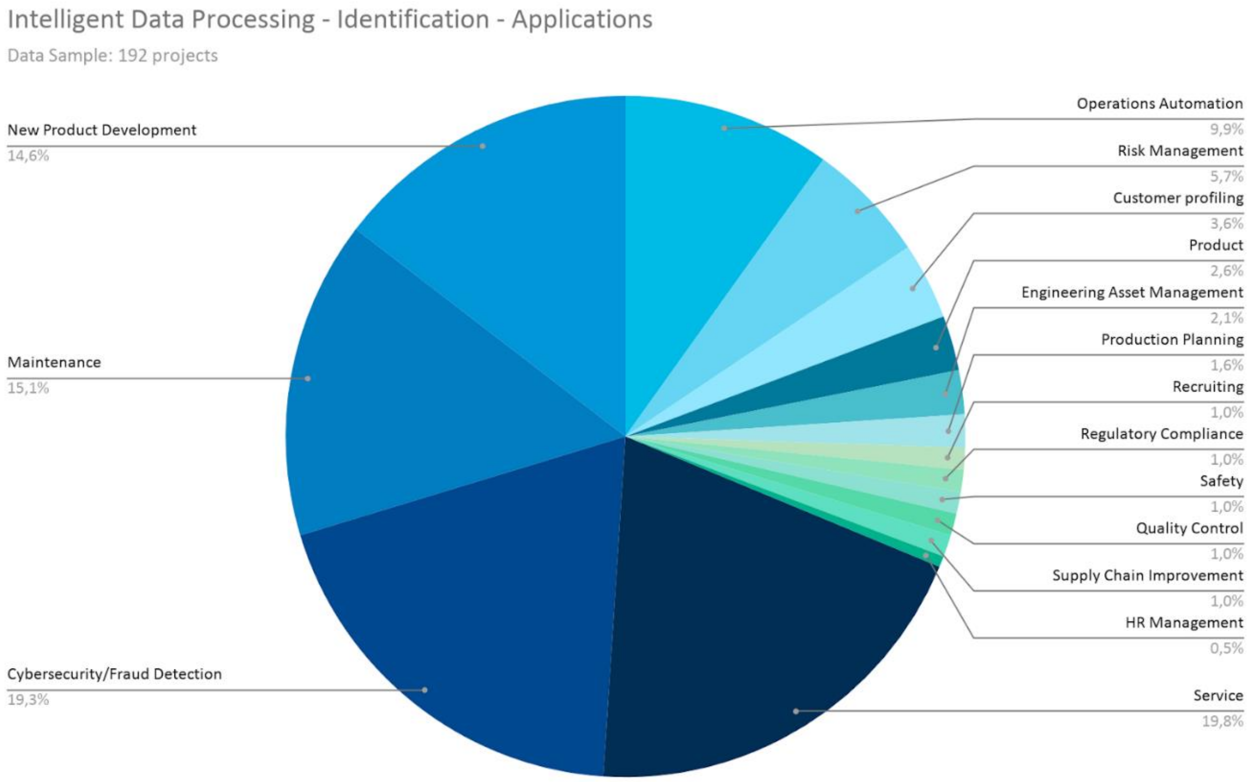

Figure 28: Distribution of Intelligent Data Processing Identification Applications

Cyber Security/Fraud Detection represents the 19,8% of the initiatives and consider solutions for both Cyber Security and Fraud Detection.

Cyber Security regards solutions aiming to protect computer systems and the network from possible cyber attacks, while Fraud Detection refers to projects aiming to prevent and protect the company or the users from possible frauds.

Within this Specification, two industries are the most interested ones in solutions for Cybersecurity and Fraud Detection: Banking & Finance and Insurance, comprising most of the projects for this Application.

For Instance, in Cybersecurity JP Morgan Chase utilizes an Artificial Intelligence algorithm for early-stage detection of malware, Trojans, phishing emails and other malicious payloads targeting bank employees. This is done by using Deep Learning algorithms to identify mass phishing campaigns and malicious URLs, comparing them with characteristics that provide indications about suspicious actors, such as traffic patterns.[4]

As said, banks are also particularly interested in using Artificial Intelligence for Fraud Detection purposes. An example is the one of Wells Fargo, currently using pattern

recognition algorithms applied to multiple data sources to look for signals and recognize money laundry and other types of frauds, and prevent them from happening. [5]
Alongside companies operating in Banking & Finance, Insurance companies are the ones showing more interest in solutions for Fraud Detection, allowing them to understand whether transactions or claims are legitimate or fraudulent. For instance, Manulife, a Canadian Insurance company, utilizes a Machine Learning algorithm based on anomaly detection to support its claim department, analyzing large volumes of daily insurance claims and identifying fraudulent ones. [6]

Similarly to the Cybersecurity/Fraud Detection Application, Services covers the 19,8% of the initiatives for the Specification Identification. In this case projects include a wide variety of different solutions, so that it is not possible to recognize prevalent groups of solutions. Also considering the distribution of projects along the industries, it results that all of them are involved, from the Aerospace & Defence sector to the Automotive industry.
For instance, considering a B2G scenario, General Dynamics, an American Aerospace and Defence Corporation, provides a software solution to the US government using Neural Networks with detection capabilities to recognize signals for Radio Frequency spectrum situational awareness. This solution allows to make sense of the variety of signals in military missions and recognize threats.[7]
Or again, in the Automotive sector, Hyundai, the Korean multinational motor company, is collaborating with the MDGo startup to integrate as on-board emergency service into its vehicles a software to alert medical institutions in real time about a driver's likely injuries after a road traffic accident. This is done by using Artificial Intelligence algorithms analysing data from vehicle sensors to determine the parameters of a crash, such as impact location, vehicle speed and which safety devices were deployed.[8]
However, the variety of Services for this Specification is impressive, with solutions that range from the Healthcare and Assistance sector, analyzing medical data to find possible pattern in specific diseases, to the Hospitality Sector, where Artificial Intelligence can be used to analyze business trends and industry peer benchmarking to help companies in seeing how they are performing compared to competitors.

Another relevant Application for Identification is Maintenance, including the 15,1% of Artificial Intelligence initiatives in this Specification.
Generally, this Application regards the internal use by companies of Artificial Intelligence algorithms based on Pattern Recognition or Anomaly Detection to identify possible failures in assets and equipment, so that operators can implement preventive actions to avoid these problems. Solutions for Predictive Maintenance based on Artificial Intelligence are applied in different sectors, but above all the Energy, Resource & Utilities industry is the most interested one, with 58,8% of these projects.
For instance, Royal Dutch Shell, the multinational oil and gas company, uses Machine Learning algorithms for analyzing data and identifying patterns to prevent in advance failures of compressors, valves and equipment in plants and onshore and offshore platforms. This system has a 24/7 coverage, and it is estimated that the boosting efficiency and the lowering of maintenance cost have generated about 2 million USD savings in Maintenance cost and Downtime. [9]

To conclude, similarly to Maintenance, applications in New Product Development cover the 14,6% of the projects. In this case, the aim is to use Artificial Intelligence algorithms to analyze data and find relevant patterns that could support the creation of a new product and services.
For this application, the interested sectors are even less, and it emerges a clear predominance of companies operating in the Pharma industry, with the 57,9% of the projects in New Product Development for this Specification.
Indeed, Artificial Intelligence technology has an important relevance and diffusion in the Pharma industry, because it helps to analyze data and identify relevant patterns in them to create new drugs or medications. Benefits are several, mainly related to reduction of time, cost and risk associated with new drug development and clinical trials.
For Instance, Tencent Holding, a Chinese Holding company that operates also in the Healthcare and Assistance Industry, developed a Pilot platform called iDrug with the aim to reduce time spent on the R&D finding potentially new active drugs. The platform uses Artificial Intelligence to find the key protein that eventually forms the disease.[10]

Forecasting

Artificial Intelligence projects under the Forecasting Specification represent the 23,2% of the initiatives in Intelligent Data Processing. Generally, the aim of these projects is to use Artificial Intelligence algorithms to provide a prediction on certain variables of interest.

Forecasting solutions are used by companies operating in all the sectors, and also in terms of Applications, as shown in Figure 29, they can be applied to a variety of contexts. Particularly, the 58,6% of Forecasting projects are covered by two Applications: Service (30,9%) and Demand Forecasting (27,7%).

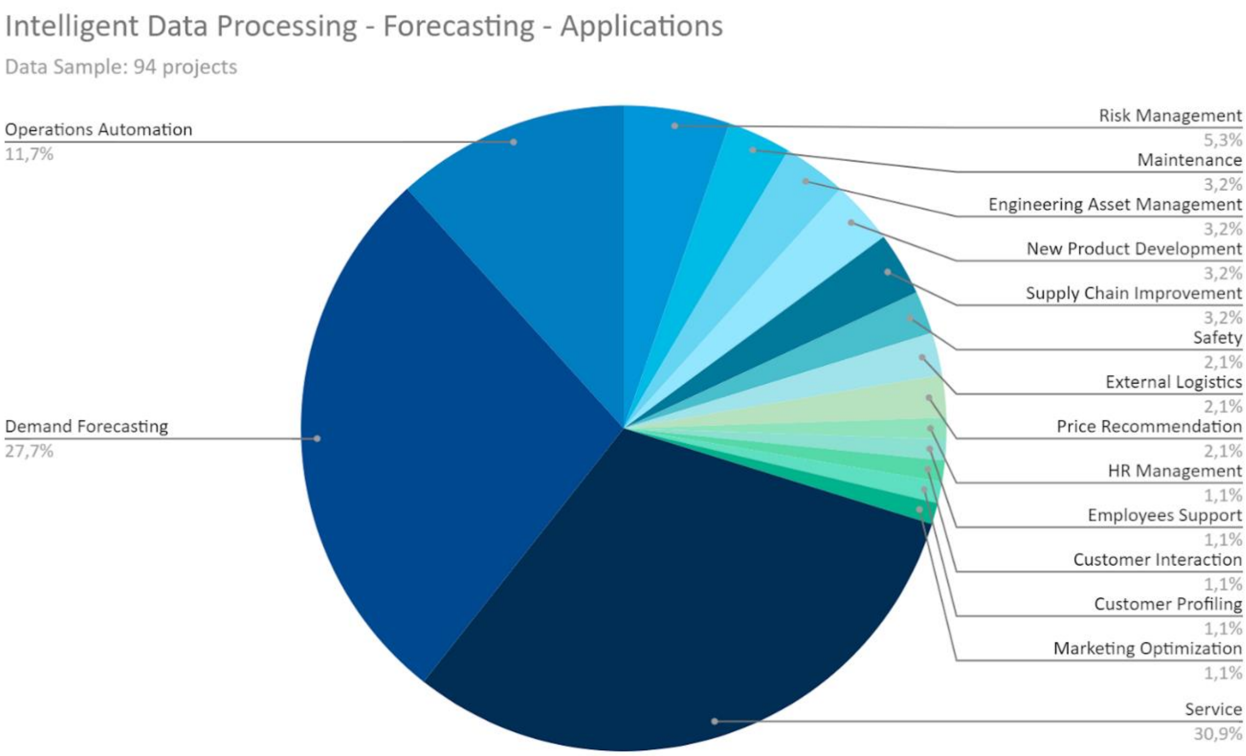

Figure 29: Distributions of Intelligent Data Processing Forecasting Applications

In detail, the Application Demand Forecasting under the Forecasting Specification includes all those projects using Artificial Intelligence to support the prediction of the demand for a certain product or service that the company is selling.

Companies can use this information to implement ad-hoc action for satisfying the future demand, such as increasing the stock quantity or modifying the service based on the users' requests. Artificial Intelligence solutions enabling Demand Forecasting are finding diffusion among most of the analyzed industries, but two sectors above all are showing great interest in these solutions: Telco and Retail.

In Telco, Demand Forecasting regards the use of Artificial Intelligence algorithms for traffic forecasting. By predicting the use of the network in the short term, enterprises can take ad hoc actions for better telecom network configuration management, improving the service and avoiding problems. For instance, Orange, a telecommunication services company, is using Artificial Intelligence for real-time traffic predictions to switch off modules, like antennas, when they are not needed, enabling optimization of the radio access network. In another project, it is collaborating with Nokia and IBM in a joint research aiming to use Artificial Intelligence for predicting demand patterns in mobile networks, allowing to plan them instead of always having to provision for peak usage. [11]
On the other side, in the Retail industry companies can use Artificial Intelligence technology to predict the future demand for products, so that they can supply the stock of the needed item. For instance, Walmart, an American retailer operating a chain of supermarkets, uses Artificial Intelligence to forecast the purchase of more than 500 million store-item combinations each week. This is done also by considering in the prediction model local events such as weather and sport events, and allows customers to always find the needed product sitting on the shelf. [12]
To conclude, it is important to underline the use of Artificial Intelligence for Demand Forecasting also in other sectors, such as Manufacturing, Automotive and Food & Beverage, to predict their products' future demand.

While Demand Forecasting solutions are internally applied by businesses, the Application Services (30,9% of projects in Forecasting) considers Artificial Intelligence solutions addressed to an external market (B2B, B2C, B2G). Again, there is a wide variety of different solutions that companies buying these services can use to predict what they need, from forecasting where the next tornados will be to predicting the possible diseases of apples in a farm.
Particularly, a major part of these Services (27,6% of the cases) can be attributed to the Banking & Finance sector. Indeed, big banks offer their clients systems to predict the movements of the market, or a portfolio, in the future. Therefore, users can make better decisions for allocating the assets of their investments. For instance, Bank of China offers its retail clients Deep Fx, an Artificial Intelligence based app able to predict short term price movements of major foreign exchange currency pairs.[13]

Healthcare & Assistance is another sector relevant to mention for Services. In this case, companies operating in the sector are providing software solutions to public or private medical institutions, that range from the use of Artificial Intelligence to predict future influenza activity, to forecasting the possible outcomes of a specific treatment.
For instance, the Intelligent Disease Prediction Project of Ping an Healthcare utilizes Deep Learning to predict influenza activities in China in advance, allowing disease prevention and control at the city level. This is done by using multi-source data, such as historical data from regional authorities as well as meteorological and environmental statistics.[14]
Similarly, in the ongoing pandemic period companies in this sector are developing systems to help medical institutions in taking actions against the COVID-19 outbreak. For instance, HCA Healthcare, an American provider of Healthcare services, is part of a project to develop the COVID-19 Public Forecasts, a dashboard with 13-day projections for hospitalization and deaths in all the 50 states of America. The dashboard relies on Machine Learning using public data to make forecasts.[15]

To conclude, despite the analysis of the Forecasting Specification was focused on Demand Forecasting and Services, it is important to stress that Forecasting solutions based on Artificial Intelligence can have very heterogeneous Applications in different sectors.
For instance, they can be used in Maintenance to predict corrosion growth rates for assets, rather than in External Logistics to predict the arrival time of each shipment, so that the process can be optimized and a better service offered to the client.

Optimization

Optimization includes the 21,2% of Artificial Intelligence projects in Intelligent Data Processing. This Specification is generally referred to projects using Artificial Intelligence to achieve the optimal value for target variables, modifying variables of the system to achieve the optimum.
As shown in Figure 30, Optimization is used for over half of the considered Applications, but three Applications are the most diffused ones: Service (20,9% of the projects), External Logistic (17,4%), Production Planning (11,6%).

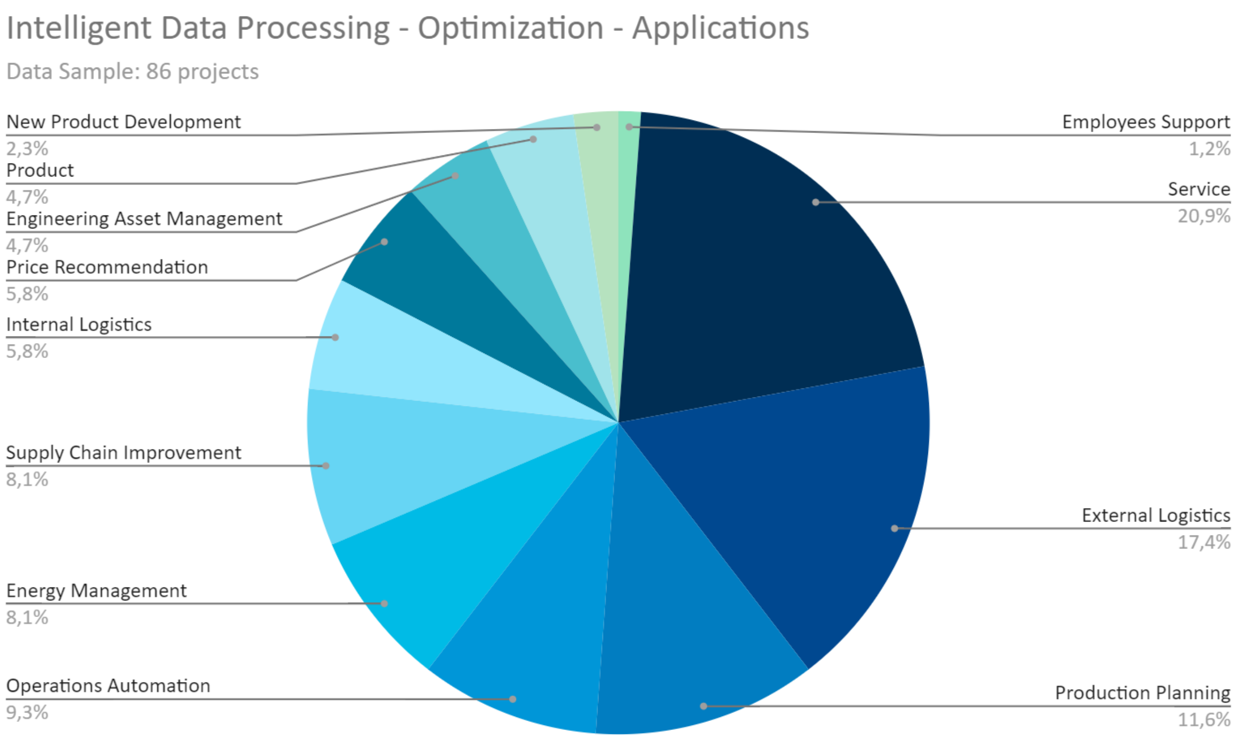

Figure 30: Distribution of Intelligent Data Processing Optimization Applications

The Service Application includes projects targeted to an external B2B, B2C or B2G market. As often happens with Services, a variety of solutions based on Artificial Intelligence exists and it is not possible to identify dominant sectors.

For instance, in the Energy, Resources & Utility sector, Royal Dutch Shell has launched Shell RechargePlus, an electric vehicle smart charging service for its customers. It uses an Artificial Intelligence based optimization algorithm to minimize the electricity cost for the clients, enabling them to move charge to moments in which solar or wind are at their highest and prices are at their lowest.[16]

In a B2B scenario, Caterpillar, a manufacturer of construction and mining equipment, has introduced a subscription-based service for ship owners, leveraging Artificial Intelligence to monitor fuel consumption of ships and optimize efficiency, providing them recommendations on how to reduce fuel operational expenses. [17]

Concluding with a B2G scenario, Volkswagen Group has recently utilized a traffic management system in a trial in Lisbon, using Machine Learning to both predict traffic flows and optimize the routes of nine public transit buses. [18]

On the other side, External Logistics (17,4% of Optimization projects) regards the use of Artificial Intelligence solutions to optimize the transportation of goods, decreasing

shipping costs and wasting time for movements between the factories, warehouses and distribution centres of a company and its customers.
As it can be expected, the Transportation and Logistics industry is particularly interested in this Application. For instance, United Parcel Service, an American package delivery company, uses ORION (On-road Integrated Optimization and Navigation), an Artificial Intelligence software that creates optimal routes for delivery drivers, considering multiple sources of data such as information supplied by customers, drivers and vehicles, weather conditions and accidents. [19]
In addition to firms operating in the Transportation & Logistics sector, this kind of solution is finding widespread adoption among a variety of industries too, providing support to each company that decides to use Artificial Intelligence for External Logistics processes optimization.
For instance, Bayer Crop Science, a subsidiary of Bayer that operates in the agricultural sector, has embedded Machine Learning into its logistics function. This allows them to reach annual savings and cost avoidance of $14 million annually, while reducing 300.000 miles and 350 metric tons of C02 from how they deliver products all over the world. [20]

To conclude, in Production Planning (11,6% of Optimization projects) Artificial Intelligence is applied to manufacturing operations, with the aim to optimize production and improve the manufacturing process efficiency. Solutions range from process optimization through the adjustment of production parameters, to production and workforce schedule optimization.
For instance, in the Energy, Resources & Utility sector, Sinopec, a Chinese oil and gas enterprise, uses Deep Learning algorithms to implement intelligent scheduling and dynamic production optimization. At the same time, Machine Learning is used within its factories for chemical reaction processes optimization in refining and production, adjusting parameters such as the volume of crude oil, fuels, and catalyst. [21]

Classification & Clustering

Classification & Clustering includes the 4,9% of Intelligent Data Processing projects, a limited number if compared with the other Specifications. This specification includes solutions aiming to assign a category to each record from a predefined set, or to divide input data into groups, or clusters.

Particularly, despite a number of projects for other Applications, the majority of initiatives are related to the use of Classification & Clustering models for Customer Profiling purposes. Figure 31 shows how the 50% of the projects for this Specification fall into this category.

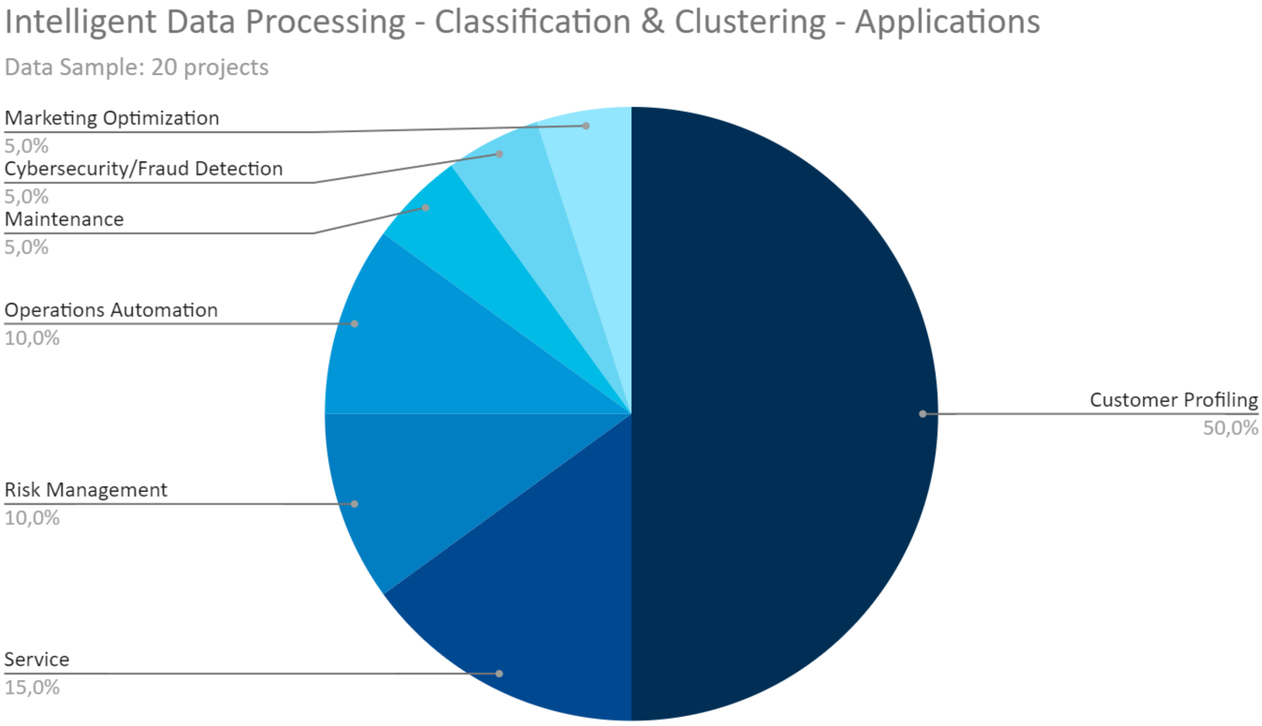

Figure 31: Distribution of Intelligent Data Processing Classification & Clustering Applications

Customer Profiling Application uses Artificial Intelligence to cluster or classify customers into segments, using several data sources for the creation of customer profiles. These ones are later used for a variety of different purposes.

For instance, Procter and Gamble uses Machine Learning that leverages data on web browsing behaviours to identify specific target groups (like mums, sport lovers, fashionistas). These information are later used to improve the marketing and advertising decisions of the company.[22]

In another project, HSBC Holdings, one of the largest banking and financial services organisations in the world, uses Machine Learning to create natural grouping for its customers. Machine Learning has enabled the identification of commonalities among customers and their subsequent clustering, with clusters going beyond the conventional market segments served by the company. The company was able to discover that the needs of some clients were very different from the ones they thought, and these information were used to recommend the right product to the right customers at the right time, based on the financial life cycle of the client.[23]

Content/Design Creation

As Classification & Clustering, Content/Design Creation includes a minority of projects in this Class of Solutions, with the 3,2% of the analyzed Intelligent Data Processing initiatives. This Specification regards the use of Artificial Intelligence algorithms and data to create new contents and design, as well as to plan new products and services. Through Content/Design Creation, companies can take advantage of Machine Learning to create new things in a more efficient way and deliver more value to customers.

As it might be expected, Figure 32 shows how most of the projects for this Specifications regards the New Product Development Application, with the 61,5% of the analyzed initiatives.

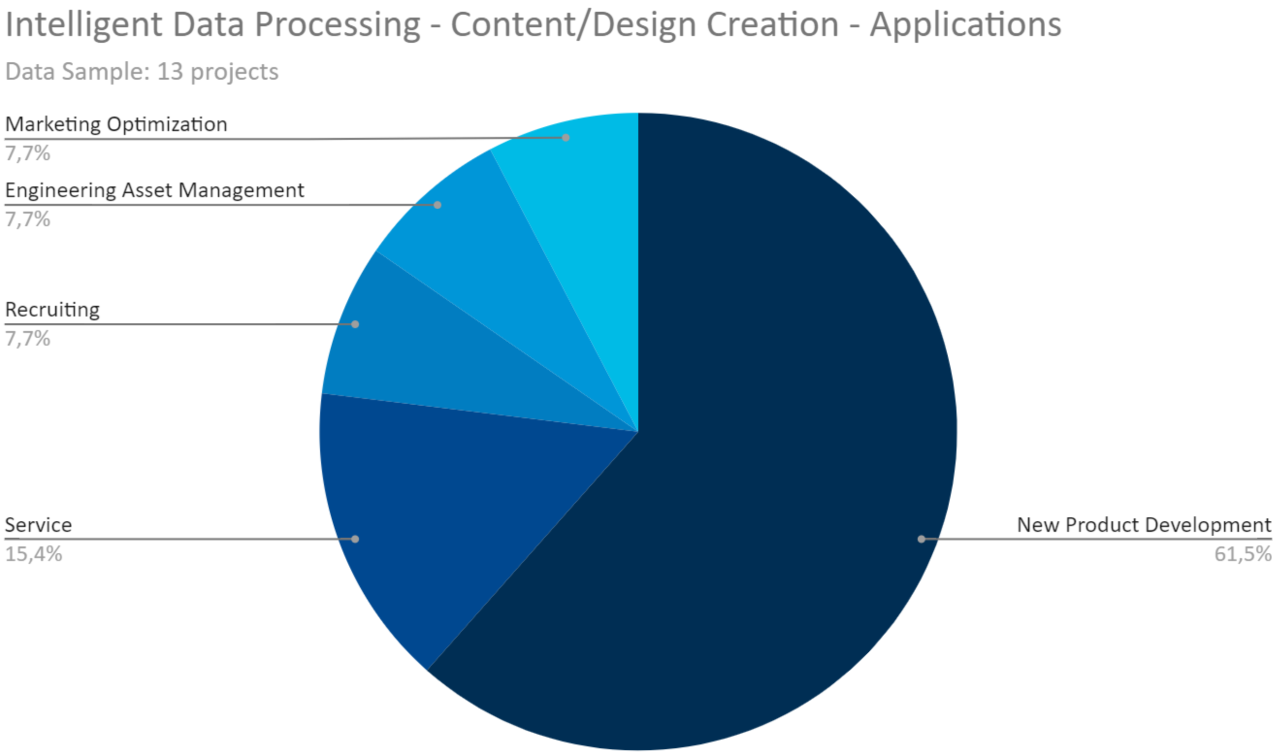

Figure 32: Distribution of Intelligent Data Processing Content/Design Creation Applications

In this case, New Product Development includes projects developed by companies that use Artificial Intelligence algorithms to analyze or process data to create a new product or service. Since the number of initiatives for this Specification is limited, also the involved industries are very few.

For instance, the German company Deutsche Telekom has launched in 2019 a project leveraging Artificial Intelligence to complete the unfinished Beethoven's 10th Symphony. This has been possible through the use of creative Artificial Intelligence for symphony composition, using Machine Learning to understand Beethoven's style and complete missing fragments with meaningful musical movements mirroring its style.[24]
Another example is a Pilot Project developed by Volkswagen in which they use Artificial Intelligence in 3D printing to produce car components, such as the steering wheel, side mirror support or wheels. By training Artificial Intelligence, engineers can use generative design programs to create new products that minimize mass and material, while meeting required engineering constraints.[25]

3.2.2 Natural Language Processing

Natural Language includes the 11,3% of Artificial Intelligence projects, underlying a good adoption rate for this Class of Solutions. Also in terms of Status of the Project, it can be noticed that the 78,1% of the projects are Operative (67,5%) or in Implementation (10,6%), pointing out how Natural Language Processing is a quite mature Class of Solutions supporting companies in their business. On the other side, the 21,9% of the projects have been classified as Project Proposals (8,9%) or Pilots (13%), suggesting that companies are now starting to embrace the potentialities offered by Natural Language Processing solutions.

When considering the industries launching Natural Language Processing projects, Figure 33 shows how most of the sectors are adopting this kind of solutions.

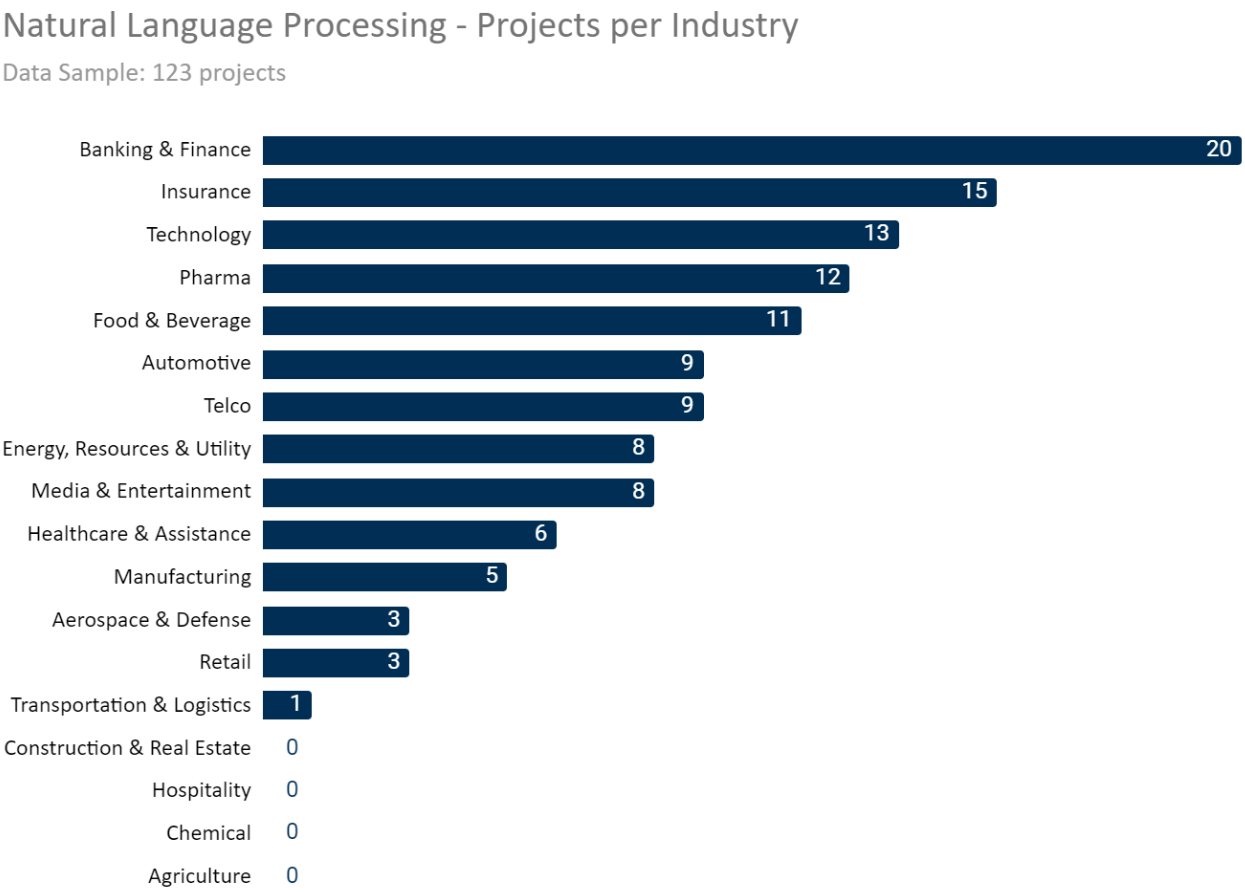

Figure 33: Distribution of Intelligent Data Processing Projects per Industry

Above all, Banking & Finance and Insurance seems to be the industries more interested in Natural Language Processing initiatives. A possible explanation is that these two industries are dealing with a huge number of documents, and Natural Language Processing can be used to process them and make more informed decisions.

When considering the destination of the projects, it can be noticed how the 67,5% of the projects in this Class of Solutions are internally used by companies (B2B Internal). If we also aggregate B2E projects (12,2%), this percentage reaches 79,7%. On the other side, only the 20,5% of Natural Language Processing solutions are commercialised for external use in a B2B (1,6%), B2C (15,4%) or B2G market (3,3%).

This information is confirmed by looking at the different Applications for Natural Language Processing projects. As shown in Figure 24, Operations Automation (20,3% of the projects), Employee Support (10,6%) and Sentiment Analysis (17,1%), that are all Applications strictly related with an internal use, constitutes a major part of the initiatives in this Class of Solutions. On the other side, Services for a B2B, B2C or B2G market, with the 19,5% of the cases, is the other Application widely diffused in Natural Language Processing.

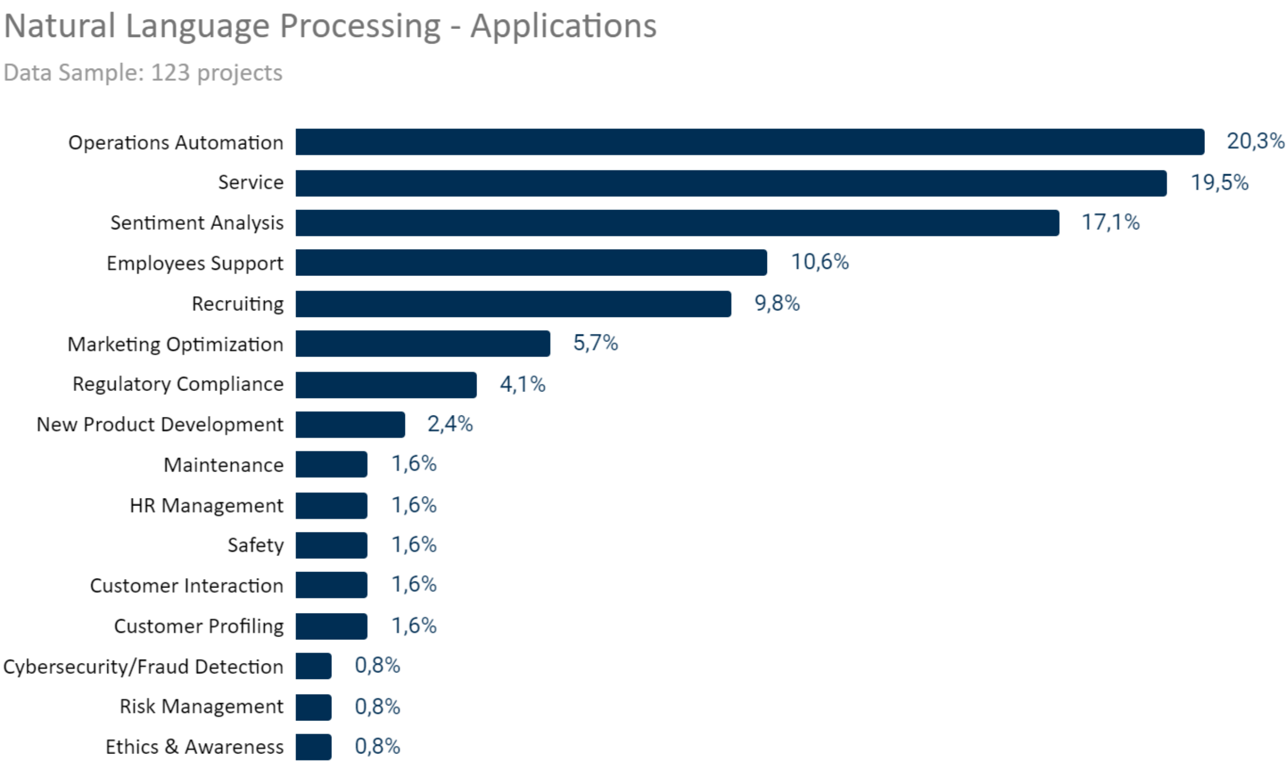

Figure 34: Distribution of Natural Language Processing Applications

As previously done with Intelligent Data Processing, a focus on the different Specifications is required to have a better understanding of the topic. Figure 35 shows the distribution of projects among the Specifications, highlighting how the most diffused Specification is Information Retrieval, with the 55,3% of the projects. The other Specifications follow, with Information Filtering (18,7% of the projects), Text Generator (17,1%), Speech to Text (6,5%) and Language Modelling (2,4%). The latter constitutes a very limited proportion of Natural Language Processing projects, because of the limited application field.

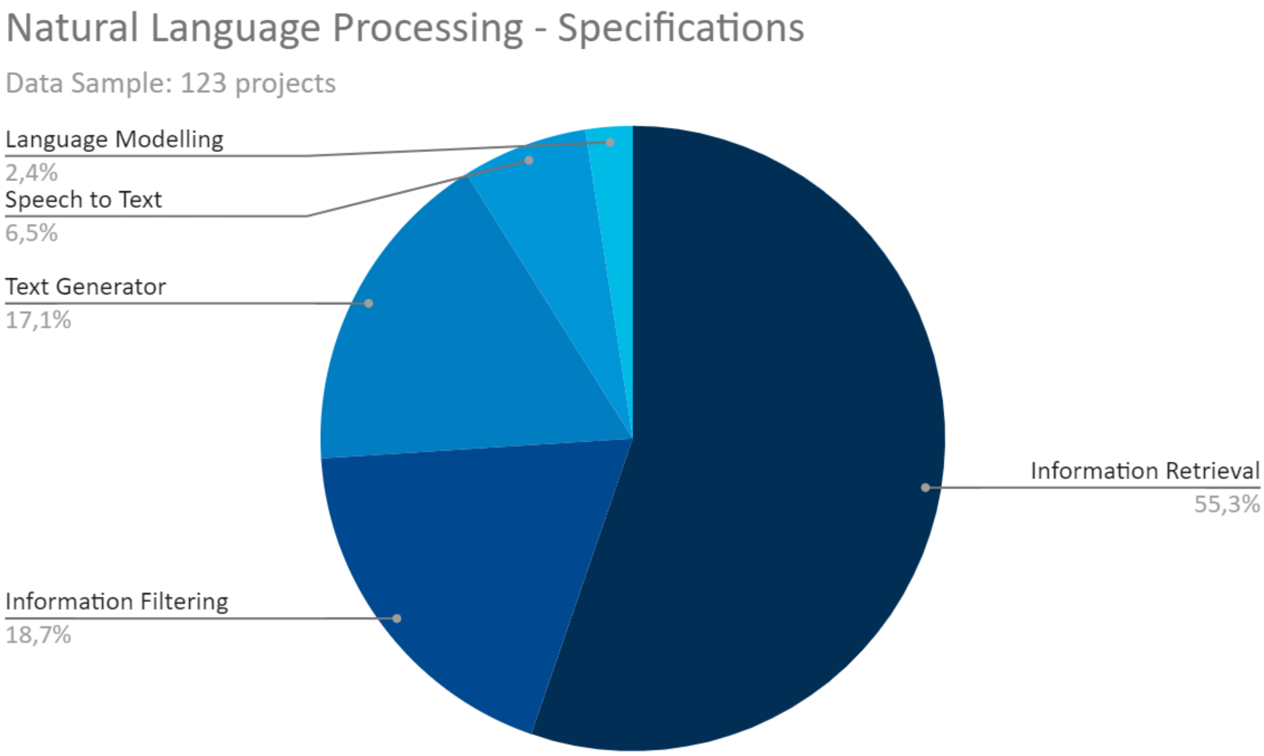

Figure 35: Distribution of Natural Language Processing Specifications

Information Retrieval

Information Retrieval is the most diffused Specification for Natural Language Processing. Figure 36 shows how Natural Language Processing initiatives are typically applied within corporations.

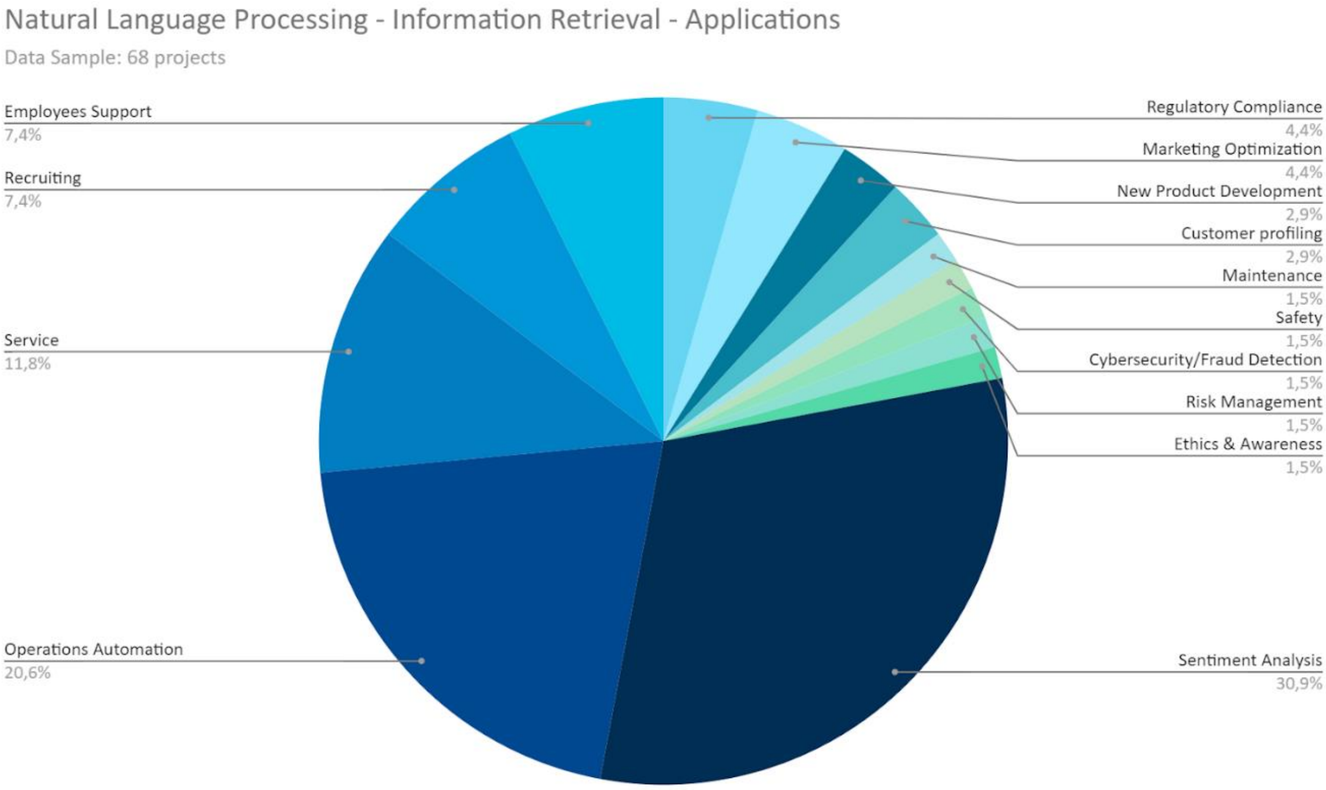

Figure 36: Distribution of Natural language Processing Information Retrieval Applications

As shown, the Applications with the highest diffusion are Sentiment Analysis (30,9% of the cases), Operations Automation (20,6%), Services (11,8%) and Employee Support (7,4%).

Sentiment Analysis is the most diffused Application in Information Retrieval, with 22 projects out of 68. It refers to the use of Natural Language Processing to identify and extract affective states, attitudes, subjective opinions and feelings of the customer through text analytics. Sentiment Analysis is relevant for businesses, since companies want their brand to be perceived positively, or more positively than the ones of competitors. Since Sentiment Analysis gives voice to customers, data are taken from social platforms or other channels, to get public opinions, suggestions or valuations that clients give to certain products or services. This allows to have an overview of the customer brand perception and his/her satisfaction along the different points of the customer journey.

Sentiment Analysis is performed by companies operating in different sectors, internally using this kind of solutions with the aim to improve the customer experience.

For Instance, Honda uses IBM Watson in its quality division to understand and extract information from the impressive amount of customer feedback they receive from several channels. Therefore, employees can have a better understanding of quality issues faced by clients and respond more quickly to them. Honda disclosed that the use of Artificial Intelligence in Sentiment Analysis helped the company to decrease the time required to understand customer feedback by 80%, underlying the importance of this solution for the success of the company.[26]

Operations Automation, on the other side, includes the 20,6% of Information Retrieval initiatives. In Information Retrieval, Operations Automation typically refers to solutions using Natural Language Processing to automate the analysis and extraction of information from unstructured text. The extracted information can then be used to support the decision making process or can be further analyzed.
For instance, Royal Dutch Shell partnered with Maana, a startup company, to use their Knowledge Platform to extract knowledge from internal documents. Natural Language Processing is used to analyze the documents of the company and to automatically extract relevant information from them. Information is then turned into recommendations for employees, to support better decision making in areas such as engineering research or safety incidents.[27]

One sector showing particular interest in this kind of solutions is Insurance, aiming to analyse and extract information from multiple sets of documents especially when it comes to customers claims. Information Retrieval is used to extract information needed for policyholders' payouts, to analyze text for policy reviews, and for the extraction of information from insurance documents of interest.
For instance, Zurich Insurance Group uses Expert System's technology Cogito, based on Artificial Intelligence, to automatically analyze documents for supporting the decision making process of its underwriters. A policy review system to read documents has been launched to improve local policies: in this way, rather than having the underwriter spending hours reviewing the document, he can receive a first analysis with relevant information from the system, that he can decide to confirm or to further investigate.[28]

To conclude, Employee Support collects the 7,4% of Information Retrieval projects.

Information Retrieval solutions for Employee Support are typically applied in call centers of companies, using Natural Language Processing to automatically extract specific knowledge from service manuals, Q&A and other unstructured data, to rapidly bring it up to digital operators. Artificial Intelligence enables them to understand the customer's crux based on voice data, and assist digital operators in timely and correctly answering queries, by providing them relevant information or guiding them to relevant help pages. For instance Anthem, the American provider of health insurance in the United States, uses Artificial Intelligence to improve its customer service: when someone calls Anthem, he often wants to know which are the covered benefits of insurance policies. Anthem uses Natural Language Understanding to understand when this is the case and to bring up relevant information for the operator, automatically populating its screen with them.[29]

To conclude, Services (11,8% of the projects) comprises a wide range of very different solutions, so that a dominant typology of solutions cannot be identified. Initiatives range from the use of Information Retrieval to detect, through textual analysis, the intention of customers in wealth management platforms, to virtual manuals for cars in which Information Retrieval allow users to extract guided answers about what needed, to military use for analysis of technical documentation and bias detection.

Information Filtering

Information Filtering is the Natural Language Processing Specification with the highest diffusion after Information Retrieval, with the 18,7% of projects in Natural Language Processing.

Information Filtering regards the use of Artificial Intelligence to filtering and screen documents and unstructured text coming in all forms based on specific criteria.

Figure 37 shows the distribution of Applications for the analyzed Information Filtering projects.

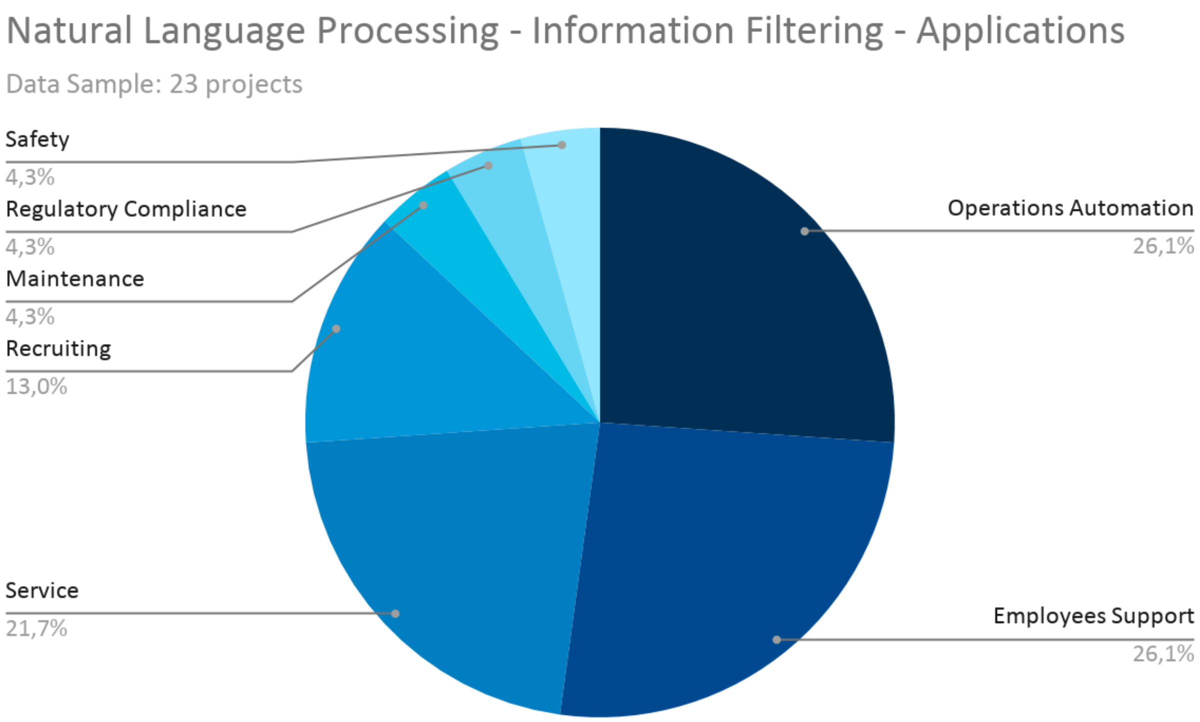

Figure 37: Distribution of Natural Language Processing Information Filtering Applications

It can be noticed how Information Filtering is characterized by a limited number of Applications, with three main Applications particularly diffused: Employees Support (26,1% of the projects), Operations Automations (26,1%), and Service (21,7%). Meanwhile, it is not possible to identify dominant sectors for these Applications, but different sectors are showing interest in Artificial Intelligence solutions for Information Filtering.

In Employees Support, Information Filtering refers to the use of smart screening solutions allowing employees to reduce the time spent in searching for relevant documents. Internal search engines based on Artificial Intelligence allow them to filter search, easily retrieve crucial documents and quickly access them.

For instance, the American conglomerate 3M is using Information Filtering to facilitate the job of its material scientists, who often need to access past research that may be useful for new research. Since finding documents is often time consuming and exhausting, 3M has implemented a smart screening system based on Artificial Intelligence through which they can use natural language queries to quickly find the needed documents. [30]

These solutions are particularly diffused in the Pharma Industry, where Natural Language Processing helps researchers to quickly go through thousands of research papers and articles, enabling them to instantly search for relevant publications.
Solutions in Maintenance (4,3% of Information Filtering projects) and Regulatory Compliance (4,3% too) refers to similar uses of Artificial Intelligence for smart screening, but have been separately classified to underline their focus on the Maintenance and Regulatory Compliance application fields. However, they can also be considered as solutions for Employees Support, pointing out even more the relevance of Information Filtering in supporting the employees of a company to quickly search documents of interest.

On the other side, Operations Automations includes the 26,1% of projects in Information Filtering.
In Information Filtering, Operations Automation generally refers to solutions used to autonomously categorize emails, documents, and other textual resources of interest into categories, through textual analysis powered by Artificial Intelligence. The automation of this task allows to later handle these textual resources more quickly and accurately, or to assign them to the right personnel with the same purpose.
For instance, Generali, the Italian Insurance Group, partnered with Expert System in 2018 to launch a Pilot Program in Spain. In this project, Artificial Intelligence is used for e-mail filtering, to classify approximately 1 million emails the company receives every year based on specific management criteria and more than 200 categories or topics. Since emails are directed to appropriate areas, this solution ensures rapid and smooth content processing.[31]

To conclude, Service represents the 21,7% of Information Filtering projects, and includes solutions for B2B, B2C and B2G markets. The underlying mechanisms of these solutions are similar to the ones previously described for Employees Support and Operations Automation, with Artificial Intelligence enabling Services through categorization of unstructured text, filtering and smart screening.
For instance, a new feature of Google News uses Artificial Intelligence to organize news. While the old system organized articles to show different sources on the same topic, the new Google News uses Natural Language Processing to categorize articles based on the

different perspectives they bring about the same topic. In this way, news is organized into storylines showing the evolution of a situation and the different impacts and reactions of people through different articles.[32]

Text Generator

Text Generator represents the 17,1% of Natural Language Processing projects and refers to solutions leveraging Natural Language Generation. It includes for instance projects using Artificial Intelligence for the automatic translation of text and documents, generation of fragments of text and summaries about the most relevant information in large datasets, text summarization.

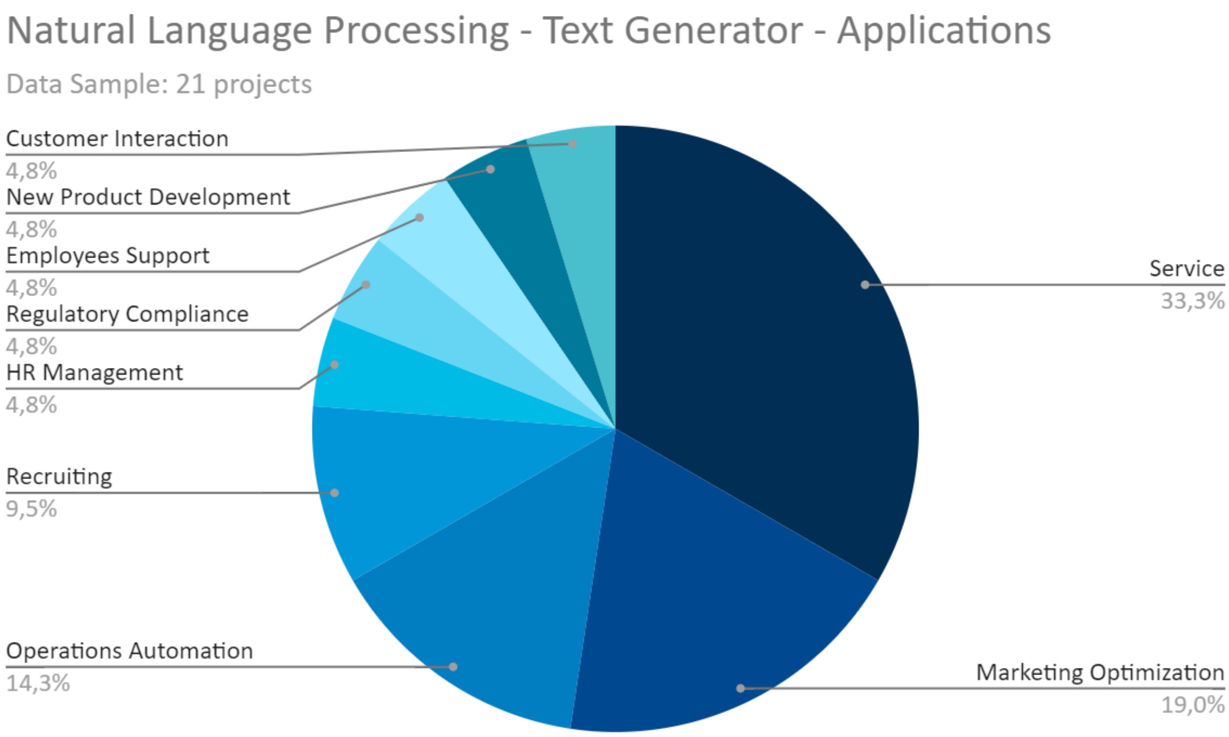

Figure 38: Distribution of Natural Language Processing Text Generator Applications

As shown in Figure 38, Text Generator solutions have different possible Applications. Among them, the most diffused ones are Service (33,3% of the projects) and Marketing Optimization (19%).

Services include solutions mainly addressed to a B2C and using Natural Language Processing for translating, summarizing or paraphrasing text. For instance, from 2016

Google Translate uses a neural machine translation system based on an Artificial Neural Network to improve accuracy and fluency of translations.[33]
Or, for example, Artificial Intelligence is used in Gmail for automatic email reply, suggesting text based on the subject line and what the user has previously written to compose effective and better replies.[34]
To conclude, Microsoft uses Artificial Intelligence in Microsoft Word Online to automatically summarize text for the user. The system is able to recognize key topics in a text, and automatically summarize them in short paragraphs. This way the user can provide relevant information in a fraction of the time.[35]

Text Generator solutions are also finding application in Marketing Optimization. This refers to solutions internally used by companies to improve the effectiveness of their marketing campaigns through text generation: companies can use Natural Language Generation to generate more effective marketing messages, and so to engage more customers. Despite the number of projects for this Application is anyway reduced, two partners are typically involved in this kind of solutions: the startups Acrolinx and Persado. For instance, Nestlè partnered with Acrolinx to use its Artificial Intelligence platform: this one leverages a linguistic analytic engine to read pieces of contents written by employees and provide immediate guidance to improve it, with more effective word choices for the message to convey.[36]

In addition to Services and Marketing Optimization, different Applications with lower diffusion exist, ranging from the automatic generation of job descriptions in Recruiting, to automatic news generation in the Media & Entertainment industry, to solutions internally adopted by businesses for language translation.

Speech to Text

Speech to Text covers only the 6,5% of projects in Natural Language Processing, including all those solutions using Natural Language Processing to convert spoken words into written text. Speech to Text solutions are often preliminary to further elaboration of text through other solutions regarding the other Specifications. For all the projects falling into this category the only other Specification has been considered, since Speech to Text is starting to be considered mainstream by many (think about its widespread use in Virtual

Assistants). As a consequence, this Specification is intended to consider only those Speech to Text solutions in which the only objective is the translation of audio contents into written words, with no further elaboration.

Since the number of initiatives is very limited, it is not possible to make meaningful reasoning on the involved sectors and the most diffused Applications. However, as shown in Figure 39, it emerges how these solutions can both be internally used by companies for Employees Support, Operations Automation, HR Management and Recruiting, as well as find external applications as Services.

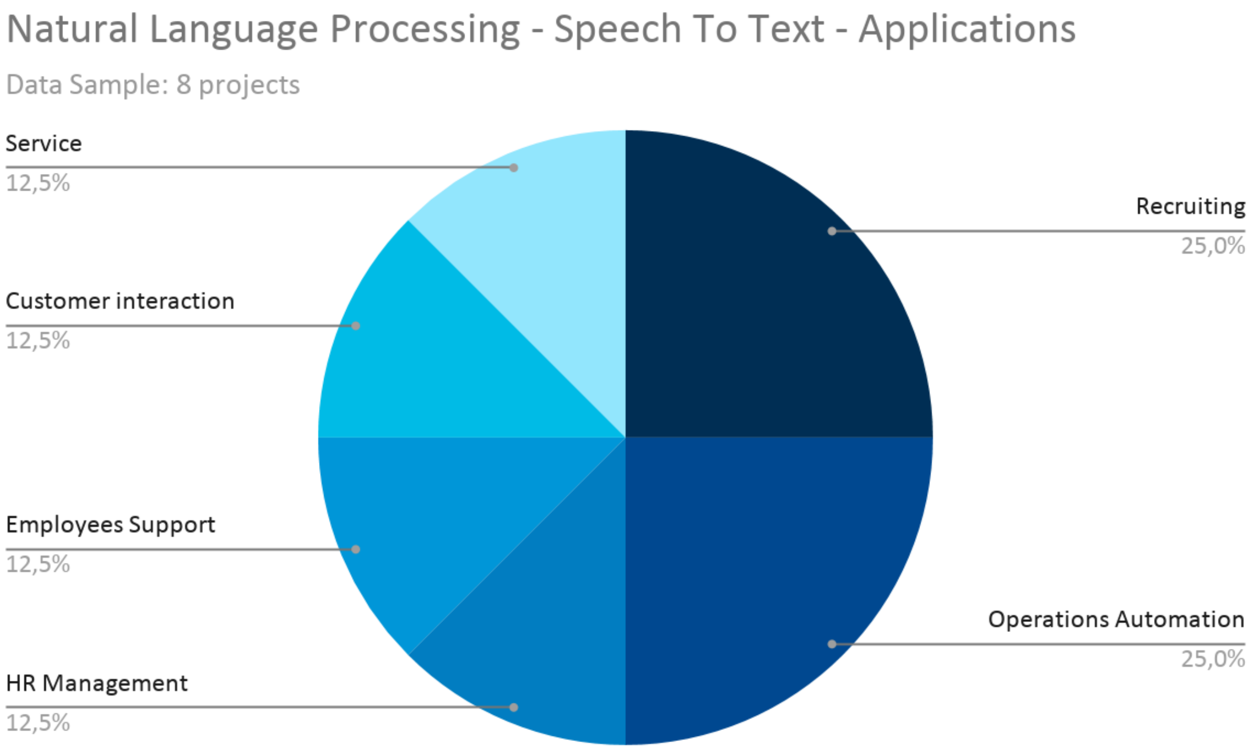

Figure 39: Distribution of Natural Language Processing Speech to Text Applications

For instance, a feature of Microsoft's PowerPoint uses Natural Language Processing to provide captions and subtitles in real time for PowerPoint's live presentations, helping everyone in the room to understand them. This solution is especially targeted for deaf people, who may otherwise have problems in following the presentation.[37]
Another example derives from the American-based Insurance company Allstate, using Artificial Intelligence to capture voice calls of customers and transcribing them. In this way, when customers report an accident, the company is able to capture all the

information provided during the call without missing any relevant detail, and to understand how to better process the claim.[38]

Language Modelling

Language Modelling represents a very limited number of initiatives in Natural Language Processing. This Specification includes all those solutions using Artificial Intelligence to perform: spell check to correct misspelled words in the text; grammar check to correct the word order or the verb tense errors; word auto completes to suggest the remaining part of the word the user is typing. The very few projects collected in the database are all Services for a B2C market. For Instance, Google uses Natural Language Processing to autocomplete the word typed by Google Search's users.[39] Meanwhile, Microsoft has developed a feature of Word to correct misspelled words and provides writing alternatives based on the context of the sentence.[40]

3.2.3 Virtual Assistant/Chatbot

The Virtual Assistant/Chatbot Class of Solutions represents the third class by number of projects, with 14% of the Artificial Intelligence initiatives. This should not be interpreted as a signal of intermediate diffusion for this Class of Solution. By analyzing the collected projects, it results that 67,3% of the companies have at least one project in the Virtual Assistant/Chatbot Class of Solutions. This impressive percentage suggests that Virtual Assistants and Chatbots have a high level of adoption and have become a widespread reality within businesses. These considerations are confirmed by information about the Status of the Project: the 82,6% of the considered initiatives are Operative or in Implementation, showing how this is a mature Class of Solutions, while the remaining 17,4% refers to Chatbots and Virtual Assistants that companies plan to introduce in the short term or have introduced as pilots. The wide diffusion of this Class of Solutions can be explained also by considering that, since Chatbots and Virtual Assistants are typically used to interact with customers, they are not depending on industry-specific operations. Therefore, they can be adopted in almost all the sectors for Customer Interaction purposes. Figure 40 confirms the applicability of conversational interfaces in a wide variety of industries.

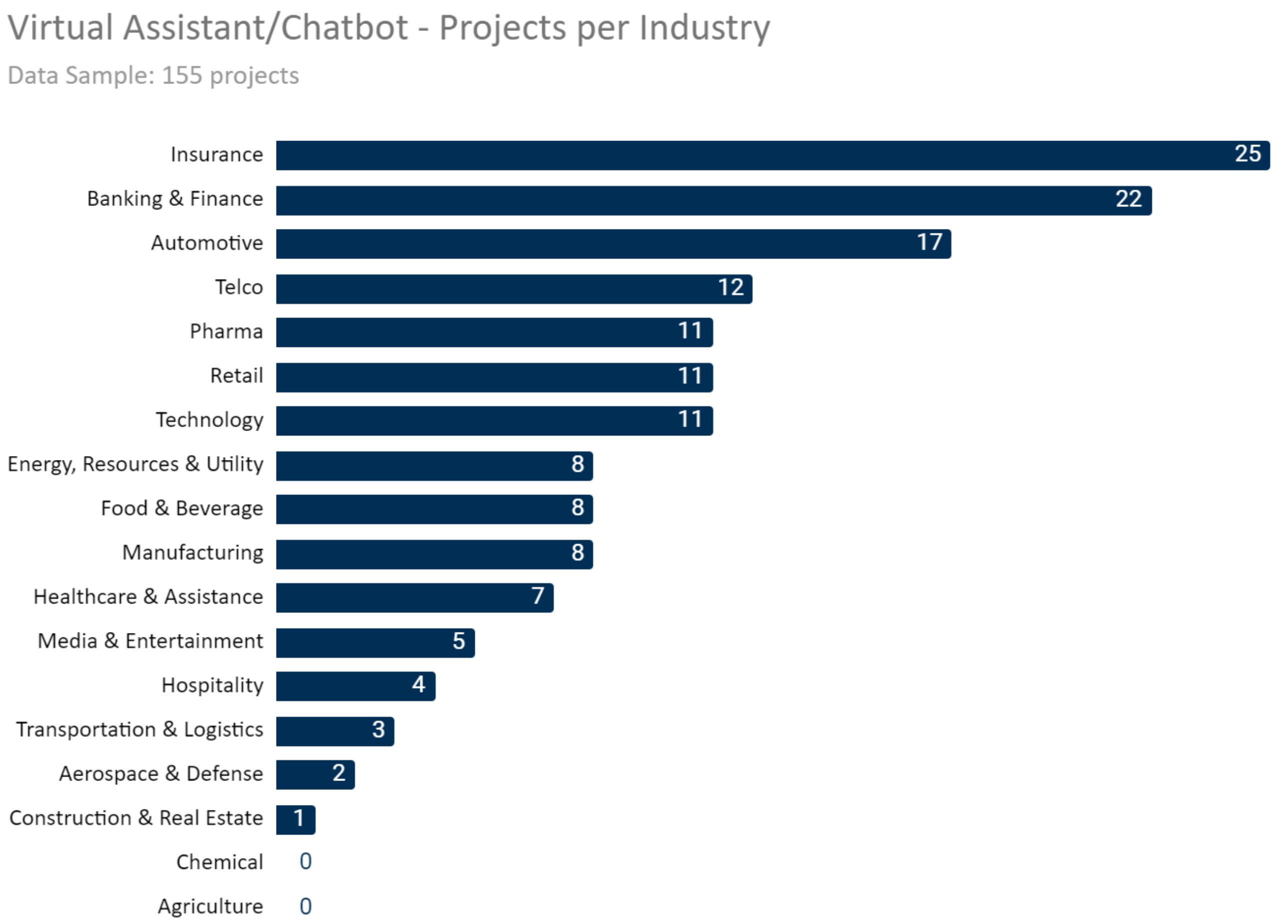

Figure 40: Distribution of Virtual Assistant/Chatbot Projects per Industry

On the other side, when it comes to possible Applications, Figure 41 shows how these solutions are mainly used for Customer Interaction (45,2% of the Virtual Assistant/Chatbot initiatives) and as Services (31%) offered to consumers, for a total of 76,2%. In parallel, internal applications for Chatbots and Virtual Assistants are emerging, with solutions for Employees Support (6,5%), Recruiting (6,5%) and HR Management (3,9%) representing in total the 16,9% of the initiatives in this class. These results are aligned with information about the solutions' destination: B2C for the 69,7% of the projects, B2E and B2B Internal for respectively the 10,3% and the 10,3% of them.

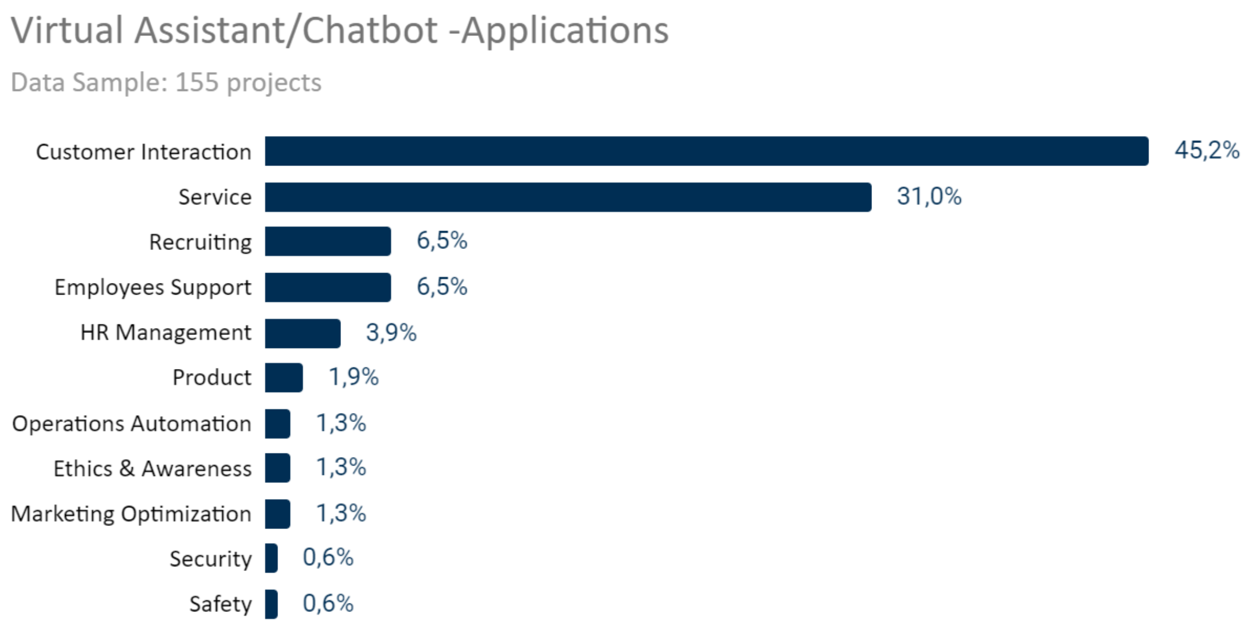

Figure 41: Distribution of Virtual Assistant/Chatbot Applications

When moving to analyze the distribution of Specifications (Figure 42), Chatbots demonstrate a higher diffusion with the 63,9% of the initiatives, but Voicebots (Virtual Assistants) are popular too, with the 36,1% of the projects.

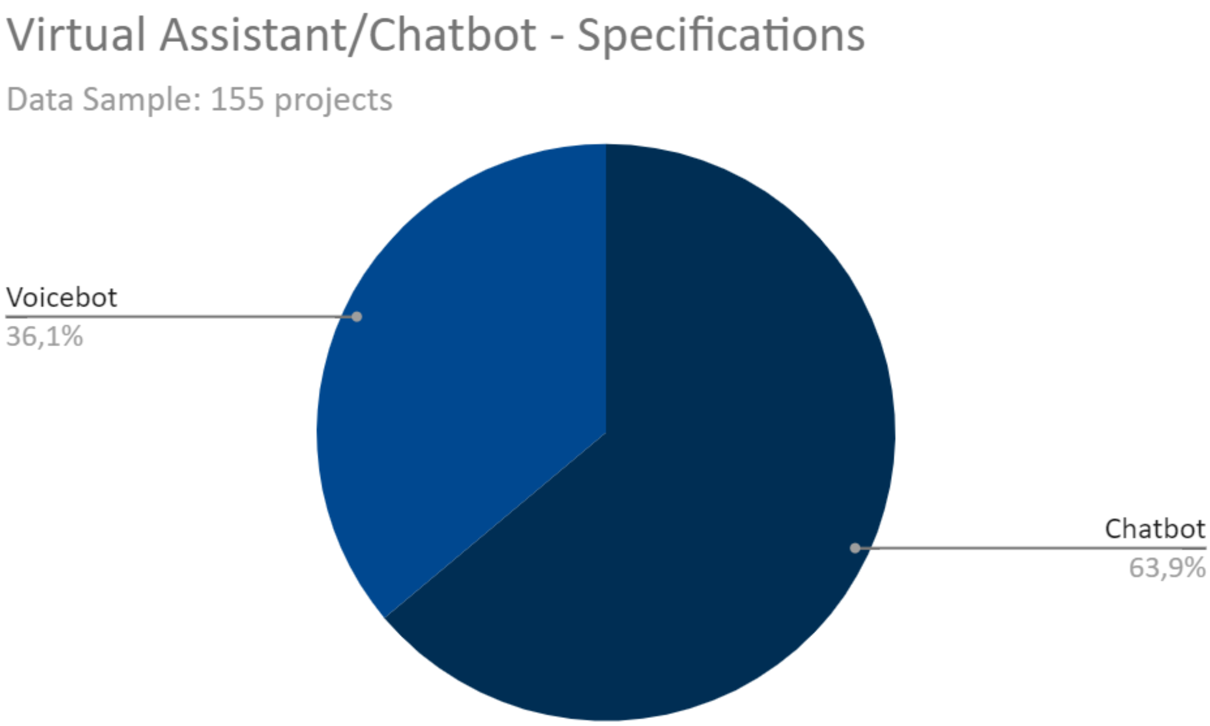

Figure 42: Distribution of Virtual Assistant/Chatbot Specification

Chatbot

Figure 43 shows the distribution of Applications for Chatbots. As shown, the predominant Applications is Customer Interaction (61,6% of the Chatbots initiatives), followed by Services (14,1%) and internal Applications with lower adoption rates.

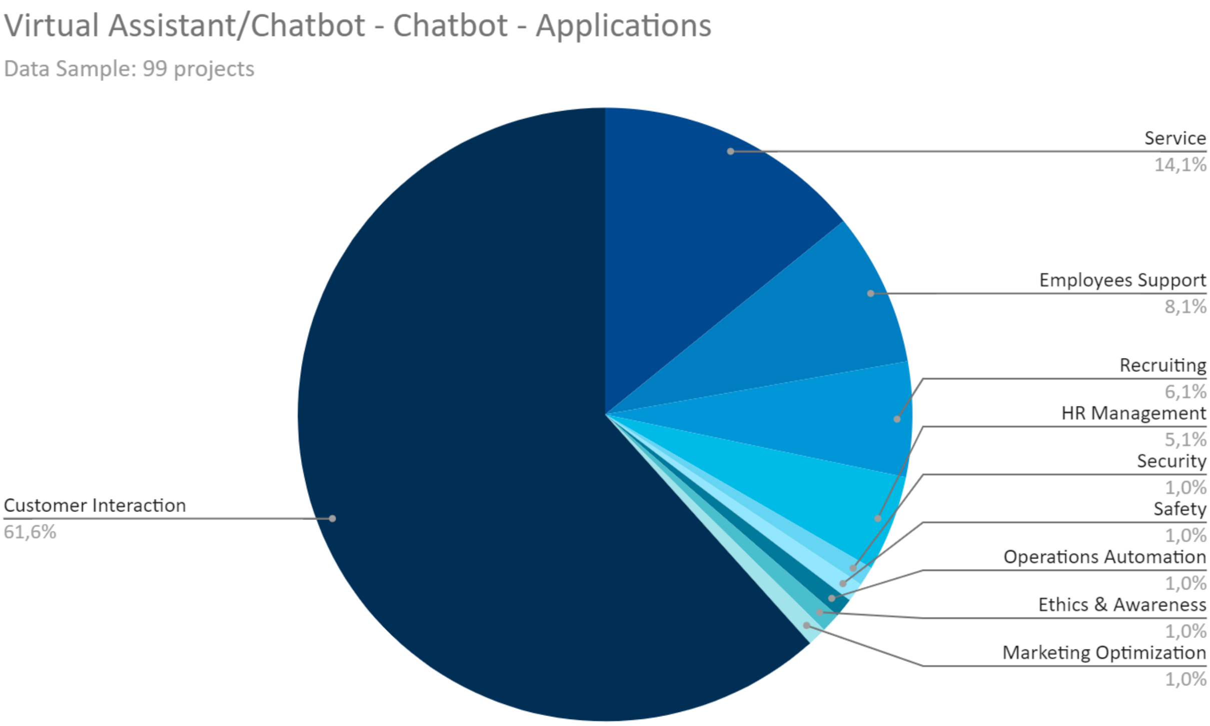

Figure 43: Distribution of Chabot Applications

In Customer Interaction, Chatbots are used to assist and support customers, handling their inquiries and providing them useful information: they are able to quickly and conveniently process customers' queries and questions, and answer them. Chatbots solutions are typically available in a variety of channels, such as SMS, WeChat or Messenger, websites and mobile apps of companies. Many of the analyzed solutions leverage Artificial Intelligence not only for Natural Language Understanding and Generation purposes, but also to learn the users' behaviour and accordingly personalize interaction over time.

As previously mentioned, Chabot solutions have been introduced in almost all the sectors; among the others, Banking & Finance and Insurance are the most interested ones. In Banking & Finance they are used to: provide information to customers; help them in managing their bank accounts and keeping track of their spending habits; provide advice on credit financial services; fulfil a multitude of requests, such as sending and scheduling payments.

For instance, China Construction Bank "Xiao Wei'' is an Artificial Intelligence based Chatbot available in WeChat, SMS and internet channels. Described as a 27 year-old girl good at ballet and piano, it can understand more than 600 different versions of one question and answer in 56 different Chinese regional dialects. Moreover, it can handle more than 80 kinds of banking services. In 2018 it handled 1,9 billion customer conversations, with advantages in labour cost, standardised quality services and improved service efficiency.[41]

On the other side, Insurance Chatbots are used to answer customers' questions about their insurance contracts, to make policy changes, to send application forms, to support through policy-buying and enquiry processes.

For instance, in 2018 Assicurazioni Generali France piloted "Leo", an Artificial Intelligence based Chatbot for its customers, to present its offer and answer their recurring questions via website. After a six month period, there were 100.000 recorded interactions with Leo and more than 2.400 calls avoided, with benefits also in terms of engagement rates, reduced support waiting times and better customer experience. The success of the pilot gave Generali the reassurance they needed to scale it also to mobile and tablet devices. [42]

When considering Services (14,1% of projects in Chatbots), the most diffused solutions are digital consultation services that companies in the Healthcare & Assistance and Pharma sectors offer to public and private medical institutions. These solutions use Chatbots powered by Artificial Intelligence to provide hospitals' patients health consultation and medical advice.

To conclude, internal applications for Chatbots are progressively taking hold. In Employees Support (8,1% of Chatbots), conversational bots powered by Artificial Intelligence capabilities are used to support employees, leveraging internal chat systems and the intranet. These solutions are used to provide workers with the information they need for executing their tasks and to answer routine questions. For instance, the French international banking group BNP Paribas has developed in 2019 a platform to create Chatbots for any internal team wishing to automate its response process to queries and questions from other internal collaborators. More than 50 teams distributed all over the world have shown interest and received tailored Artificial Intelligence based Chatbots.[43] Other interesting projects can be found in Recruiting (6,1% of Chatbot projects) and HR Management (5,1%), such as Chatbots to support the early stages of the recruiting process by handling questions from job seekers and Chatbots acting as career advisors, helping employees to identify skills gaps and suggesting materials to read accordingly.

Voicebot

When dealing with Voicebots, or Virtual Assistants, the main Applications are shown in Figure 44.

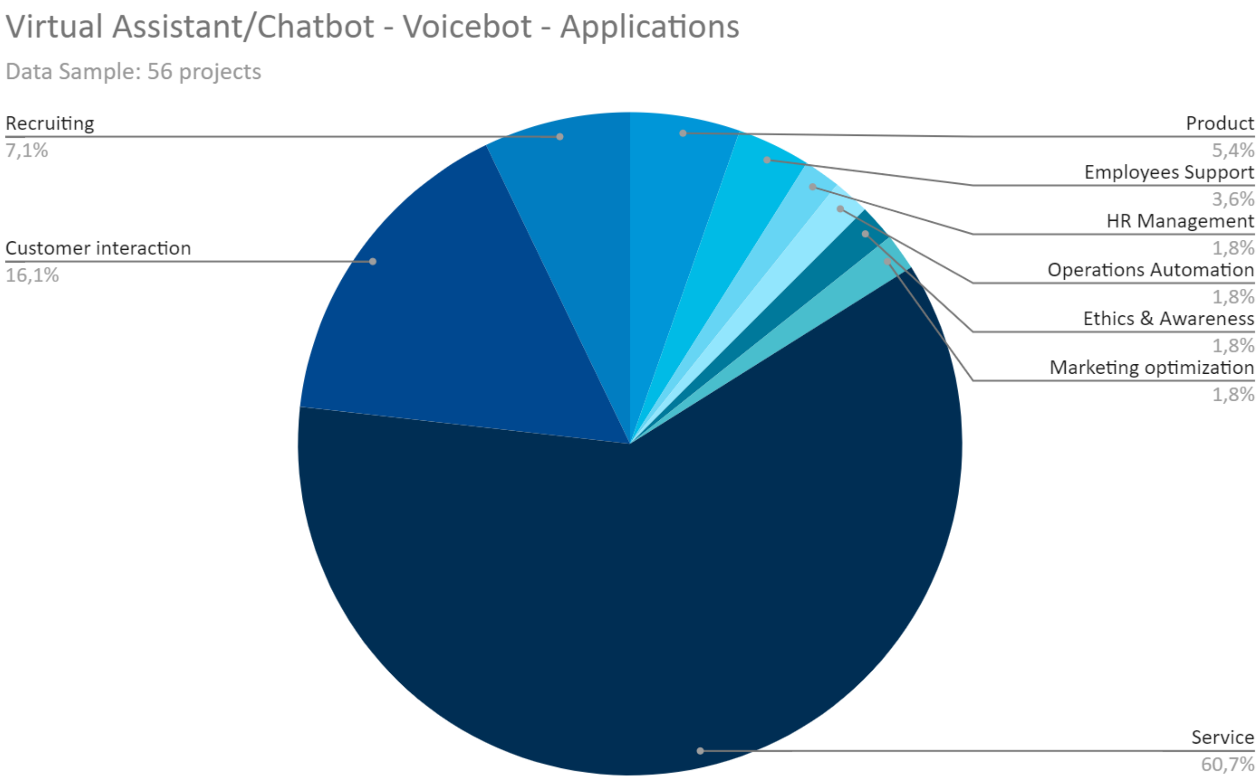

Figure 44: Distribution of Voicebot Applications

As shown, Voicebots are typically applied in Services, with the 60,7% of the Voicebot projects. By analyzing the initiatives in this Application, the solutions can be grouped in three main clusters.

A first group includes companies operating in the Technology sector that have developed popular and well-known Virtual Assistants such as Siri, Google Assistant, Cortana, Alexa, Samsung's Bixbi, Tencent's Xiaowei and many other ones. These Virtual Assistants are able to understand spoken commands and questions and accordingly answer or take actions, learning the habits and preferences of the users and adapting to them over time. They can be connected with phones, computers, TVs, speakers and other devices to have a conversation.

A second group contains companies operating in the Automotive sectors, such as Toyota Motor, Volkswagen Group, BMW Group, Honda, Mitsubishi and Fiat. These firms are starting to integrate Virtual Assistants into the operating system of their new vehicles, as additional onboard service. These systems enable a variety of tasks, such as searching routes and navigating the driver, modifying vehicles settings, providing information about weather and traffic conditions. Typically, the Virtual Assistants they are integrating in their vehicles are the ones of the Technology companies, mentioned in the first group. In some cases, automobile companies are developing their own Virtual Assistant, but this is almost always done in collaboration with external Technological Partners. Therefore, independently from the case, Technological Partners are typically involved for this category of solutions.

To conclude, a third group of solutions includes companies from the most various sectors that are starting to introduce voice-based services on the Virtual Assistants of the companies mentioned in the first group. For instance, the American retailer Walmart has launched as a pilot a "Walmart Voice Order" service, that enables users to order groceries by using devices such as tablets, watches, TVs and smartphones connected with Google Assistants, adding them directly to the Walmart Grocery cart. Anthem, a provider of health insurance in the United States, has developed a skill to support patients in maintaining their treatment plan through Alexa-enabled devices or Amazon Echo, assisting them with prescription refills, renewals and checking of orders.[44] EDF Energy, one of the UK's biggest energy providers, uses Amazon's Alexa to allow consumers to manage their energy accounts by just using their voice, delivering functionalities in areas such as account balance and meter reading submission.[45] Many other use cases can be found in different industries.

Moving to Customer Interaction (16,1% of the cases), the aim is similar to the one discussed with Chatbots. In this case, Voicebots are implemented in the call systems of companies, to understand the reason of calls and automatically answer to callers, or in customer service apps and other channels, to provide assistance for customer queries by understanding their intent and providing tailored responses.
For instance, Citigroup, an American multinational investment bank, has launched in 2020 a Voicebot solution powered by Artificial Intelligence within its call centers for conversations with the clients. The solution is now being extended in other regions, after that the pilot has resulted in enhanced customer satisfaction.[46]

To conclude, applications can also be found in Recruiting (7,1% of the cases), where Artificial Intelligence powered job interview systems are used to accelerate the hiring process by autonomously interviewing candidates. For instance, in 2018 PepsiCo's HR managers needed to fill 250 vacancies in two months in the sales department, as drivers and as workers: a Virtual Assistant solution developed by the Russian startup Stafory was used to phone and interview job applicants.[47]

In conclusion, several benefits resulting from the adoption of Chatbots and Virtual Assistant solutions to support the customers or the employees of the company can be identified, considering both the perspective of the user and the one of the company. The benefits resulting from the analysed projects have been collected, generalized and summarized below. They can be used to both justify the widespread diffusion of Chatbots and Virtual Assistants, as well as to summarize why companies that have not adopted these solutions yet should be interested. The most relevant advantages resulting from successful implementations of Chatbots or Voicebots to support customers or employees are:

- Customer support is provided 24/7, extending the hours of the customer service and allowing to resolve issues when they happen.
- Implementation results in better and standardized service, with improvements in the accuracy of the interaction.

- Reduced customer support waiting time and faster response time to customer queries and FAQs.
- As a result of the aforementioned benefits, better customer/employee experience and enhanced satisfaction
- Positive reactions from customers, with increases in customer engagement rates. The same applies for employee satisfaction in internal solutions.
- Reduction of service cost and improved service efficiency
- Decrease in volume of calls to call centers of the company and in the number of conversations handled by human agents, easing the pressure of manual customer service.
- Because of savings in working time and reduction in agent workloads, the use of these systems frees up an impressive amount of time for customer agents, making them free to assist people with the more complex queries or to focus on other opportunities.

3.2.4 Computer Vision

Computer Vision projects constitute the 17,7% of the mapped Artificial Intelligence initiatives, making Computer Vision the second Class of Solutions in terms of diffusion and pointing out a good adoption rate among companies. This Class is characterized by an intermediate level of maturity: the 61,5% of the initiatives are Operative or in Implementation, the 38,5% Pilots or Project Proposals; this also suggests that several companies are starting right now to explore the potentialities of Computer Vision solutions for practical applications. The wide diffusion of this Class of Solutions is also justified by its possibility of application in a variety of contexts, so that companies operating in very different industries can use Computer Vision solutions at their advantage for several purposes. Figure 45 confirms how Computer Vision is finding adoption in all the analyzed sectors.

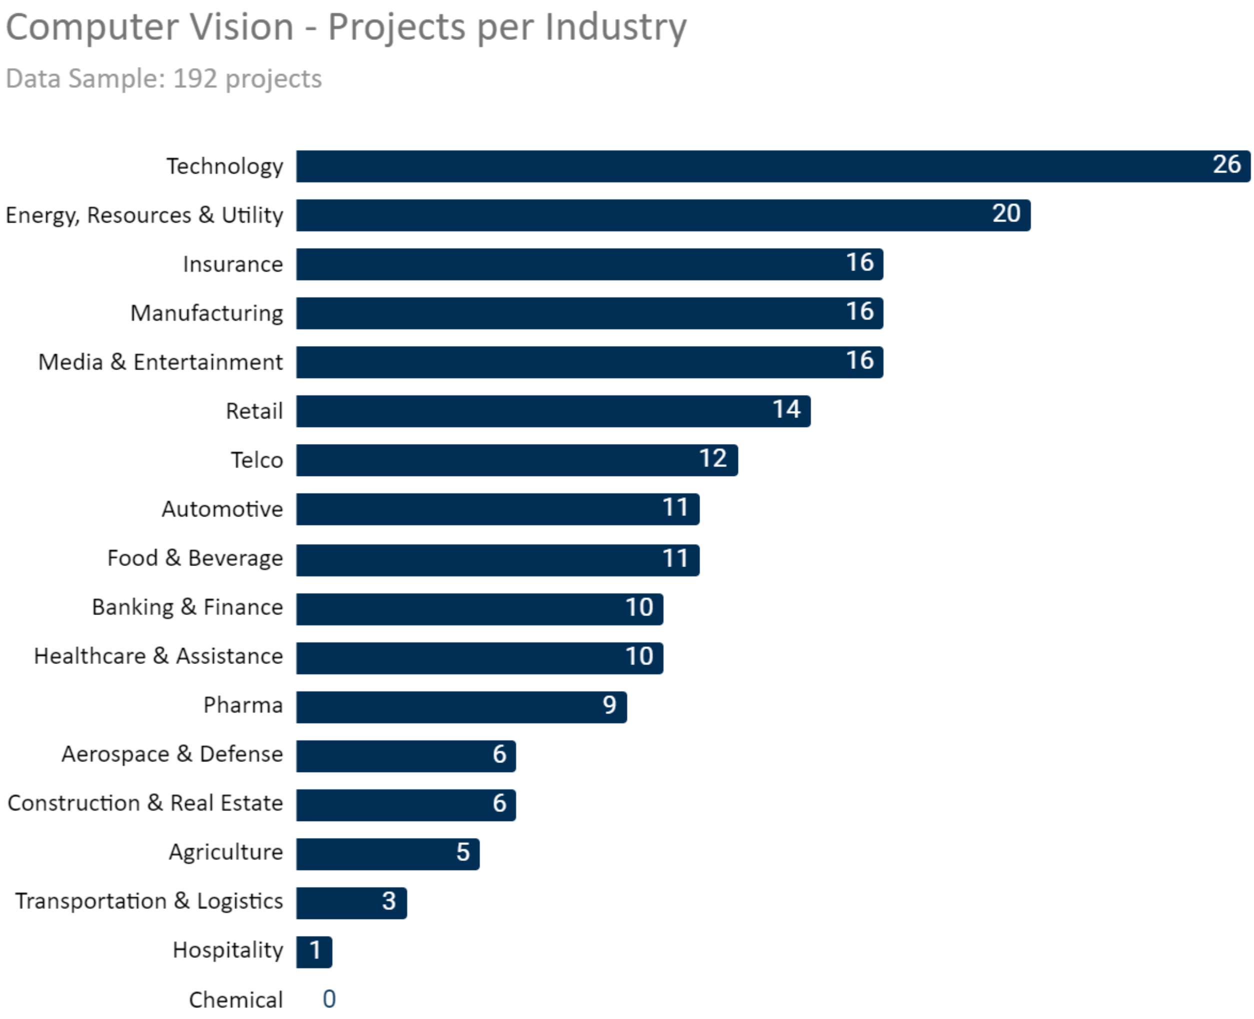

Figure 45: Distribution of Computer Vision Projects per Industry

If we move to analyze the distribution of the different Computer Vision's Specifications, results are shown in Figure 46. The most diffused typology of Computer Vision solutions regards Image Analysis, with the 64,1% of the mapped projects, corresponding to 123 initiatives out of 192.

Biometric Recognition (12,5%), Video Analysis (12%), and Image & Video Editing (11,5%) follow, with similar percentages among them. Since this level of analysis does not allow to obtain meaningful insights about the application of Computer Vision solutions yet, hereafter each Specification is analyzed in detail.

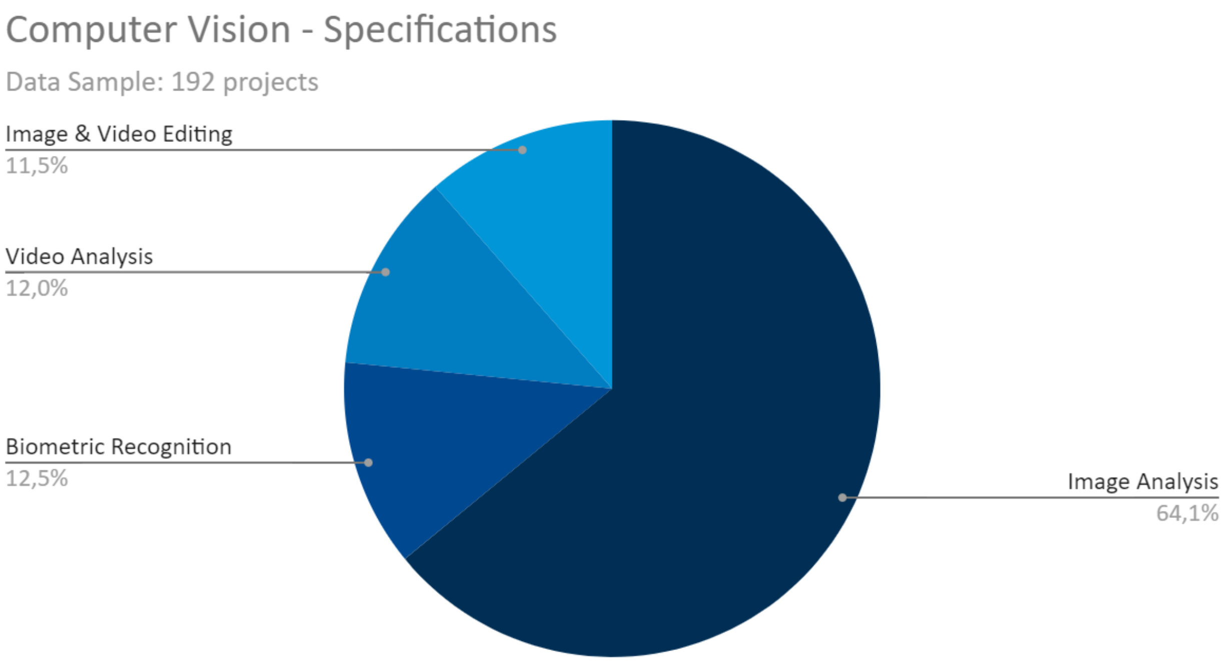

Figure 46: Distribution of Computer Vision Specifications

Image Analysis

Image Analysis, the most diffused Specification in Computer Vision, is characterized by a variety of possible practical Applications, as shown in Figure 47. Among the others, three are the most common ones: Service (29,3%), Quality Control (14,6%) and Operations Automation (13%).

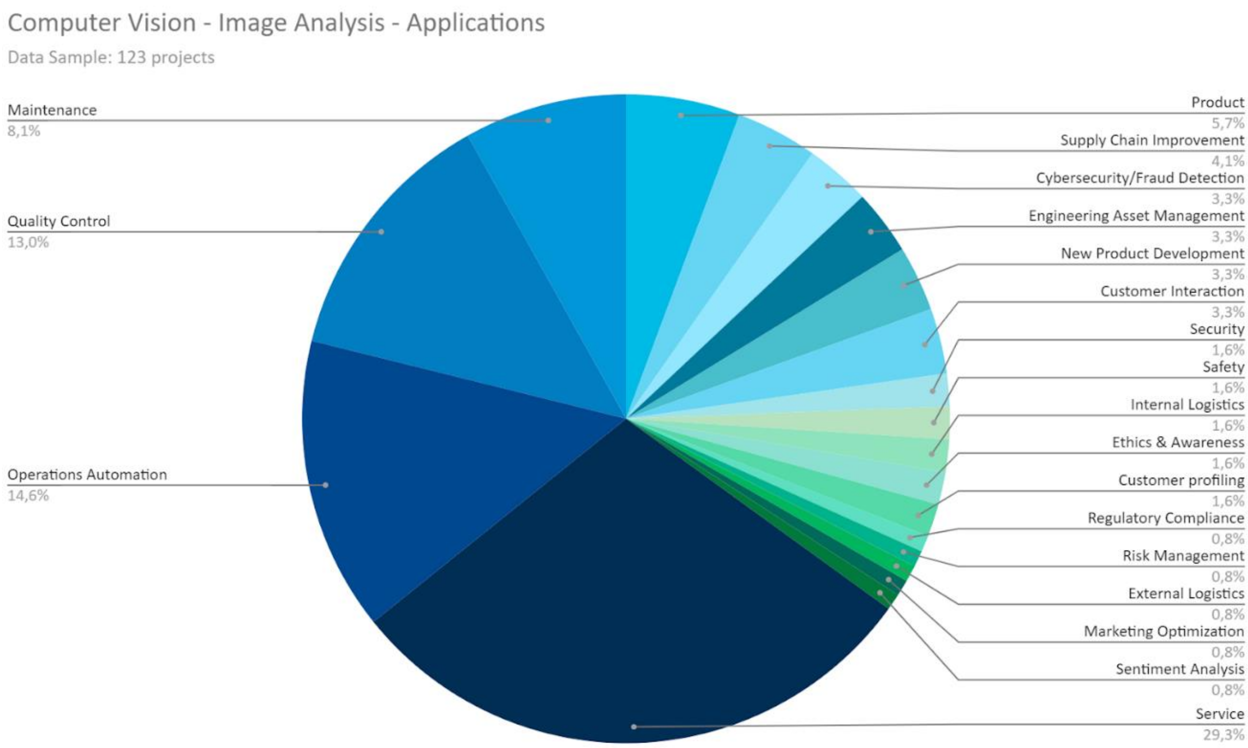

Figure 47: Distribution of Computer Vision Image Analysis Applications

When dealing with Services, about half of them are image-reading software systems that companies in Healthcare & Assistance, Pharma and Technology sectors are developing and selling to private or public medical institutions. These systems represent a quite mature use of Artificial Intelligence for the analysis of medical images, and deploy Convolutional Neural Networks to support disease screening, diagnosis and evaluation of treatments. Analyzing a variety of image formats, from computed tomography scans to ultrasound images, they can for instance support detection and early diagnosis of several typologies of cancer, and have helped Chinese medical institutions during the ongoing COVID-19 global pandemic. The advantages of adopting these solutions are notable, above all the accuracy and speed of detection and the possibility to support institutions in the analysis of challenging cases and diagnosis double checking. For instance, Ping An Smart Healthcare, Chinese subsidiary of the Insurance giant Ping An, has developed over 40 medical imaging models leveraging Artificial Intelligence, covering nine major human body parts and supporting disease screening and diagnosis. These systems have proven to be helpful not only in city hospitals but also in China rural areas, notoriously characterized by a shortage of medical personnel. [48]

Other Services are variously distributed between B2B, B2C and B2G markets. For instance, in a B2B scenario several digital systems are proposed to farmers and vegetable growers, analyzing fields images and returning key insights about them; they can be used to monitor the color and height of crops, to detect diseased plants or if they are infested with insects, to identify the boundaries of fields.

In a B2C market, for instance, Google's photo storage service Google Photos enables you to search photos by taping in names: if you want to find pictures with your dog, typing "dog" Deep Learning algorithms are used for Object Detection tasks, providing as a result only images classified as "dog".[49] Or again, L'Oréal's Perso mobile app analyzes face pictures that users can take with their smartphone camera. Artificial Intelligence is used to evaluate the user's skin condition, considering deep wrinkles, fine lines, eventual dark spots, and pore visibility.[50]

To conclude, B2G solutions range from Computer Vision analytics to monitor parking spot occupancy, to updating the congestion status of tourist spots by counting pedestrians, to identification of targets in military missions.

Quality Control, on the other side, includes the 13,0% of initiatives in Image Analysis. In this case, as for Operations Automation, the use is exclusively internal. Quality Control implies the use of inspection systems leveraging Artificial Intelligence to recognize manufacturing or operational defects and anomalies along production and assembly lines. This enables the automatic inspection of products and parts and the optimization of quality controls: avoiding conventional inefficient and inaccurate inspections performed by human eye, improvements in efficiency and accuracy are achieved. Companies in the Automotive and Manufacturing sectors are the most interested in internally adopting this kind of solutions.

For instance, in the Automotive industry, BMW Group is piloting several projects: one of them uses a camera to take pictures of components, and then an Artificial Intelligence software is used to distinguish components with micro-cracks from those without. In other projects, Computer Vision is used to check if the hazard warning triangle has been correctly fitted to the car, or to control if the correct model designations are attached, comparing vehicle order data with images of the model designation of a car under production.[51]

In the Manufacturing sector, Midea Group, a Chinese manufacturer of home appliances, is implementing a Deep Learning based inspection system on the cloud, with the aim to automatically identify defects in tested products, such as a missing screw, a missing nameplate, or a logo silk screen defect. The solution has allowed the company to improve fault detection rates and product quality, as well as to reduce material costs by 30% in their microwave oven visual inspection project.[52]

To conclude, Image Analysis in Operations Automation represents the 14,6% of the initiatives. Among the others, companies operating in the Energy, Resources & Utility and Insurance sectors are the most interested ones.
In the Insurance sector, Computer Vision algorithms are used to assess severity of damages to buildings, cars, and other assets, transforming a traditionally manual process into an automated one. In this way not only cost decreases, but also productivity and efficiency are improved, increasing customer satisfaction.
For instance, the Japanese insurer Tokio Marine Holdings is piloting the use of Computer Vision to analyze vehicles damage in its claim process, using Machine Learning trained on millions of photos of car damages to assess damage and estimate both the needed repair operations and labour hours. Considering that in Japan it can take 2-3 weeks to determine the amount to pay, the objective is to shorten this time from days to minutes.[53]
On the other side, companies in the Energy, Resources & Utility sector are using Computer Vision technology for the automation of very industry-specific tasks, such as the analysis of microscopic images of rock sections and subsurfaces images, to obtain additional information for faster and more effective oil and gas fields' exploration and assessment.

To conclude, several other Applications with lower rates of adoption exist. Describe all of them is not possible, but they include: the internal use of Artificial Intelligence for Maintenance purposes, analyzing images to assess the conditions of assets; Computer Vision's use in consumer products, such as refrigerators identifying food inside through cameras, or washing machines recognizing color and style of clothes and making related suggestions; solutions analyzing images of shelves and determining any out-of-stock item that needs replenishment in Retail.

Biometric Recognition

Biometric Recognition constitutes the 12,5% of Computer Vision initiatives and is becoming increasingly adopted in China, where it is used more and more for payments in shops and supermarkets, and for commercial transactions. This is the reason why a relevant part of Biometric Recognition initiatives have been launched in China and are already Operative.

As Figure 48 shows, the most diffused Application for Biometric Recognition is Customer Interaction (66,7%), considering the way consumers can interact with companies and their services. In this field, facial recognition solutions enabled by Artificial Intelligence are used to perform authentication of people when accessing online services and making payments, allowing verification of user identities. These solutions entail improvements in terms of both security and customer experience, with convenient and fast identity verification.

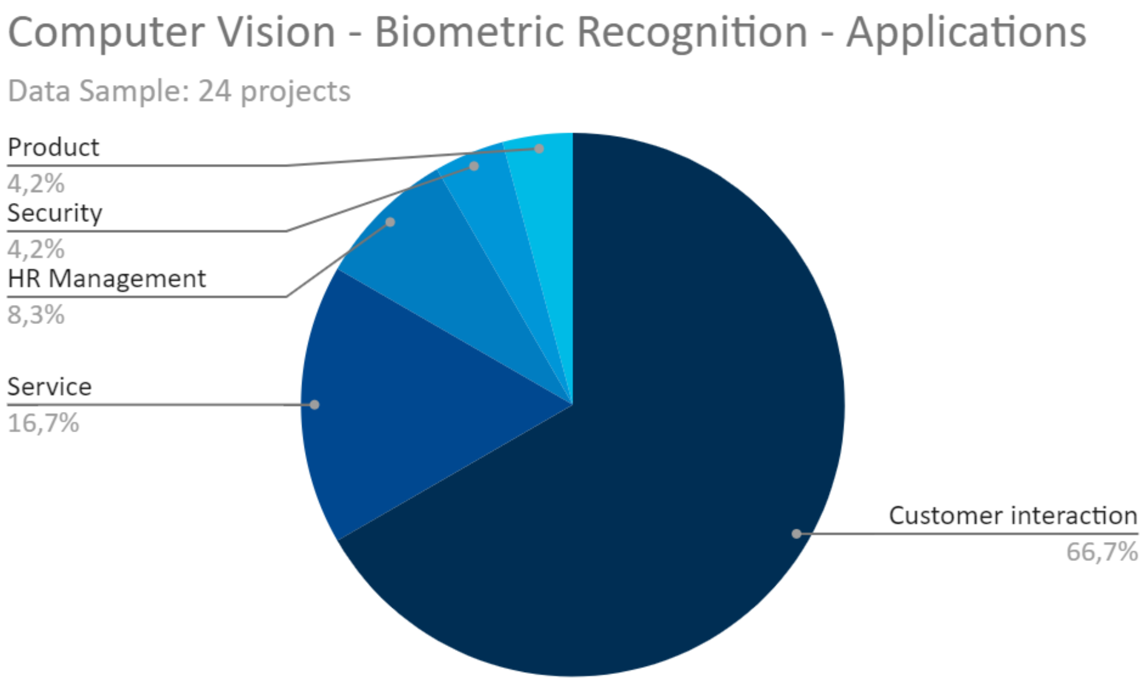

Figure 48: Distribution of Computer Vision Biometric Recognition Applications

One sector particularly interested in the use of Biometric Recognition for Customer Interaction purposes is Banking & Finance: Chinese banks are extensively adopting Artificial Intelligence solutions for authentication, introducing facial recognition to withdraw cash and make payments at ATMs or in mobile phone banking applications.

The use of Biometric Recognition eliminates the need for typing lengthy personal ID and security codes along the authentication procedure, allowing a simpler and much more secure identification.
For instance, the Agricultural Bank of China has launched in 2017 a withdrawal service based on facial recognition at over 20.000 branches across China; clients have to verify their identity through Biometric Recognition at the ATM before withdrawing cash.[54]

Even if the great majority of initiatives are in Customer Interaction, Biometric Recognition can find application in other contexts too. For instance, if considering B2C Services, facial recognition is used in Google Photos to automatically organize photos according to people's faces.[55] Similarly, Facebook uses face recognition to automatically recognize people in pictures and suggest tags.[56] Furthermore, in HR Management Biometric Recognition can be used to ensure that employees attend meetings instead of asking their friends to sign them with their staff cards, as well as to use data from identity verifications to improve the fine management of attendance and payments.

Video Analysis

Moving to Video Analysis (12% of Computer Vision projects), two main Applications can be identified: Safety and Security. Indeed, despite results shown in Figure 49, many of the B2B/B2C/B2G Products and Services mapped are Smart Surveillance solutions for final consumers, enterprises, or municipalities, with Safety and Security purposes. These solutions often consist in the only surveillance service, with a cloud-hosted video management system that can support existing cameras; in other cases, also smart cameras are sold, capturing high-quality videos and transmitting them to the cloud for further analysis. Meanwhile, the percentages of Security and Safety Applications in Figure 49 consider just the internal adoption of Video Analysis solutions based on Artificial Intelligence. By further elaborating the results to consider Safety and Security purposes, independently from their sale as Products/Services or internal use, Safety solutions reach the 21,7% of the initiatives, Security the 4,3%, with a total percentage of 26%.

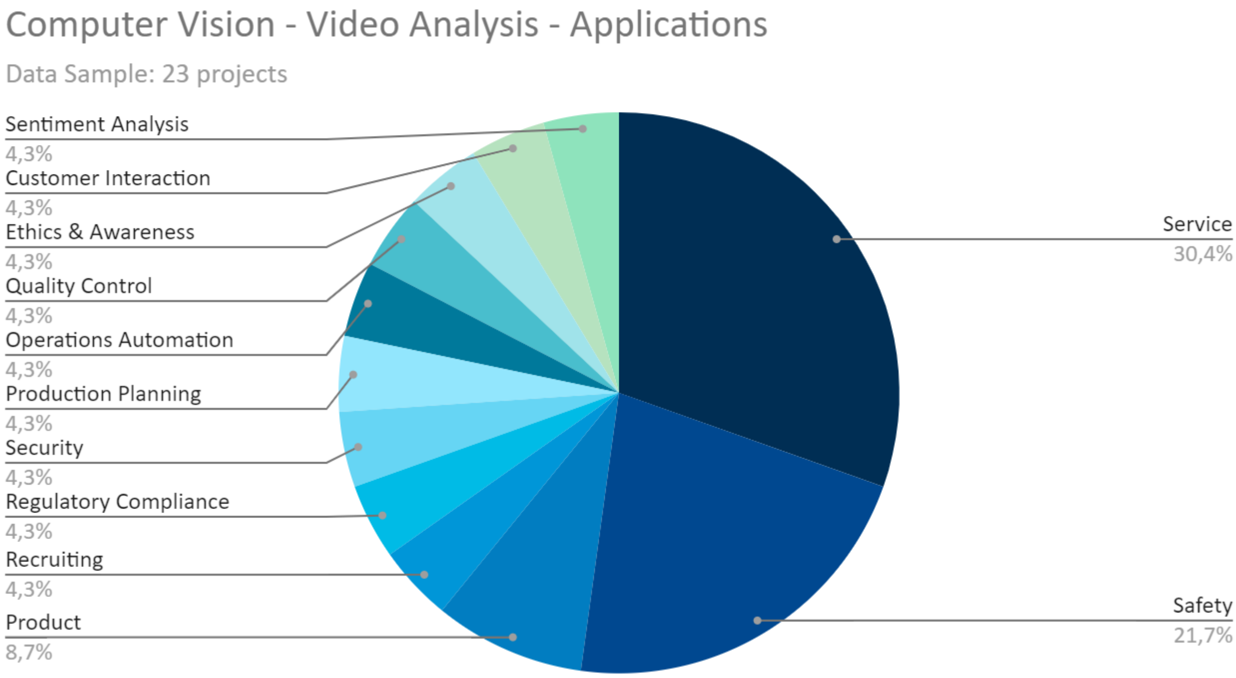

Figure 49: Distribution of Computer Vision Video Analysis Applications

In both Safety and Security Applications, Computer Vision solutions are used to analyze videos with the aim to identify unusual and abnormal behaviours and trigger alerts.

Safety solutions are typically applied in a business context, with particular interest from companies operating in the Energy, Resources & Utility and Construction sectors. Artificial Intelligence vision algorithms are used to analyze real-time videos from cameras to monitor safety conditions in plants and construction sites, detecting and reporting eventual safety risks. Benefits result in improved safety compliance, prevention of accidents and incidents, faster alert, and lower responses time in case of accident.
For instance, Saudi Arabian Oil Company has installed video cameras in fixed positions of its plants to monitor movements of machinery and personnel.[57] The real-time video is analyzed by an Artificial Intelligence based engine to autonomously detect and report safety risks. In another project, Royal Dutch Shell is piloting a cloud-based Deep Learning solution to analyze videos from surveillance cameras, with the aim to ensure safety of employees and customers at its Shell retail sites: events of interests, such as vehicles hitting someone, thefts, smoking, wrong refuelling behaviours are identified and employees automatically alerted.[58]

Security solutions, on the other side, regard not only businesses, but can be also be used for the Smart Home in a B2C context and the Smart City in a B2G one. The aim, independently from the market, is to improve security, by alerting in case of unusual situations' detection.
For instance, Walmart, the American multinational retail corporation, uses Artificial Intelligence in thousands of its stores: videos from surveillance cameras are analyzed for real-time detection of thefts at registers and self-checkout kiosks. If a customer places an item in a bag without having it previously scanned, the system recognizes it and automatically informs store associates.[59]
In a B2C context, Google Nest Cam Outdoor security camera uses Artificial Intelligence to perform scene recognition tasks, understanding if living beings wandering around the property are animals or individuals and sending alerts if needed.[60]

While the 26% of the projects have Safety and Security purposes, the remaining part is made of individual initiatives in a variety of industries and applications, underlying once again the multitude of potential application scenarios for Computer Vision technology. Initiatives includes: the early exploration of video analytics for hospital room monitoring, to understand group dynamics, behaviour, intent and to develop situational awareness;[61] a pilot by Disney to detect facial reactions of movie audiences, such as smiling and laughing in cinemas, for audience analysis purposes; [62] the use of Artificial Intelligence classifiers by YouTube to identify and flag objectionable videos supporting racism and terrorism, removing them.[63]

Image & Video Editing

Image & Video Editing, as previously mentioned, currently represents a minority of initiatives (11,5%) in the Computer Vision Class of Solutions, pointing out a still limited diffusion for this Specification.
Figure 50 summarizes the main results in terms of Applications, suggesting that this typology of solutions currently has not only a limited diffusion, but also a limited number of application fields.

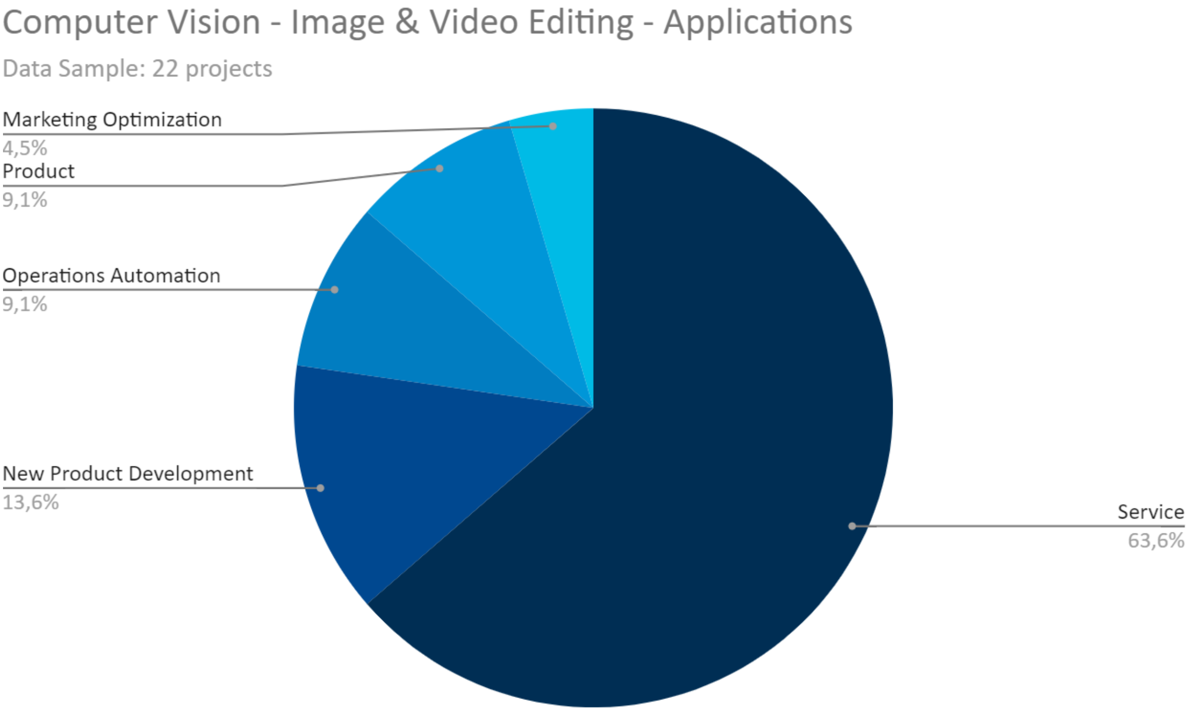

Figure 50: Distribution Computer Vision Image & Video Editing Applications

However, the most relevant Application for Image & Video Editing solutions based on Artificial Intelligence is in Services (63,6% of the projects), typically intended for a B2C market.

The dominant typology of solutions is represented by mobile applications, or specific functionalities of services, developed by Big Tech companies and allowing the users to edit photos and videos. For instance, Machine Learning can be used to detect subjects in a black and white photo and automatically color their skin and clothing, to find patterns in images and amplify them to create trippy versions of the same image, to transform live videos into something resembling the work of Van Gogh, Picasso and other artists, or to change videos' backgrounds.

Other use cases in B2C regard the use of Artificial Intelligence to virtually interact with products using Augmented Reality. For instance, in 2018 the American retailer Target has introduced a functionality for virtually trying on makeup and cosmetics. Using Artificial Intelligence, high resolution images of products are applied to the user's face, scanned through the smartphone camera or the computer webcam.[64] Similarly, L'Oreal has

launched a mobile app powered by Artificial Intelligence, StyleMyHair, that allows users to virtually try on different haircuts and colors.[65]
Services can also be intended for a B2B market: for instance, medical software for medical institutions in the Healthcare & Assistance industry have been developed, using Artificial Intelligence to modify medical images and scans to make them more intuitive, or analyzing them for building 3D anatomic reproductions of body parts.

Apart of the abovementioned solutions, the few other mapped projects are individual projects that range from the use of Artificial Intelligence to integrate branded contents into movies and TV series, by substituting for instance a cup of coffee with a branded cup of coffee, to its use for improving the quality of games and movies, removing noise or applying various effects.

Critical Issues
To conclude, the peculiarity of this Class of Solution, with the use of image and video data and its particular applications, allows to outline some critical issues specific for Computer Vision projects, representing either obstacles in their implementation or general concerns that could limit their adoption. General Critical Issues for Artificial Intelligence projects are, instead, discussed in Section 3.3
As regards the implementation of Computer Vision solutions, the main Critical Issues identified are:

- Huge amount of Training Data needed to train the model. While the availability of a sufficiently large database for training the algorithm is a requisite for any Artificial Intelligence project, this becomes even more relevant, for Computer Vision tasks: an impressive number of labeled images is required to train Convolutional Neural Networks.

- Images and videos to be processed may contain a lot of noise or be affected by ambient factors, such as lighting; similarly, items can be obscured by shadows or be at an odd angle, or their identification in front of a complex background can be difficult. Therefore, false detections may occur. This issue is partially related to the previous one, meaning that a wider amount of training data enables the system to learn better features invariant to changes.

On the other side, some Critical Issues are represented by general concerns that could limit the adoption of Computer Vision solutions.

- Reliability of the solution: in some of the analyzed projects, adoption of Computer Vision solutions is discouraged because of concerns about their reliability and frustration in case of low accuracy. While motivations for this can be attributable to the implementation issues previously described, a regular number of occurring false detections may strongly limit the diffusion of Computer Vision solutions. For instance, the Video Analysis solution of Walmart previously described has been nicknamed by employees “NeverSeen” because of its frequent mistakes: they claim that it frequently misidentifies innocuous behaviour as theft, while failing to stop actual instances of stealing.[66] Similarly, face recognition can be prone to failures in case of people with beard and wearing glasses.

- Privacy concerns: especially in China where facial recognition is becoming part of daily life, consumers are raising concerns about data being held on them. On one side, they are worried because of the possibility of hacking and massive data breaches, already happened in some situations. In other ones, the fear regards the exploitation of these data for unfair uses by businesses or government, such as tracking the population. Privacy is also the reason why Facebook will pay $550 million to a group of users in Illinois, since its Biometric Recognition solution for facial tagging has violated the state’s privacy laws.[67]

- Ethics concerns: to conclude, some concerns regard the ethical implications of specific solutions based on Computer Vision. For instance, in Image & Video Editing, face swapping technology could be used to create credible fakes of something that did not really happen. Similarly, in Insurance, pilot projects analyzing facial micro-expressions for lie detection of borrowers are raising concerns not only about their reliability, but also about the ethics of using similar solutions.

3.2.5 iRPA

The iRPA Class of Solutions includes the 3,8% of Artificial Intelligence projects, suggesting a still limited diffusion, if compared with other Classes of Solutions. This means that most of the companies are still using Programmed RPA, and the shift to Artificial Intelligence Supported, or Driven, RPA is still far from being reality.

As shown in Figure 51, iRPA have been applied for automation of repetitive tasks in different industries, but two sectors above all seem to be particularly interested in the adoption of iRPA solutions: Banking & Finance and Insurance, gathering the 41% of iRPA projects.

This can be explained by considering that Banking & Finance and Insurance are some of the most data-intensive sectors, and so they can particularly benefit from the adoption of solutions for Intelligent Automation.

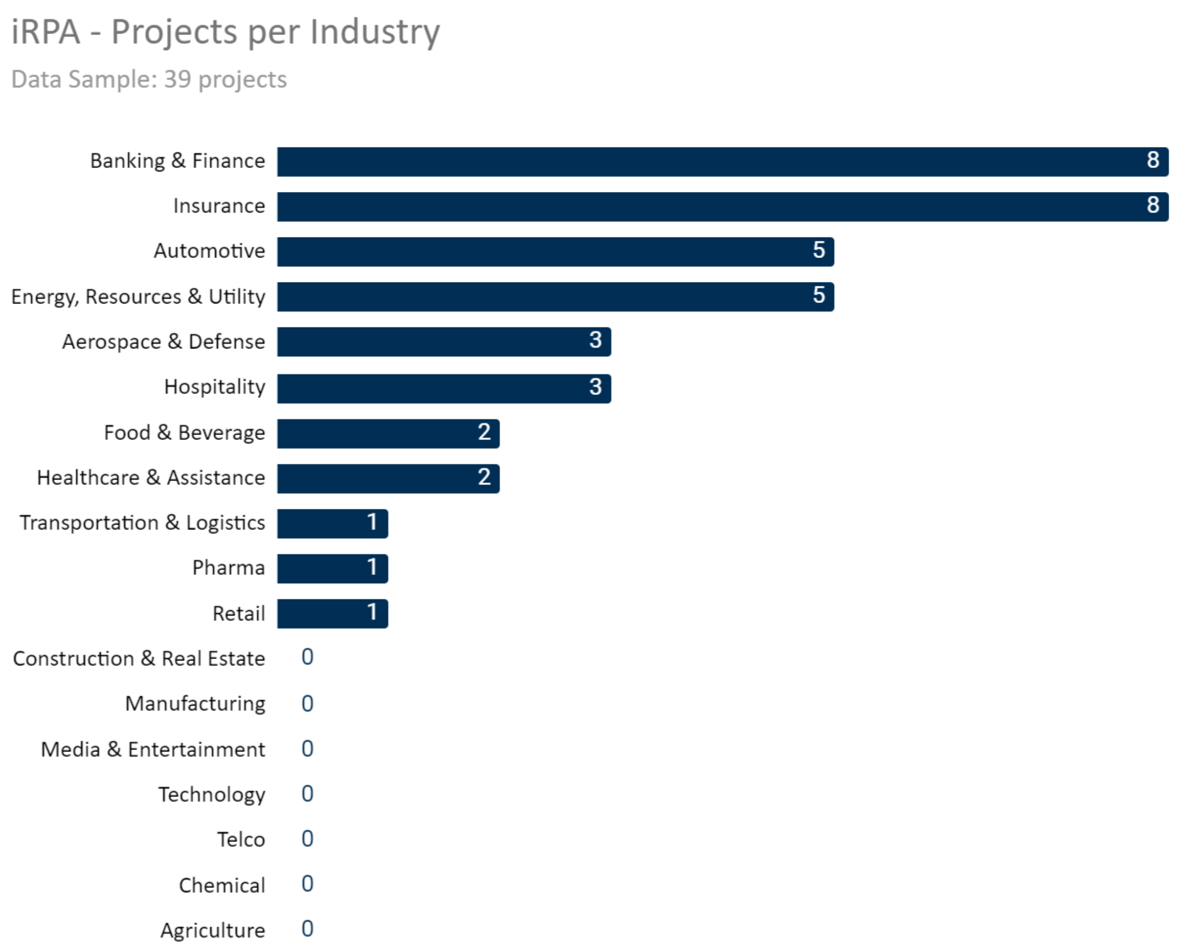

Figure 51: Distribution of iRPA Projects per Industry

When considering the Status of the Project, the majority of the projects are Operative (64,1%) and in Implementation (15,4%), while the remaining 20,5% are Project Proposals (12,8%) or Pilot Projects (7,7%), suggesting that new companies are starting to embrace solutions combining traditional RPA and Artificial Intelligence.

The typical enterprise use of iRPA solutions is internal, for the automation of back-end processes, so that the 89,7% of the analyzed projects are B2B Internal.

When considering the Applications for iRPA solutions, as it can be expected, the most diffused Application is Operations Automation, with 82,1% of the projects (Figure 52). If considering also iRPA software solutions (Services) sold to other companies or the government, this percentage reaches the 92,4% of Intelligent RPA initiatives. This is because iRPA is, because of its intrinsic nature, strictly related to the automation of tasks and processes. Only a very limited number of projects fall into other Applications, such as HR Management (2,6%) and New Product Development (5,1%). However, these solutions still refer to the use of Intelligent Automation for automation purposes, just applied in HR Management and New Product Development processes.

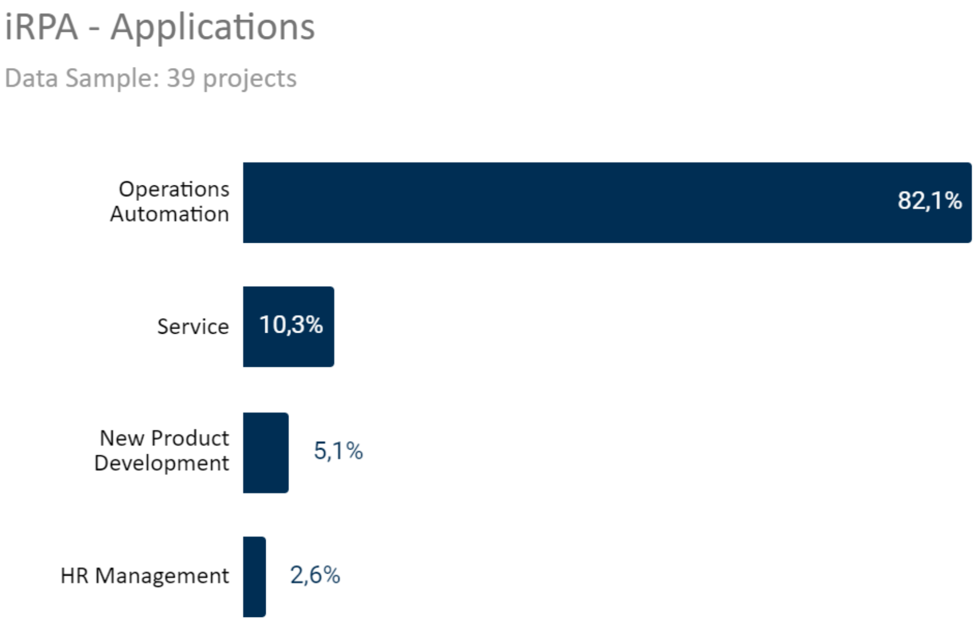

Figure 52: Distribution of iRPA Applications

Particularly, Operations Automation in iRPA includes all those projects combining traditional RPA with Artificial Intelligence, to enable the intelligent automation of repetitive tasks and processes that otherwise would be carried out by employees. The analyzed solutions are strongly dependent on the specific task to automate, so that it is not possible to provide a general description of their functioning, but just interesting or self-explanatory use cases.

For Instance, BNP Paribas, the French international Banking group, has automated the processing of asset servicing documents using a combination of Natural Language Processing, Machine Learning and RPA. The system is able to automatically detect, extract and classify relevant data from documents such as fund prospectuses, order confirmations, corporate events announcements and tax documentation. The resulting dataset is then fed into the bank operational system through the use of RPA technology. This solution has allowed the company to automate the processing of 500.000 documents in a year, with better straight-through processing rate and reduced services turn-around times.[68]

Another example is provided by Électricité de France, one of the world's biggest producers of electricity. The company partnered with the startup EmailTree AI to take advantage of its productivity tool for automation of the email management process, increasing the velocity of response. Their solution uses Natural Language Processing, Machine learning and Robotic Process Automation to faster reply to emails, automate email composition and trigger automatic actions in databases, enabling Intelligent Automation for the whole email management process. For instance, the solution is able to detect the intention of emails, identify the needed tasks to perform, prioritize them and use RPA to carry out the needed tasks.[69]

To conclude, the benefits coming from the implementation of iRPA solutions are several, considering both the perspective of the company and the one of the customers. One of the main reasons to use Intelligent Automation is that the automation of a process leads to savings in cost and time needed to perform the process. Additionally, the time that employees save in not carrying out anymore a specific task, can be used to focus on more value-added activities relevant for the business growth. To conclude, automation also leads to better performances in terms of quality and accuracy, that is translated into a better experience for the final customer.

3.2.6 Recommendation

The Recommendation Class of Solution includes the 5,2% of Artificial Intelligence projects. This diffusion is reasonable, if considering that Recommendation Systems have a limited applicability in some industries, because of their intrinsic characteristics. Figure 53 shows the distribution of projects among the different sectors.

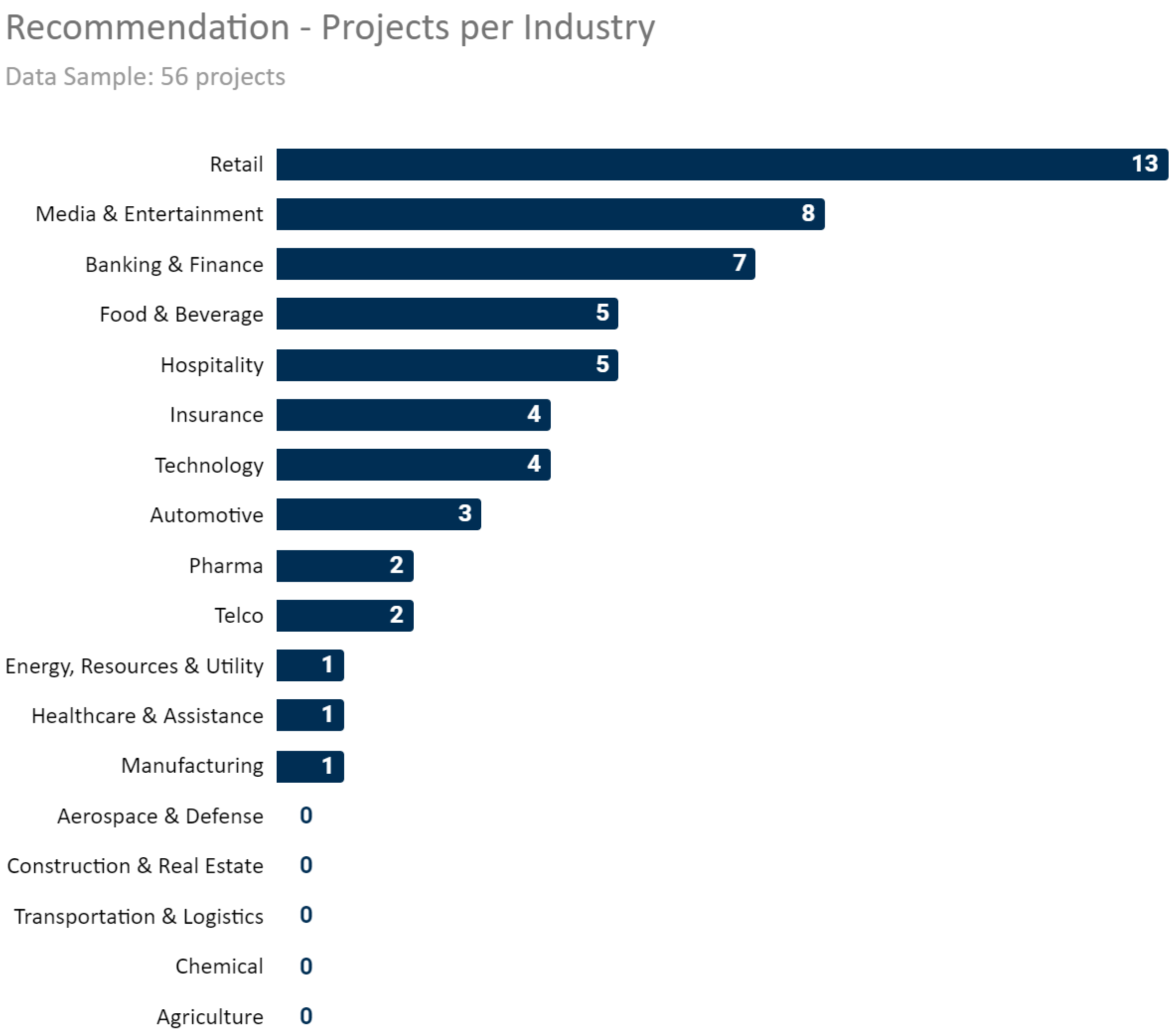

Figure 53: Distribution of Recommendation Projects per Industry

As shown, Retail and Media & Entertainment are the industries in which Recommendation Systems are most diffused, because of the possibility to use Artificial Intelligence for respectively Purchasing and Contents Recommendations. By considering these industries in which Recommendation Systems have a great applicability, it results that the 75% of Retail companies and the 71,5% of Media & Entertainment companies have at least one project for this Class of Solutions. As a consequence, it can be said that Recommendation Systems are characterized by a very high diffusion in those sectors particularly relevant for their applications.

On the other side, when it comes to analyzing the Status of the Project, it results that this Class of Solutions is the one with the highest maturity, along with Virtual

Assistant/Chatbot. Indeed, the 82,1% of the projects are Operative (71,4%) or in Implementation (10,7%), while Project Proposals (8,9%) and Pilot Projects (13%) constitute the 21,9% of the remaining cases.

To conclude, the distribution of Applications for Recommendation Systems is shown in Figure 54. As shown, the most relevant Applications are Marketing Optimization (50,0% of the projects) and Service (37,5%), followed by other Applications with lower diffusion. However, it is more meaningful to analyze this information at the Specification level.

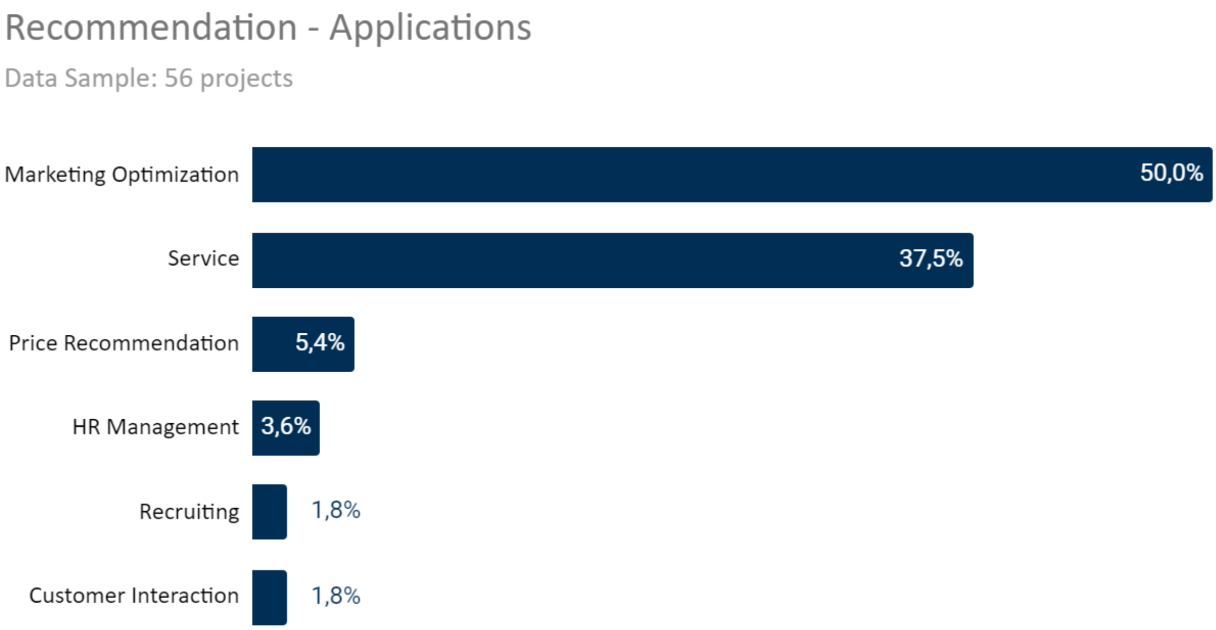

Figure 54: Distribution of Recommendation Applications

When considering Specifications, Figure 55 shows how the most diffused Specifications are Purchasing Recommendation (55,4% of the analyzed projects) and Contents Recommendation (30,4%), while only a minority of project regards Online Advertising (8,9%) and Dynamic Pricing (5,4%).

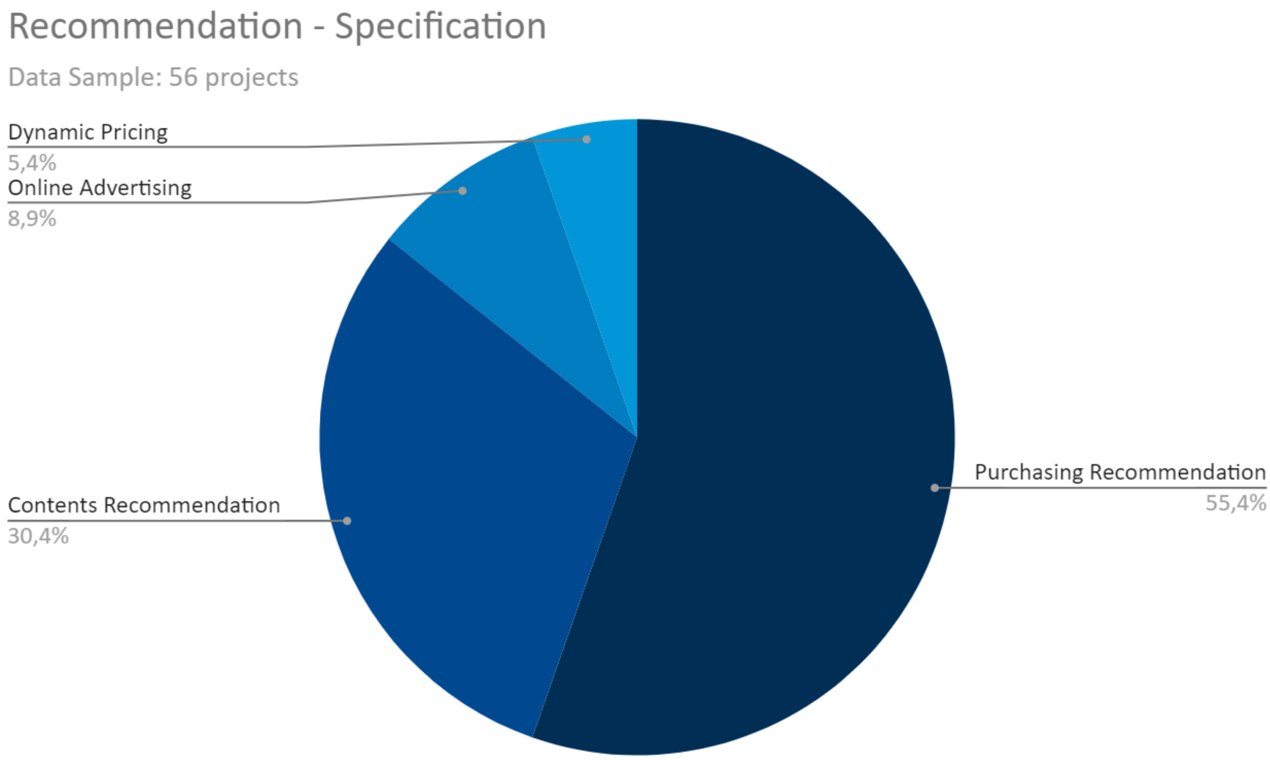

Figure 55: Distribution of Recommendation Specifications

Purchasing Recommendation

Generally, Purchasing Recommendation refers to the use of Recommendation Systems powered by Artificial Intelligence to suggest a user the purchase of products or services aligned with their interests. Purchasing Recommendations enable companies to increase their sales with less marketing effort required, since products and services are directly offered to interested clients. As shown in Figure 56, Purchasing Recommendation projects are associated with only two possible Applications, that are Marketing Optimization (71% of the projects) and Services (29%).

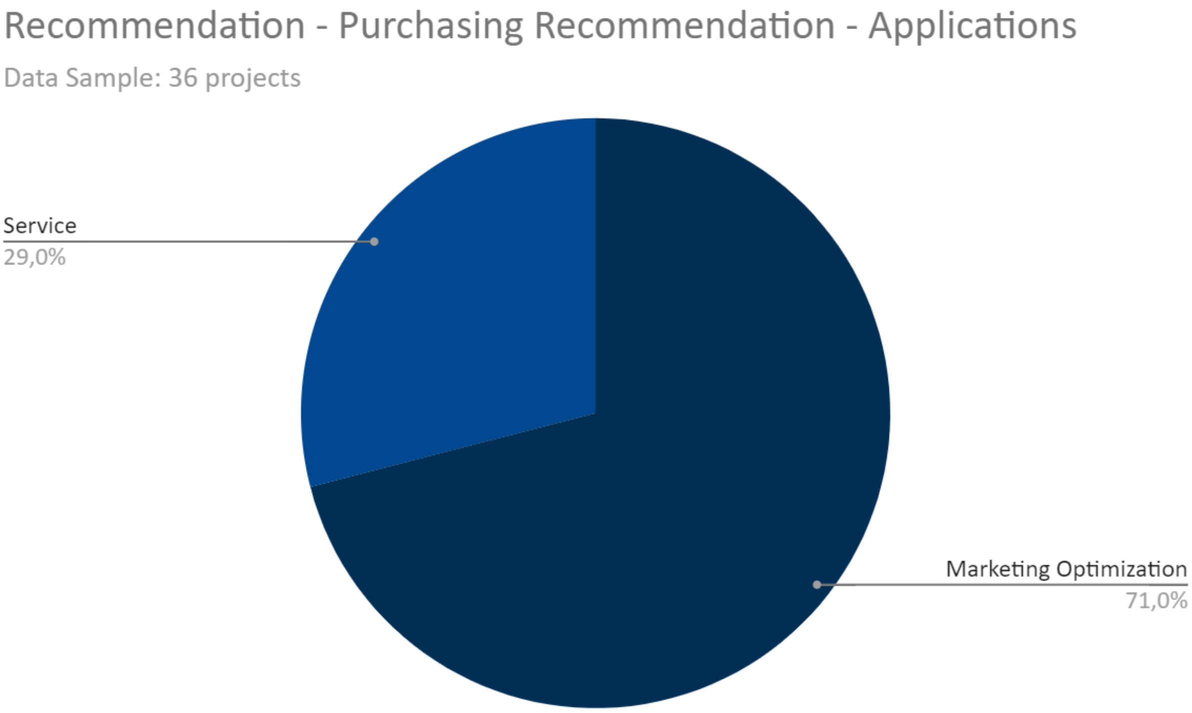

Figure 56: Distribution of Purchasing Recommendation Applications

Particularly, Marketing Optimization refers to the internal adoption by companies of Purchasing Recommendation Systems to enable Marketing Optimization.

For instance, in the United States McDonald uses a Recommendation System powered by Artificial Intelligence at thousands of drive-throughs. The system can automatically alter individual drive-through menu panels, based on factors such as the time of day, weather, popularity of certain menu items and the length of the wait. In addition to this, when the user ends his order the menu panel shows a new list of recommendations, fostering the customer to order more: for instance, if the client has ordered a burger, the systems will suggest a meal upgrade, while if he has ordered a salad the system will recommend a related product such as a bottle of water. Using this technology McDonald has been able to increase drive-through sales, automation in restaurants, speed up delivery and reduce inaccuracies in order taking.[70]

In another project Walmart, the American retailer, upgraded its website to serve customers Purchasing Recommendations. In this case, product recommendations are based on the past behaviours of customers and the location in which they connect to make the purchase.[71]

On the other side, the Application Service in Purchasing Recommendation includes all those solutions offered to a B2B or B2C market and based on Purchasing Recommendations algorithms. An industry particularly interested is Banking & Finance, with the development of personal wealth management platforms and smart investments advisory solutions; in this case, Purchasing Recommendations are investment recommendations based on the preferences and characteristics of the investor. For instance, Ping An Insurance Group' Lufax is a wealth management platform using Artificial Intelligence to understand the customer preferences in terms of risk, return, cash flow and suggest him with a tailored portfolio.[72]

Content Recommendation

Content Recommendation includes the 30,4% of the initiatives in the Recommendation Class of Solutions. This Specification refers to the use of Recommendation Systems to suggest a user contents aligned with his interest. Contents could be of every typology, such as films for a movie platform or media contents for social networks. Figure 57 shows the distribution of possible Applications.

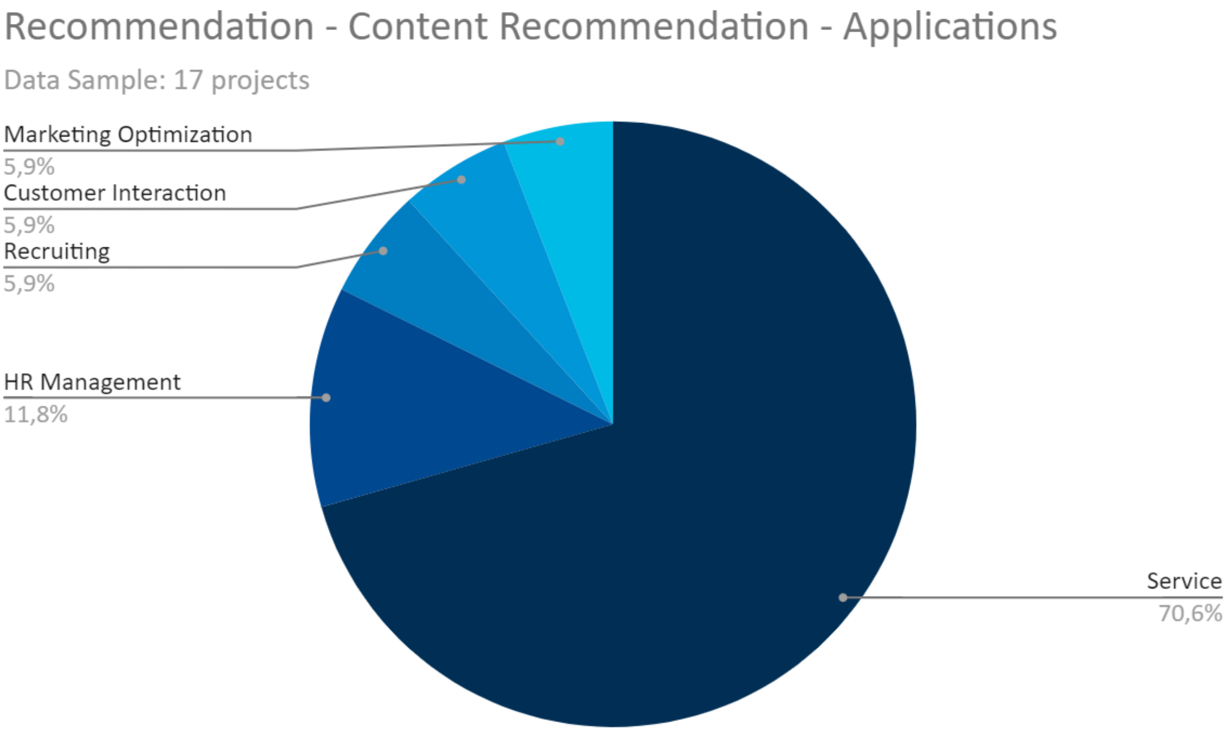

Figure 57: Distribution of Content Recommendation Applications

As shown, Content Recommendations are typically applied to the Services offered by companies (70,6%), to suggest more tailored contents to the users of these Services.
The industry that can better exploit these solutions is the Media & Entertainment one, with more than half of Content Recommendation projects.
For instance, a great example of Content Recommendation in the Media & Entertainment sector is provided by Google: YouTube uses an Artificial Intelligence based algorithm to study viewers' habits and preferences as they stream contents, and work out what would keep them tuned in. Similarly, Google Play Store is based on personalised recommendations for customers, using Artificial intelligence techniques to recommend apps that users will more likely download and enjoy. [73]
Of course, Content Recommendation Systems are not limited to the Media & Entertainment industry: for instance, CVS Health, an American Healthcare & Assistance company, has launched an app for its clients using Artificial Intelligence to make personalized therapy exercises' recommendations for users experiencing anxiety and stress.[74]

To conclude, Content Recommendation solutions are finding adoption also in Applications such as HR Management and Recruiting. For instance, IBM uses an internal platform to offer to its employee jobs currently available within the company. An Artificial Intelligence algorithm, using data like employee's location, pay grade, job role, experience and other factors, is used to recommend jobs tailored for the specific employee. Using this solution, 1500 employees have shifted to new jobs internally.[75]
Similar solutions can be found in Recruiting, where Artificial Intelligence is used to provide the best job matches for individual applicants, based on their profile, aspirations, skills.

Online Advertising

Online Advertising includes the 8,9% of projects in the Recommendation Class of Solutions, referring to the broader variety of ways through which a company can use the Internet to provide promotional marketing messages to customers.

Since Online Advertising is closely related to Marketing Optimization, all the projects are associated with this Application.

For instance, Facebook uses an Artificial Intelligence Recommendation System to decide which adverts to show and to which user on the company's platform. The Recommendation System algorithm uses a collaborative methodology, so that it can show the same adverts to users with similar interest.[76]

In another project, LVMH Moët Hennessy Louis Vuitton uses Artificial Intelligence to deliver tailored and personalized emails to each customer based on his shopping history.[77]

Dynamic Pricing

The minority of the projects in the Recommendation Class of Solution regards Dynamic Pricing, with the 5,4% of initiatives. This Specification considers the internal use of Recommendation Systems to suggest the selling price for product or services based on market demand.

Dynamic Pricing is strongly linked to the Price Recommendation Application; this is the reason why all projects are associated with this Application.

For Instance, Walt Disney uses Machine Learning for the optimization of ticket pricing: Disney has developed a dynamic pricing model for show tickets, predicting both demand and the maximum price acceptable by customers for each ticket.[78]

3.2.7 Autonomous Robot

Autonomous Robots includes only the 4% of analyzed projects, highlighting a still limited diffusion for this Class of Solutions. This can be attributed to the low maturity level of this Class of Solutions: with the 86,1% of projects classified as Pilots or Project Proposals, Autonomous Robot is, along with Autonomous Vehicle, the Class of Solutions with the lowest maturity. As a consequence, further advancements in Artificial Intelligence and Robotics are needed before reaching a wider diffusion of Autonomous Robots.

However, the opportunities offered by Autonomous Robots are tremendous and they can find application in a multitude of industries, as shown in Figure 58.

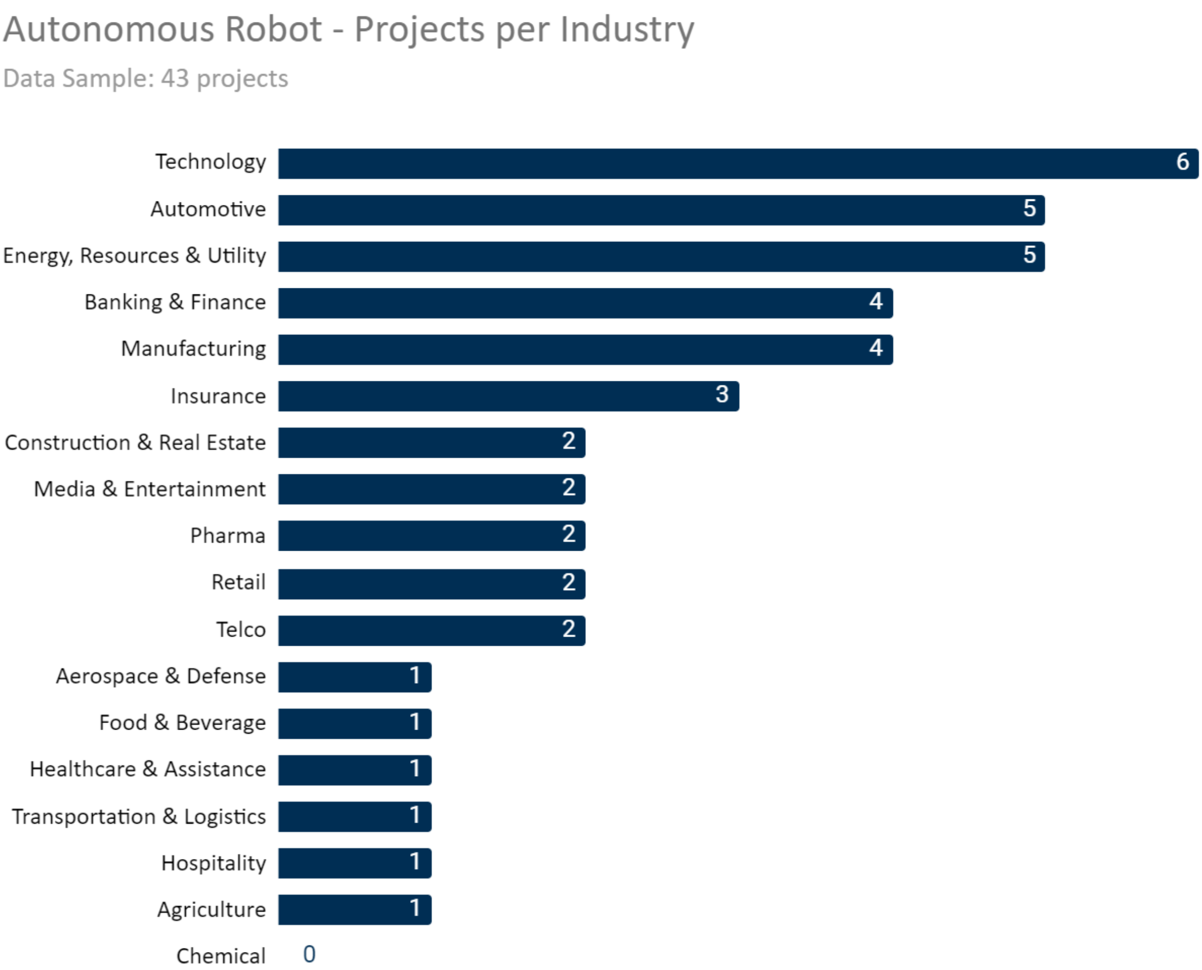

Figure 58: Distribution of Autonomous Robot Projects per Industry

As regards the possible Applications, the main results are shown in Figure 59.

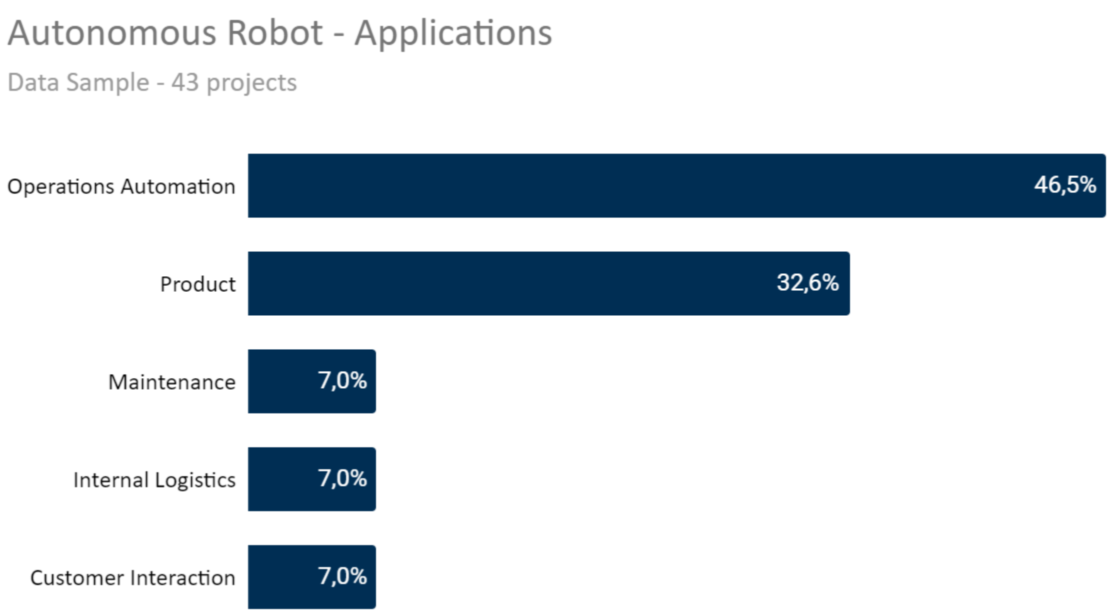

Figure 59: Distribution of Autonomous Robot Applications

As shown, the most diffused Application is Operations Automation (46,5% of Autonomous Robot projects), underlying the relevance of these solutions for the automation of internal business operations. As it can be expected, 62,8% of the projects have been so classified as B2B Internal, suggesting that Autonomous Robot solutions are conventionally internally utilized by companies. Maintenance (7%) and Internal Logistics (7%) are two other possible Applications in case of internal use.
On the other side, Autonomous Robots sold as Products to a B2B or B2C market represent the 32,6% of the solutions. Since robots sold to a B2B market are basically the same solutions internally used by companies, just sold to other companies, the analysis of this Class of Solutions will mainly focus on possible internal uses of Autonomous Robots, providing at the end a dedicated focus on solutions sold to a B2C market.

Focusing on Operations Automation, it includes projects using Autonomous Robot with Artificial Intelligence capabilities to automate operations, autonomously completing entire processes, substituting the workers only in small parts of processes, rather than collaborating with them in production. The Automotive and Manufacturing sectors are, above all the sectors, showing particular interest in Autonomous Robots for Operations Automation.
For instance, a Pilot project of Volkswagen uses a self-learning robot arm to perform different activities in its production plant, such as mounting loudspeakers in doors, applying pads or installing power window regulators. Machine Learning is used to improve its reliability, since the more movements are shown to it, the more its performances improve. The final objective is to create a manufacturing site in which employees have only to show robots what they have to do, and then can focus on more gratifying tasks. Another project of Volkswagen regards Mirco, an Autonomous Co-Worker robot intended to work in collaborative hybrid teams of human workers and robots.[79]
In another project, Mitsubishi Electric is developing a cooperative Autonomous Robots for human-machine collaboration, using Inverse Reinforcement Learning to make robots learn the actions of workers and imitate them. Again, the aim is to create a mixed work environment in production sites, to speed up operations and relieve employees from heavy tasks.[80]

Another relevant internal Application for Autonomous Robots is Internal Logistics. For instance, in a Pilot project in Germany, BMW Groups uses four different types of material handling robots powered by Artificial Intelligence to better handle logistics. Among them, the SplitBots are stationary robots using Artificial Intelligence to detect and take full plastic boxes from pallets in the incoming goods area and place them on a conveyor. Another material handling robot is the PickBot, used to collect different small pieces from supply racks and using Artificial Intelligence to understand the right grip point.[81]

Apart from these two relevant Applications, Autonomous Robots are starting to generate a certain interest also in the Banking & Finance and Energy, Resources & Utility sectors.
In Banking & Finance, some Pilot Projects are introducing humanoid Autonomous Robots in banks, to welcome guests and provide them relevant information. For instance, in a Pilot Project of HSBC Holding, they have introduced in some of banks' floors "Pepper", a humanoid robot able to read emotions and cognitive states using Artificial Intelligence. Its main duties are welcoming clients and teaching them how to open accounts, cracking jokes, relaying credit card details and more.[82]
On the other side, in the Energy, Resources & Utility industry an application for Autonomous Robots' is emerging in Maintenance. The aim of these solutions is not only to reduce inspection costs, but also to improve the safety of employees operating in risky environments such as offshore sites for the extraction of oil.
For instance, Total launched a Pilot project in collaboration with the Technische Universitaet to use an Autonomous Robot able to work in offshore sites, to perform visual inspections, read dials, level gauges and valve positions.[83]

To conclude, in recent years the interest for using Autonomous Robots in a domestic environment has increased. These solutions are used for a variety of tasks and are the solutions previously introduced as Products addressed to a B2C market.
For instance, Sony and Carnegie Mellon University are collaborating in a joint research project at the forefront of Artificial Intelligence and Robotics for the development of kitchen robots able to support humans by autonomously performing tasks such as cooking and serving. [84]

In another project, Samsung is showcasing home care robots, service robots leveraging complex Artificial Intelligence capabilities to help care for family members, especially kids and the elderly, and perform valuable tasks in a domestic environment.[85]

In conclusion, Autonomous Robots have a still limited diffusion.
Further advancements and developments are needed to reach a higher diffusion for these solutions. At the same time, the emergence and adoption of Autonomous Robots is raising concerns about their reliability and their implications on job, with human workers afraid of losing their job in favour of machines.

3.2.8 Autonomous Vehicle

Autonomous Vehicle projects represent the 5,7% of the analyzed Artificial Intelligence initiatives, with 62 projects out of 1089. This Class of Solutions, alongside Autonomous Robot, is the one with the lowest level of maturity: only the 20,8% of the projects are Operative and in Implementation, while the 79,2% of them are still Pilots and Project Proposals. The low maturity of this class is mainly due to the complex Artificial Intelligence capabilities required by Autonomous Vehicles; as a consequence, further advancements are required to achieve a higher maturity and diffusion for these solutions.

By analyzing which are the industries more interested in this Class of Solutions, results are shown in Figure 60. As it could be expected, the Automotive industry is the one with the highest number of projects; however, also other sectors are showing a certain interest in Autonomous Vehicles: above all, Energy, Resources & Utility, Manufacturing, Transportation & Logistics.
This interest can be explained by looking at Figure 61: although the majority of Autonomous Vehicle projects regard Products (33,9%), with a multitude of self-driving cars for a B2C market under development, and Services (24,2%), with ADAS as additional on-board services, internal applications (B2B Internal) for Internal Logistics (14,1%) and External Logistics (21%) are emerging. This is confirmed by considering the destination of the projects: the 41,9% of the solutions are B2B Internal, underlying how Autonomous Vehicles can also play a role within a company's boundaries. More details about this kind of solutions will be later provided.

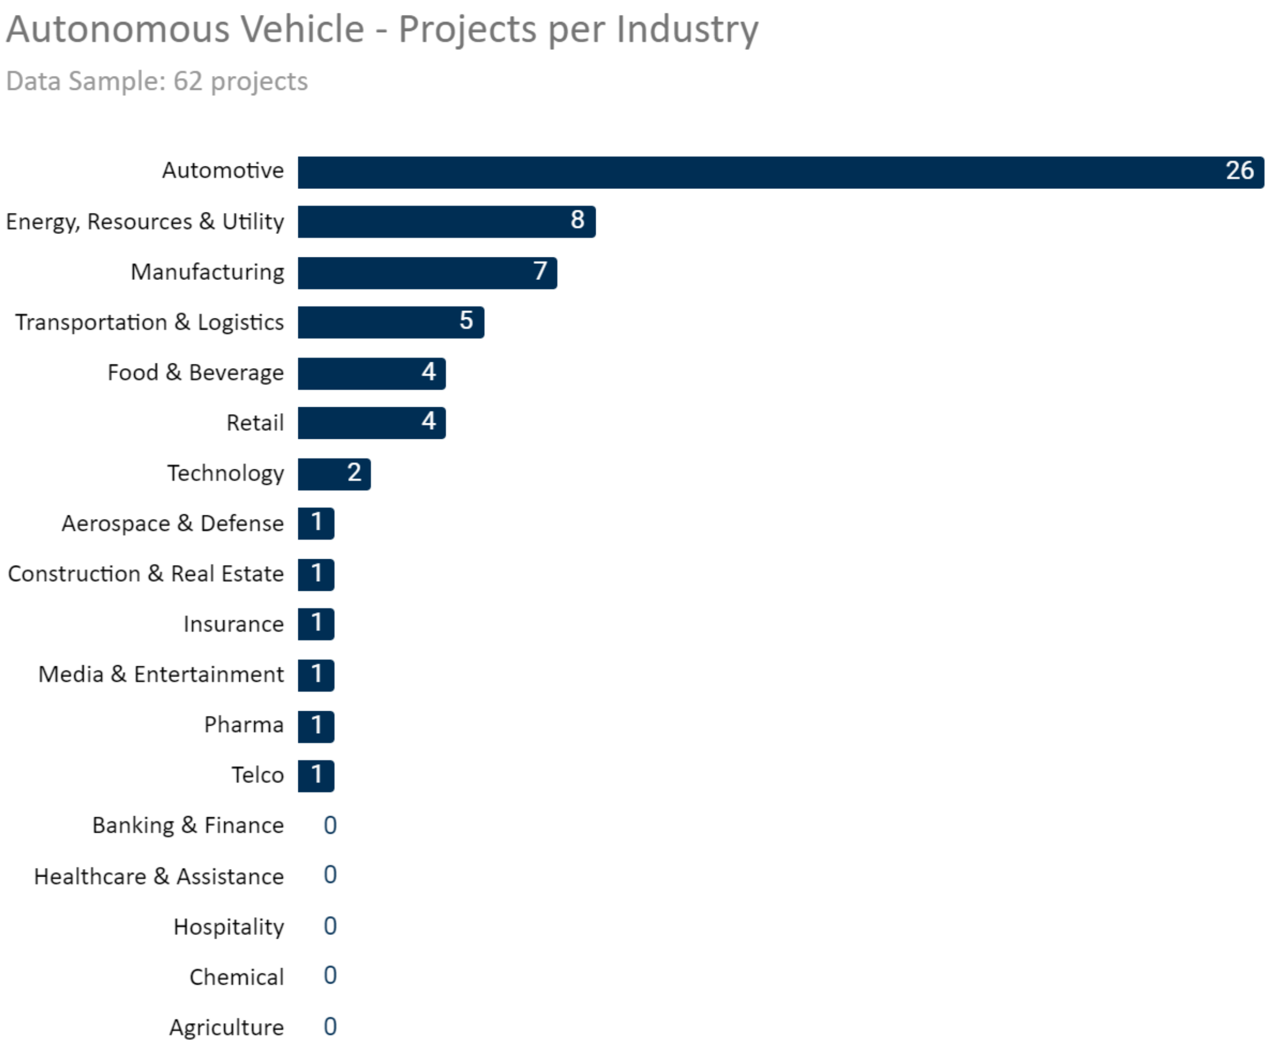

Figure 60: Distribution of Autonomous Vehicle Projects per Industry

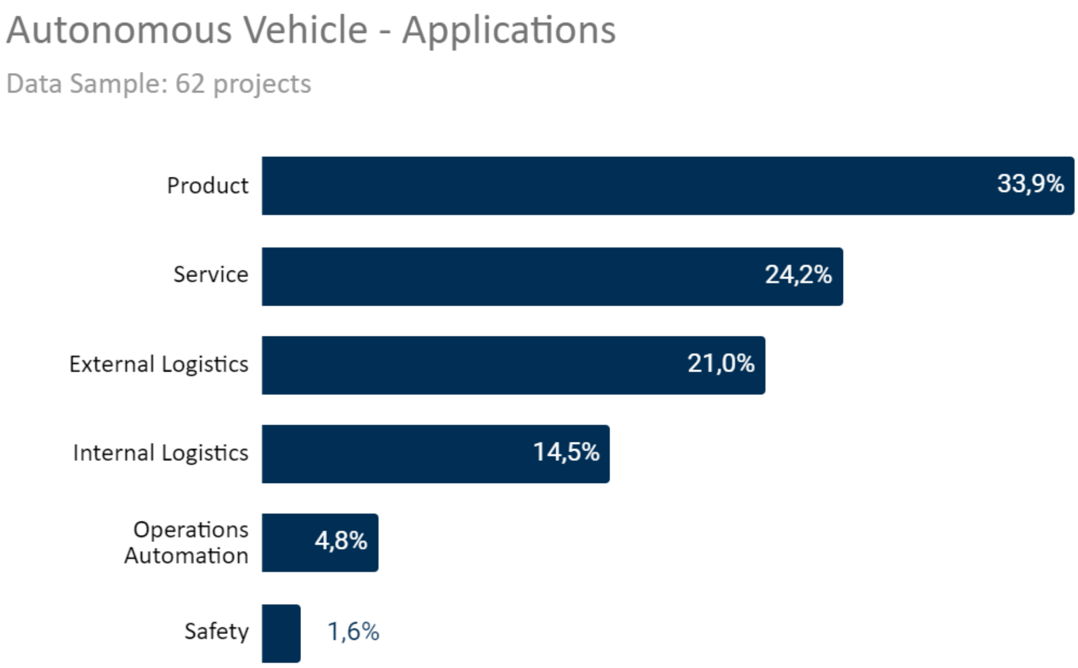

Figure 61: Distribution of Autonomous Vehicle Applications

When analyzing how the projects are distributed in Specifications (Figure 62), Autonomous Systems along a not defined path represent the majority of initiatives (43,5%), mainly because of the several projects of Automotive companies to develop the next generation of autonomous cars. Advanced Driving Assistance Systems follow, with the 29% of the initiatives. To conclude, Autonomous Systems along a defined path represent the 27,4% of the solutions, with an emerging number of solutions applied in Internal Logistics.

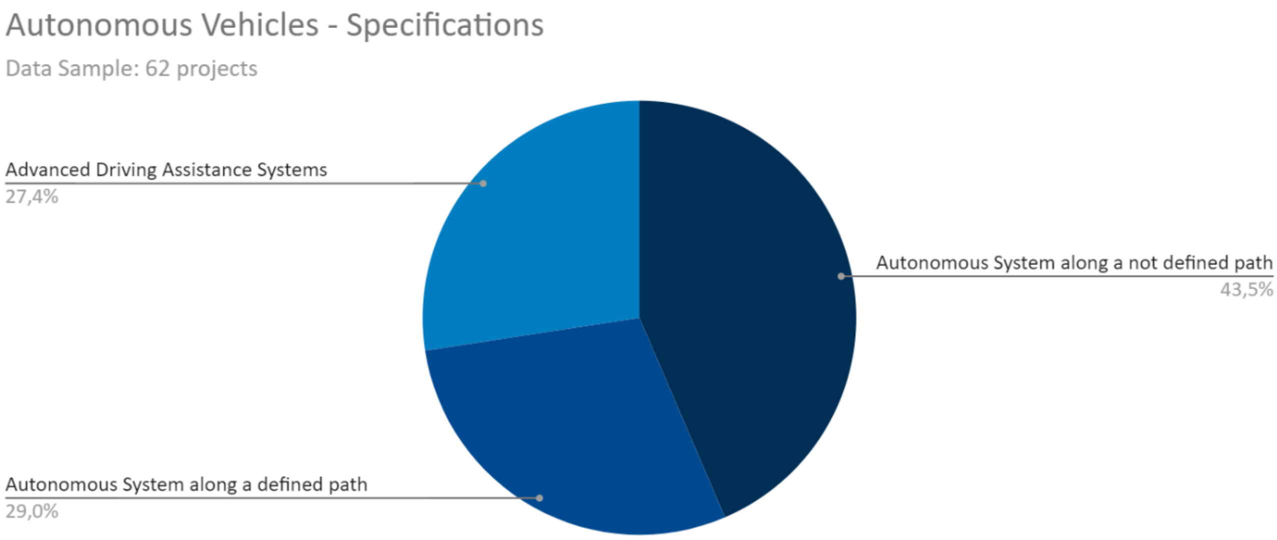

Figure 62: Distribution of Autonomous Vehicle Specifications

Autonomous Systems along a not defined path

As previously mentioned, this is Specifications with the highest number of projects in the Autonomous Vehicles Class of Solution. However, it also has the lowest level of maturity, with just 3,7% of the solutions Operative and the remaining 96,3% including Project Proposals (33,3%) and Pilots (69,3%).

The main Applications are shown in Figure 63.

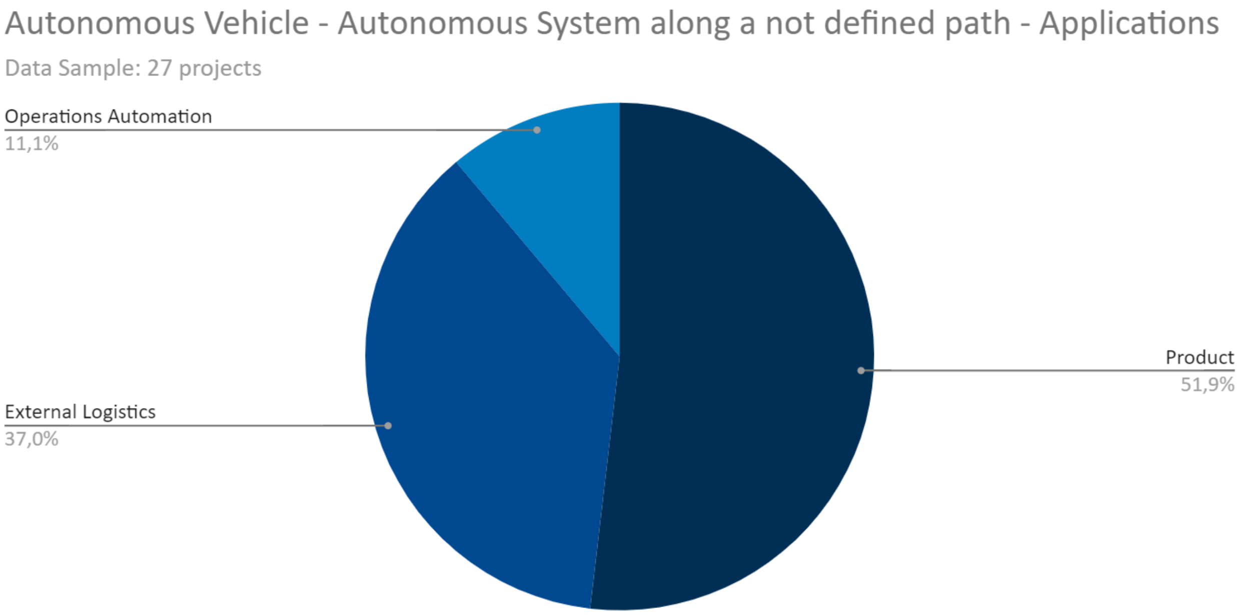

Figure 63: Distribution of Autonomous System along a not defined path Applications

As shown, the most diffused Application is Product (51,9% of the projects): this mainly refer to self-driving cars that companies operating in the Automotive sector are developing for a B2C market, aiming to increase road safety and reduce accidents. BMW, Fiat Chrysler, Honda, Hyundai, Mitsubishi, Toyota, Volkswagen and Volvo, among the others, are all doing research in this field, trying to revolutionize mobility through the large scale deployment of Autonomous Vehicles. The focus of these projects is full autonomy, leveraging Artificial Intelligence to achieve Level 4 and 5 of autonomous driving. Anyway, companies are still developing and testing these solutions. Furthermore, these projects typically take the form of joint research, in which other automotive companies and external Technological Partners, mainly hardware providers such as Intel and Nvidia, and Startups, are involved.
For instance, the automobile manufacturer Fiat Chrysler is part of a joint self-driving research and evaluation group formed by Fiat Chrysler itself, BMW, Intel, and Intel's subsidiary Mobileye, aiming to develop Level 4 Autonomous Vehicles. At the same time, it has several other self-driving partnerships underway, such as the one with the startup Voyage for the development of a self-driving software.[86]

The other relevant Application field for this Specification is External Logistics (37% of the projects). In this case, Autonomous Vehicles are not sold, but internally used by enterprises. In External Logistics, self-driving vehicles, typically trucks, are used to move goods between factories, warehouses, distribution centers and customers, automating the transportation of goods and packages delivery. Especially shipping companies operating in the Transportation & Logistics sector are interested in these solutions, aiming to achieve Level 4 and 5 of autonomous trucking. For instance, the world largest package delivery company United Parcel Service is collaborating with autonomous driving startup TuSimple for the development of level 4 autonomous driving, testing self-driving trucks to improve service and efficiency in its network. At the same time, they are partnering with Alphabet's subsidiary Waymo to test package deliveries through driverless minivans. [87]

To conclude, Operations Automation (11,1%) refers to the use of Autonomous Underwater Vehicles, or AUVs, endowed with Artificial Intelligence capabilities and exploited for ocean exploration and data acquisition across the ocean floor. For instance, the oil and gas international company Royal Dutch Shell is piloting the use of FlatFish, an AUV to acquire a 3D model of underwater structures. Artificial Intelligence is extensively used in its guidance system: the Flatfish follows a pipeline to reach the target area, but then visual homing algorithms are used to return back at the starting point of the mission.[88]

Autonomous Systems along a defined path

Autonomous Systems along a defined path represent the 29% of Autonomous Vehicles projects. Figure 64 shows the typical Applications for this Specification.

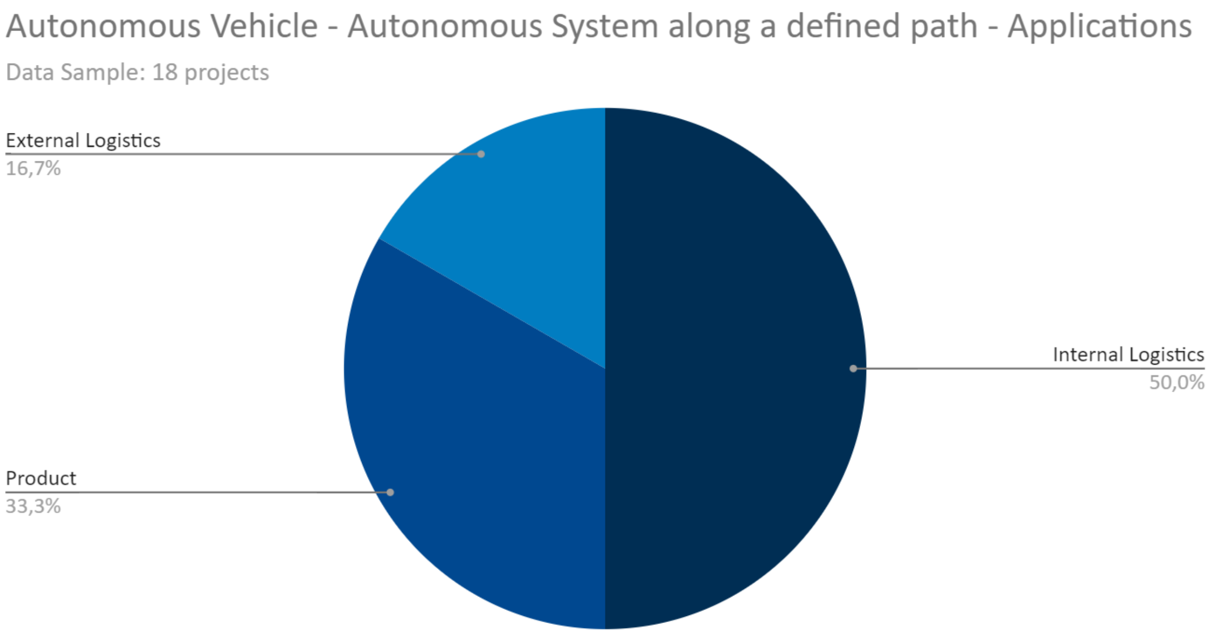

Figure 64: Distribution of Autonomous System along a defined path Applications

When considering Autonomous Systems along a defined path, the most diffused Application is Internal Logistics, with the 50% of the projects. Again, the application is internal to the company but, differently from the previous Specification in which companies from Automotive and Transportation and Logistics sectors were dominant, these solutions are applied in a variety of industries. In this scenario, Autonomous Vehicles are used into a predetermined area within the boundaries of the company, such as a warehouse, to automate the movement of goods.

Here Autonomous Vehicles consist of automated material handling systems, such as forklifts and pallet drones, or trucks, operating in predetermined areas, such as warehouses and production and distribution centers for the transportation of goods.

Artificial Intelligence allows these systems to cooperate to autonomously move with the others in the workplace, recognizing the environment and obstacles such as forklift trucks or humans and navigating through optimal routes. Benefits are not only in safety, but also in productivity and efficiency, with reductions of waste and energy leading to significant operating cost savings.

For instance, Amazon's fulfilment center in Kent, Washington, uses a fleet of autonomous vehicles powered by Artificial Intelligence, each one with nine rows of product-packed shelves. When someone in the Pacific Northwest purchases something on Amazon.com, one vehicle enters in action and is able to autonomously move around the others to reach a station at the edge of a fenced-off robotic field, where a worker picks the item in question.[89]

On the other side, Products (33,3%) refer to the same solutions for Internal Logistics, with companies developing handling systems based on Artificial Intelligence and providing them to other companies. To conclude, External Logistics (16,7%) refers to the use of autonomous vehicles for delivery, typically of food and drinks, inside predetermined areas, such as university campuses, parks, and office buildings, using Artificial Intelligence to manoeuvre around visitors.

Advanced Driving Assistance System

To conclude, Advanced Driving Assistance Systems represent the 27,4% of the analyzed initiatives. Figure 65 resumes the Applications for this Specification.

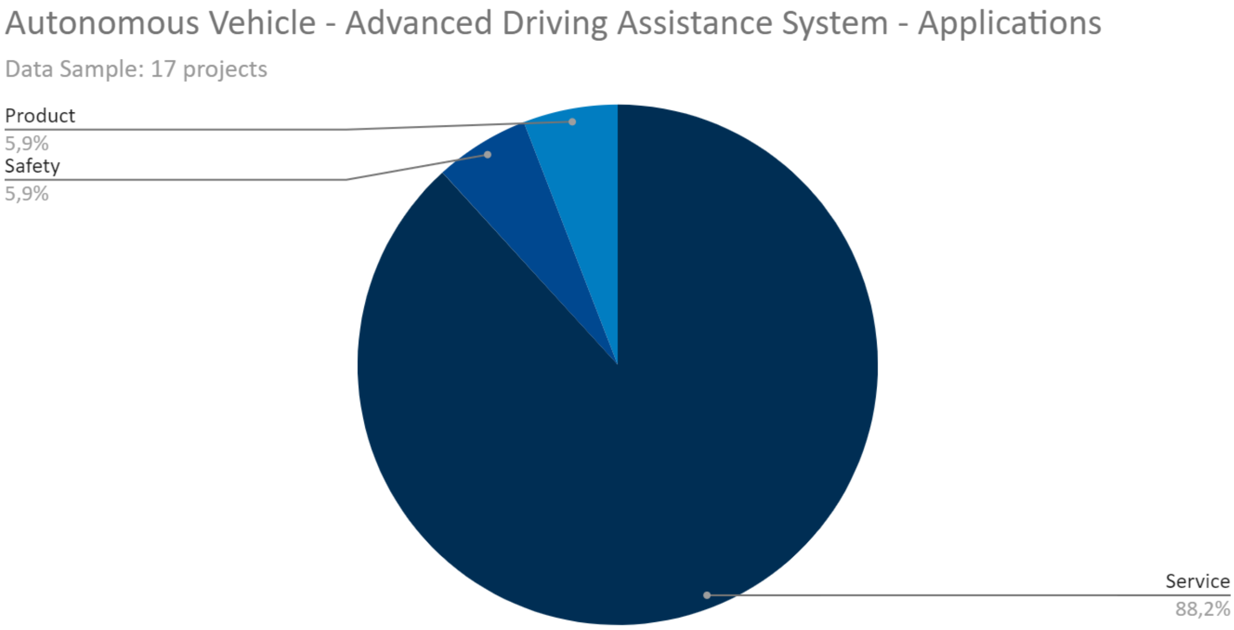

Figure 65: Distribution of Advanced Driving Assistance System Applications

As shown, for this Specification Services represent the most diffused Application, with the 88,2% of the initiatives.

The typical scenario for this Application regards companies operating in the Automotive sector, developing Advanced Driving Assistance Systems integrated as Services in cars for a B2C market, or trucks for the Transportation & Logistics industry.

These systems have not the aim to replace the driver but are intended to amplify the human performance behind the wheel, detecting problems and providing a variety of interim risk management solutions.
Again, the expected benefits mainly regard safety, helping in the prevention of accidents. Some of these solutions have already found widespread diffusion (23,5% of the projects are already Operative), and for this reason they may have had a lower visibility in secondary sources. For instance, Level 2 and Level 3 intelligent driving functions such as Lane Keeping Assist, Pedestrian Detection and Smart Parking Assist have been integrated into several car models.
Other solutions are just Project Proposal (17,7%) or solutions that companies are starting to pilot in one model (52,9%) or implement in all the new produced vehicles (5,9%).
For instance, many companies are working on development of technology to identify hazardous driving behaviours of the driver, such as drowsy driving, smoking, looking at phones, or to check his/her physical state by learning from facial expressions, and accordingly take actions.
For instance, Toyota Motor has two main projects under development. One is a fully autonomous self-driving car, the other one is the Guardian, a driver-assist system to support the driver. It is expected to monitor the environment around the vehicle, alerting in case of potential hazards, and to take control to avoid accidents when needed. It also monitors the driver's behaviour to recognize signs of drowsiness or distraction.[90]

3.2.9 Intelligent Object

Intelligent Object is the Class of Solutions with the lowest adoption rate among companies: only 14 projects out of 1089, representing the 1,3% of the analyzed initiatives, are included in this class. This suggests that Artificial Intelligence algorithms rarely run directly on the devices' hardware, because of limitations in computing, memory and storage capacity. As a consequence, data are typically still elaborated on dedicated cloud platforms and the number of projects in this Class of Solutions is still limited.

The few projects analyzed are distributed in different industries, vary in terms of Status of the Projects, and even when considering the destination of these solutions, they are similarly distributed between B2B Internal, B2C and B2B. Also, analysis of the

Applications does not provide relevant information: the 71,5% of the projects refers to Products sold in a B2C or B2B market, but this is predictable since the class in question is the one of Intelligent Objects.
As a consequence, it is not possible yet to point out an emerging application scenario for this Class of Solutions, but discussion is limited to the presentation of individual use cases. Figure 66 shows the distribution of Specifications within the Intelligent Object class.

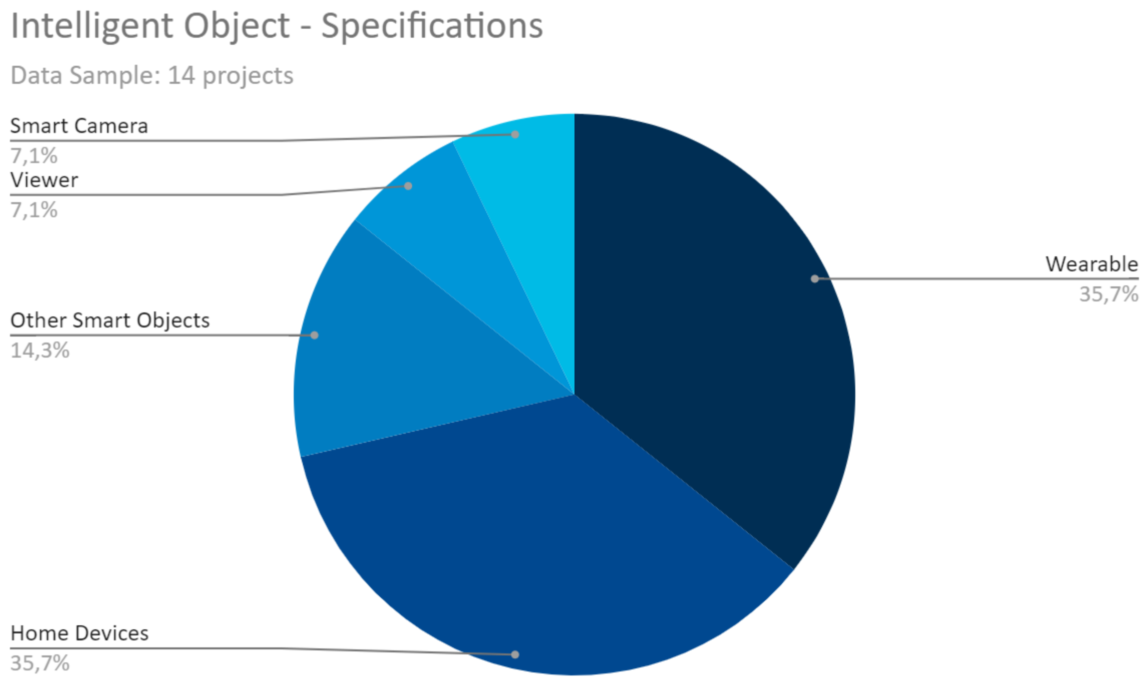

Figure 66: Distribution of Intelligent Object Specifications

As shown, Home Devices and Wearables lead, both with the 35,7% of projects in this class. Home Devices projects include microwave ovens, televisions, and air conditioners, suggesting the role that Intelligent Objects can play in the Smart Home context. On the other side, Wearables include solutions such as smart glasses and smart caps that businesses can internally adopt to support a variety of tasks. Home Devices and Wearables are followed by Other Smart Objects, with the 14,3% of the projects, Viewer (7,1%) and Smart Cameras (7,1%).

Introducing a few use cases, an example of internal adoption is the one of the world-leading resources company BHP Billiton, using Artificial Intelligence and Intelligent Objects to improve mining safety in the monitoring of drivers. Since drowsiness of drivers is one of the main causes of mining accidents, BHP has launched in 2017 a pilot project for using special smart helmets in one of its mines in Chile. These ones analyze the driver's brainwaves, captured through sensors, and alert them if they are falling asleep. The technology has been so successful that it has been implemented at scale.[91]

Another project regards DeepGaze, an Artificial Intelligence camera that Tencent provides to retail businesses. Using Intel's Movidius Myriad technology, Deep Neural Networks are able to run directly on the camera, performing object detection. In this way, information about the number of shoppers in the store along the day and their movements can be gained.[92]

Lastly, in a Smart Home scenario, Midea Group, a Chinese manufacturer of home appliances, has previewed in 2019 a microwave oven able to understand and respond to spoken commands. Voice requests are not sent to the cloud, never leaving the device: Artificial Intelligence algorithms completely run on the microwave oven processors, doing speech processing at the edge.[93]

To conclude, because of the challenges and barriers described in Chapter 1.5.9, the number of projects in this Class of Solutions is still limited and it is still too early to recognize dominant directions. Anyway, it is expected that once challenges are overcome, the amount of applications with Deep Learning algorithms directly hosted on the device will rise.

3.3 Critical Issues, Challenges and Concerns

To conclude, the aim of this Section is to provide an overview about Critical Issues encountered in Artificial Intelligence projects, relevant challenges to overcome to foster the spread of Artificial Intelligence applications, and raising concerns that could limit a future wider adoption of this technology. These Critical Issues have been collected during the creation of the database, and later categorized and organized into the following list.

a. Huge Amount of Training Data

One critical issue in Artificial Intelligence projects is that Machine Learning typically requires a huge amount of operational training data, to learn and imitate human tasks. This implies time and cost burdens in projects; furthermore, not always there is a sufficient availability of training data. Consequently, new techniques are emerging to reduce the amount of data required or to increase their availability.

b. Quality of Training Data

For application purposes, it is required not only to have huge amounts of data to feed the algorithm, but data should be reliable too. Inaccurate data, biased data, or data based on subjective evaluations, may have detrimental consequences on the accuracy of the algorithm and lead to poor performances and reliability.

c. Artificial Intelligence bias

Another drawback of Artificial Intelligence is that, if not correctly done, it can amplify problems of bias and discrimination. Artificial Intelligence bias is a situation occurring when an algorithm generates outcomes that are continuously prejudiced. This can happen because of incorrect assumptions of the model: since algorithms learn to take decisions based on training data, if they contain biases or reflect different types of discrimination, the assumptions of the model will be biased as well.

For instance, the American Health Insurer UnitedHealth Group uses an Artificial Intelligence algorithm to identify the most vulnerable patients and assign them additional support, such as automatic prescription refills or home visits by nurses. The algorithm was under investigation when a study found that black patients with the same risk level of white ones were actually sicker. According to the study, the algorithm reduced the number of black patients eligible for additional support more than a half.[94]

d. The Black Box problem

One of the downsides of using Deep Neural Networks is that the algorithms work as a “black box”: you have visibility into the inputs and the outputs of the model, but not onto the intermediate layers. Therefore, it is not well understood what the algorithm is doing and the logic that it follows to reach a specific decision, and so how it actually works. For this reason, biases and mistakes are more difficult to detect, and it is difficult to understand and explain how decisions are made.

e. Reliability of Solutions

In a small fraction of the analyzed cases, further adoption of Artificial Intelligence may be limited by the reliability of the solutions. Some solutions were found to be ineffective and to fail in the execution of their task, with frequent mistakes. Consequently, in a number of cases further advancements are needed to have a wider diffusion of the technology. For instance, Autonomous Vehicles are still inadequate in performances, and need further advancements to become a reliable solution.

In one pilot project, HSBC, one of the world's largest banking and financial services organisations, introduced in some of its banks Pepper, an Autonomous Robot with hosting duties, such as teaching clients how to open an account and relaying credit card details. Anyway, its performances were found to be still limited, answering bizarre responses when confused or transferring complicated questions to humans.[95]

f. Unfair Use of Artificial Intelligence

One common concern regards the unfair use of Artificial Intelligence solutions, or data being held on customers, by commercial companies, government or other actors. For instance, data about customers could be exploited by governments to track the population. In Image & Video Editing, face swapping technology could be used to create credible fakes of something that did not really happen. Similarly, Samsung's Neon, a human-like video Virtual Assistant, is raising questions about its possible use for deep fakes, or to manipulate videos to show people doing or saying something they did not.[96] Or again, experts fear that solutions to ensure that patients are not skipping their medications could be used by insurance companies to cancel policies, or to avoid paying patients if they take prescriptions incorrectly or miss to do so.

g. Ethical Concerns

Some Artificial Intelligence solutions are raising ethical concerns too. Above all, Autonomous Vehicles are raising ethical considerations, such as who should the vehicle save between passengers and pedestrians in case of unavoidable road accidents. Similarly, applications in a military context are raising ethical concerns about connecting an artificial intelligence to weapons systems.

h. Privacy Concerns

   In some projects, consumers are raising privacy concerns about data being held on them, worried by the possibility of breaches in data and hacking, as well as unfair uses of data.

i. Implications on Job

   Artificial Intelligence is also raising concerns about job displacement and risk of mass unemployment. Autonomous Robots are often considered by the human workforce a major threat, worried by the idea of machines taking over their job. The typical answer to these concerns is that a shift of mindset is necessary: Artificial Intelligence is not there to replace them but to help them, making their job easier, allowing them to focus on more relevant tasks and, even if some jobs may disappear, some new job opportunities will be created.

j. Necessity of Regulation

   Some Artificial Intelligence solutions, such as Autonomous Vehicles, may require ad hoc regulation by lawmakers to be adopted at a large scale. Or, for instance, in China robo-advisory for investments require regulatory approval to be launched, proving that the algorithms create rational investment strategies and remind investors of the risk associated with their use.

To conclude, this Chapter 3 was intended to analyze the diffusion of Artificial Intelligence enterprise applications at a global level, starting at a high-level and then focusing on each Class of Solutions. The discussed results point out how Artificial Intelligence is finding more and more applications in a business context, with an impressive number of projects launched in the last few years. Particularly, the benefits offered by these solutions are so evident that no company can ignore this and stay out of the ongoing Artificial Intelligence revolution. This becomes even more relevant by considering that, in some fields, we have just started to explore the opportunities offered by this technology.

Meanwhile, some concerns about Artificial Intelligence adoption are rising, and a further diffusion of this technology at a global level requires both overcoming a series of practical challenges and further developments.

4. Industry Adoption of Artificial Intelligence

The objective of this last Chapter 4 is to provide an overview about how Artificial Intelligence is being adopted in different sectors, analyzing the created database to answer *RQ2: How Artificial Intelligence adoption changes from industry to industry and what is the potential contribution for the sector?*

In Chapter 3, the focus mainly was on the international adoption of Artificial Intelligence and the analysis of each Class of Solutions, how they are used and for which Applications they are adopted. In this Chapter 4, the perspective shifts from the point of view of the Classes of Solutions to the point of view of the Industry: the main objective is to understand how Artificial Intelligence is applied by organizations operating in a specific Industry, and how this technology can help them in creating competitive advantage. This allows to provide an overview about how Artificial Intelligence is being adopted in a certain industry and how it is transforming the business.

Particularly, this sectoral analysis is focused on 4 different industries: first of all the Food & Beverage, an industry particularly interesting since it is the sector with the highest average number of projects per company in the analysis. This suggests a great interest of the industry in Artificial Intelligence solutions and a wide applicability of the technology in the sector. Then the Manufacturing sector, interesting to analyze due to the possible implications of Artificial intelligence on manufacturing plants and production factories, to better understand the role that the technology can play in the digital transformation of the factory. The third industry analyzed is Retail, interesting since it differs from other sectors for the fact that retailers are not creating products, but just purchasing them from wholesalers or manufacturers and selling them to consumers at a mark-up. Therefore, the sector is very focused on the shopping experience, and it is in direct contact with the ultimate consumers. To conclude, the Banking & Finance industry is the one with the highest number of projects in the database and have widely adopted Artificial Intelligence in recent years, with several solutions becoming a standard in the industry.

To support these sectoral analyses, a Qualitative Framework has been specially developed. The aim of this model is to both support the analysis of how Artificial Intelligence is being applied within a certain industry and a structured and orderly presentation of the results.
The framework is not intended as a tool to identify possible gaps in an industry or to suggest where companies in the sector should invest; on the contrary, it intended as a support to analyze and describe the adoption of Artificial Intelligence in a sector and for which Applications the technology is being deployed in the industry.

The developed framework is shown in Figure 67: as shown, it is organized in 3 different Levels, in which the possible Artificial Intelligence Applications have been grouped based on their proximity with the specific customers of the industry. Particularly, the model can be represented with a pyramidal scheme that places on top the customer.

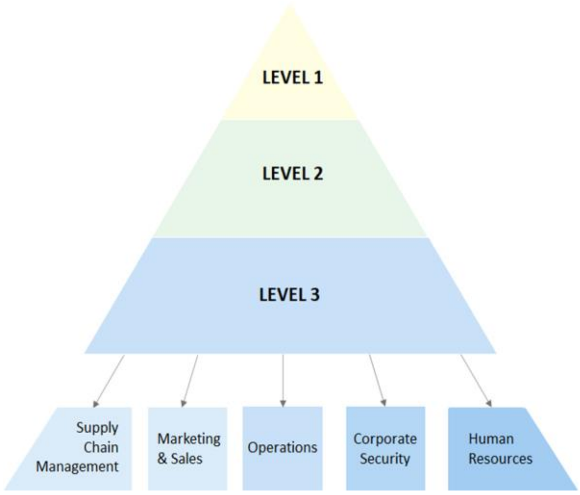

Level 1	AI - Enabled Products & Services and Customer Interaction	Service Product Customer Interaction
Level 2	Marketing, Sales and NPD	Customer Profiling Demand Forecasting Marketing Optimization New Product Development Price Recommendation Sentiment Analysis
Livel 3	Supply Chain Management	External Logistics Internal Logistics Supply Chain Improvement
	Business & Financial Control	Regulatory Compliance Risk Management
	Operations	Employees Support Engineering Asset Management Maintenance Operations Automation Production Planning Quality Control
	Corporate Security	Cybersecurity/Fraud Detection Safety Security
	Human Resources	HR Management Recruiting

Figure 67: Visual Representation of the Qualitative Framework

As regards the three Levels of the framework, each of them includes possible Applications based on the proximity of these solutions to the customer. In particular, they have been defined according to the following criteria.

Level 1 – AI Enabled Products & Services

At this level, Artificial Intelligence is used as an enabler of the Product and Services that companies offer to their clients (B2B, B2C or B2G market), or is at the basis of the interaction between the clients and the company.
Therefore, Applications included in Level 1 are: Service; Product; Customer Interaction.

Level 2 - AI in Customer-Oriented Processes

At this level, Artificial intelligence is used internally by companies. Data elaborated by Artificial Intelligence algorithms are customer data, or the customer has a direct visibility on the outcomes of Artificial Intelligence solutions.

Therefore, Applications included in Level 2 are: Customer Profiling; Demand Forecasting; Marketing Optimization; New Product Development; Price Recommendation; Sentiment Analysis.

Level 3 - AI in Enterprise-Oriented Processes

At this level, Artificial intelligence is still used internally by companies, but generally data elaborated by Artificial Intelligence algorithms are not customer data, or the customer has not a direct visibility on the outcomes of Artificial Intelligence solutions (even if he could indirectly benefit from their implementation).

Since Level 3 includes a variety of Applications, it has been furtherly organized into sub-Levels to enable a more structured analysis, grouping Applications based on their field of application.

Therefore, sub-Levels (and related Applications) at Level 3 are:

a. Supply Chain Management: External Logistics; Internal Logistics; Supply Chain Improvement.
b. Business and Financial Control: Regulatory Compliance; Risk Management.
c. Operations: Employees Support; Engineering Asset Management; Maintenance; Operations Automation; Production Planning; Quality Control.
d. Corporate Security: Cybersecurity/Fraud Detection; Safety; Security.
e. Human Resources: HR Management; Recruiting.

To avoid possible misunderstandings about the classifications of some Applications, the less intuitive ones are briefly clarified. While at Level 2 (AI in Customer-Oriented Processes) Customer Profiling, Sentiment Analysis, Price Recommendation, Demand Forecasting and Marketing Optimization are solutions typically based on customer data or information about the behaviour of the clients, this is not always true for the New

Product Development application. Consider, for instance, Intelligent Data Processing solutions widely used in the Pharma Industry. However, the outcomes of these solutions (new products) are visible to customers since the resulting products are bought from them.

Furthermore, Employee Support at Level 3 has been classified in the Operations sub-Level, and not in the Human Resources one, since this Application included all those solutions supporting employees in the execution of their daily tasks.

The developed framework has then been applied to each of the industries previously mentioned, supporting a structured analysis of how and where companies in a sector are adopting Artificial Intelligence solutions, as well as an orderly presentation of the results. After a short contextualization of the industry, the model is applied to understand the high-level distribution of Artificial Intelligence projects in the different levels and comment on it. Then, the analysis goes into detail to each Level, describing which are the most relevant sub-Levels and Applications and providing, when needed, relevant business cases. The analysis ends with a summary of the analyzed initiatives in each industry, combining the developed Qualitative Framework with the different Classes of Solutions, and summarizing the most important uses in the industry. (Figure 68)

		IDP	NLP	VA/CHATBOT	CV	iRPA	RECOMM.	AR	AV	IO	
AI-Enabled Product & Services		13	1	19	9		5	1			41,7%
AI in Customer-Oriented Processes		3	7				2				10,4%
AI in Enterprise-Oriented Processes	Supply Chain Management	1									47,8%
	Business & Financial Control	7	1								
	Operations	6	10	2		8		3			
	Corporate Security	14			1						
	Human Resources		1	1							
		38,3%	17,4%	19,1%	8,7%	7,0%	6,1%	3,5%	0,0%	0,0%	

Figure 68: Application-Class of Solutions Matrix applied to a generic Industry

In the following sections, an overview about the adoption of Artificial Intelligence solutions in each of the selected industries is provided, applying the developed Qualitative Framework.

4.1 Food & Beverage Industry

The Food & Beverage industry includes all those companies involved in the production (the processing of raw materials into food and beverage products), packaging and distribution of edible goods via several channels. Food products considers both prepared and packaged goods, beverage products both alcoholic and nonalcoholic beverages.
Foods directly produced via farming have not been considered in the Food & Beverage industry, characterized by the processing of raw food materials into more value-added foods and beverages, but have been associated with the Agriculture Industry.
To conclude, in the Food & Beverage industry distribution regards the transportation of finished products to a variety of actors, including retail outlets, restaurants, bars and cafés, hotels, pubs, or direct delivery to end consumers. In addition, the availability of vending machines including food products and beverages in private and public locations must be considered.

The analysis of the industry has been carried out by considering the Food & Beverage companies in the database, including: Anheuser-Busch InBev (AB InBev); Coca Cola; Kraft Heinz Company; Mondelez International; Nestlé; PepsiCo; Unilever. The limited number of companies in the sample, due to the reduced number of Food & Beverage companies in the first 235 positions of the Forbes Global 2000, should not be interpreted negatively. On one side, all these firms are multinational enterprises with an impressive number of global and local brands around the world. For instance, PepsiCo's brands portfolio includes not only Pepsi, but also other 22 brands of snacks and beverages such as Gatorade, Tropicana, Kipton, Frito-Lay and Quaker. Nestlé's brand portfolio covers all the possible food and beverage products, with thousands of global and local brands including Nestlè, S.Pellegrino, Maggi, Buitoni, Nesquik, Nespresso, Kit Kat, Smarties. In the alcoholic beverage market, AbInBev portfolio includes more than 200 beer brands, such as Corona, Stella Artois, Becks. On the other side, the Food & Beverage sector has been the one with the highest average number of projects per company (9,86; total number of projects in the Industry: 69) in the analysis, underlying a great interest of the industry in Artificial Intelligence applications. Therefore, the analysis of the Food & Beverage industry was particularly interesting.

Figure 69 shows how the 58% of the Artificial Intelligence initiatives in the industry are Operative, and the 10,1% in Implementation, pointing out that the adoption of some Artificial Intelligence solutions has been successful, so that they have been implemented at scale. Meanwhile, the number of Pilot projects (24,6%) and Project Proposals (7,2%) suggests that Food & Beverage companies are still exploring the potential of Artificial Intelligence in the industry, experimenting the use of this technology with a variety of new solutions.

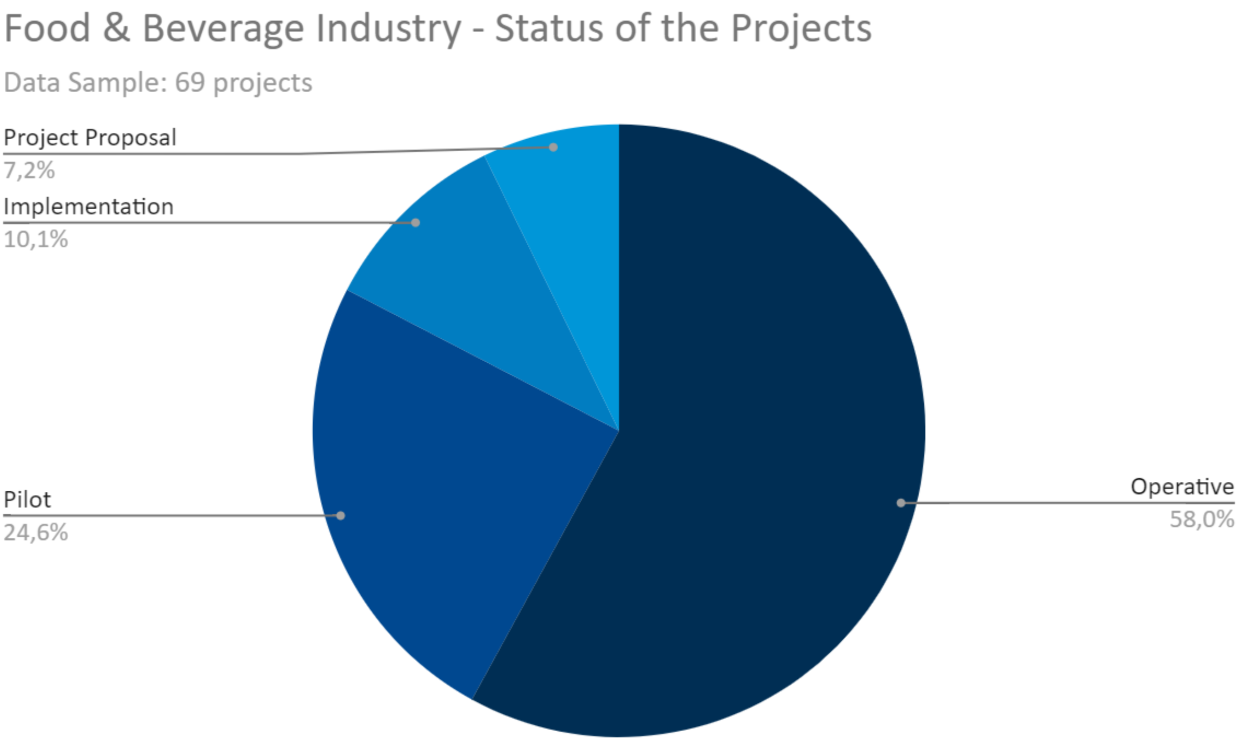

Figure 69: Status of the Projects in the Food & Beverage Industry

Particularly, Food & Beverage companies are leveraging Artificial Intelligence and Machine Learning in a variety of different ways. To better understand where Artificial Intelligence is currently applied in the industry, the Qualitative Framework has been applied. Results are shown in Figure 70.

Level 1	Service	5,8%
	Product	0,0%
	Customer Interaction	8,7%
Level 2	Customer Profiling	10,1%
	Demand Forecasting	4,3%
	Marketing Optimization	4,3%
	New Product Development	0,0%
	Price Recommendation	4,3%
	Sentiment Analysis	4,3%
Level 3	External Logistics	4,3%
	Internal Logistics	2,9%
	Supply Chain Improvement	11,6%
	Regulatory Compliance	1,4%
	Risk Management	0,0%
	Employees Support	0,0%
	Engineering Asset Management	0,0%
	Maintenance	5,8%
	Operations Automation	4,3%
	Production Planning	4,3%
	Quality Control	4,3%
	Cybersecurity/Fraud Detection	0,0%
	Safety	0,0%
	Security	0,0%
	HR Management	5,8%
	Recruiting	13,0%

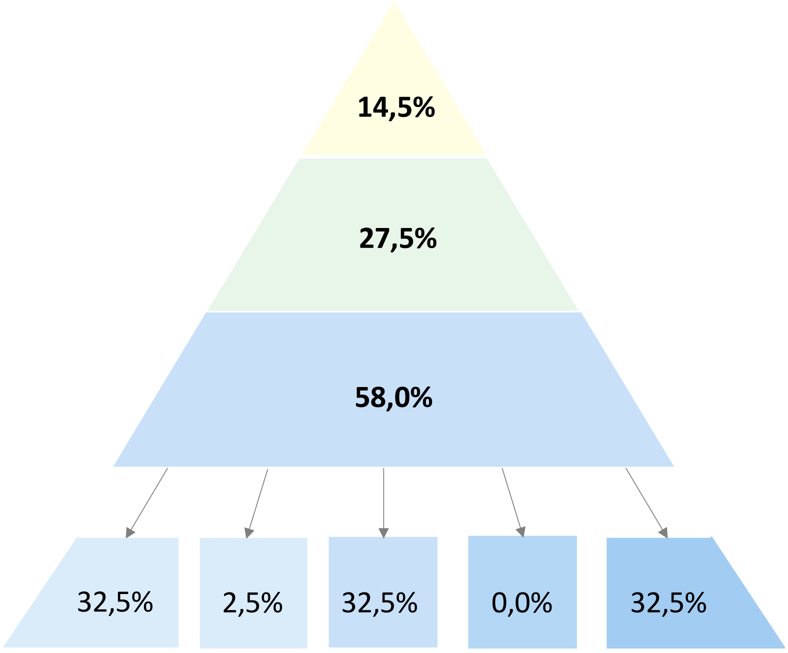

Figure 70: Qualitative Framework applied to the Food & Beverage Industry

As shown, the analyzed Artificial Intelligence initiatives mainly regard Level 3 (AI in Enterprise-Oriented Processes), with the 58,0% of projects in the industry. As a consequence, Artificial Intelligence is finding adoption especially within the boundaries of Food & Beverage enterprises, in supporting the processing and packaging of food and beverage products (Operations: 18,7% of the projects), in Supply Chain Management (18,8% of projects). At this sub-Level, initiatives are mainly related to the improvement of the whole Food & Beverage Supply Chain and the distribution of products. Human Resources is relevant too, with 18,8% of total projects. Artificial Intelligence is finding a good level of adoption in the industry also at Level 2, AI in Customer-Oriented Processes, with 27,5% of the projects. At this Level, Food & Beverage companies are using it for Customer Profiling and Sentiment Analysis, to gather insights about customers' behaviour and extract their sentiment from social media. The obtained results are then used to develop new products or to personalize marketing interactions, reaching customers in the most appropriate way. To conclude, Level 1 (AI Enabled Products & Services) is characterized by a limited adoption of Artificial Intelligence, with just the 14,5% of initiatives. This can be partially explained by the inherent characteristics of the industry, since products sold to customers are food and beverage products and the number of services in this industry is generally limited. However, Artificial Intelligence is being applied also at this level in interactions with the customers (Customer Interaction: 8,7% of projects).

In the following sections, each Level will be briefly discussed, to provide information about the way Food & Beverage companies are adopting Artificial Intelligence.

Level 1 - AI - Enabled Products & Services

At this level, Artificial Intelligence seems to have a limited applicability in the Food & Beverage industry.

This can be explained also by considering that products sold are food & beverage products, so that they cannot be endowed with Artificial Intelligence capabilities. Meanwhile, also the number of offered services in the industry seems to be limited. However, some Services enabled by Artificial Intelligence are emerging: they are mainly directed to end consumers (B2C market), but also solutions for bars and restaurants (B2B market) have been found.

Services for end consumers typically consist of applications or programs to support the clients of the company in planning their meals, providing them tailored recommendations based on their needs, diets, or preferences for taste.

For instance, in Japan Nestlé is piloting the "Wellness Ambassador", a nutrition program at the subscription cost of $600 per year, with already 100.000 participants. Artificial Intelligence is at the basis of the program: users can upload photos of their meals through a dedicated app and Computer Vision is used to recognize foods and understand the users' nutrition habits. Artificial Intelligence allows to analyze these data to provide users with suggestions about lifestyle changes and personalized dietary supplements to take (Nestlé's capsules to make teas, smoothies and snacks fortified with vitamins).[97]

Customer Interaction seems to be the scenario in which Artificial Intelligence can find the highest applicability at Level 1, and mainly refers to the integration of Chatbots and Voicebots into corporate websites, Facebook Messenger, WhatsApp and other channels to handle customer requests and provide them a superior customer experience. In a B2C scenario, they are used to provide consumers information, to suggest new products, to answer queries on nutrition; in a B2B scenario, they are used to provide business customers like restaurants, bars and other vendors information about the status of their orders and to answer FAQs. Alongside these solutions with similar applications also in other sectors, new opportunities related to the way customers interact with the brand, or its products and services, are emerging. For instance, an innovative project of Coca Cola aims to develop a Virtual Assistant to be integrated in vending machines, changing its behaviour based on the specific location and allowing users to customize their choices by mixing flavours[98].

Level 2 - AI in Customer-Oriented Processes

Level 2 includes the 27,5% of the projects in the industry, suggesting that data about customers or their behaviours are often internally used by Food & Beverage companies in Customer-Oriented Processes.

Particularly, at this level Food & Beverage companies are adopting Artificial Intelligence solutions to better know their customers: Intelligent Data Processing and Natural Language Processing are widely used in the industry for Customer Profiling and Sentiment Analysis purposes. These kinds of solutions have found widespread diffusion in the Food & Beverage industry for a few years. In Customer Profiling, Artificial Intelligence scans millions of datasets to highlight dominant patterns, extract insights about buying behaviours, or to define customer segments. On the other side, in Sentiment Analysis Artificial Intelligence analyses social media, conversations and feedback from customers to understand what people are saying about brands and their products, as well as which situations prompt them to talk about them.

The results of Customer Profiling and Sentiment Analysis are then used to bring new food and beverage products to the market (supporting New Product Development) or for better marketing and promotion of products (supporting Marketing Optimization), targeting customers with relevant and personalized communications.

For instance, Unilever is using Artificial Intelligence to analyze data coming from several sources, like social media and CRM, to find insights. In an interesting project, they applied Artificial Intelligence to analyze sources of influence for people, such as dialogues in films and lyrics in songs, to spot emerging trends and behaviours of people. They found that more than 50 songs in the public domain talked about “ice-cream for breakfast”, and that

people were actually starting to consume ice-cream for breakfast. Based on this discovery, in 2017 one of their brands producing ice cream launched a range of breakfast flavours, like Fruit Lorry and Frozen Flakes, and two years later several competitors have started to imitate them[99].
In another project, Coca Cola used Artificial Intelligence to understand customers' needs and preferences in terms of flavours in certain regions. To do so, they analyzed data about the customers' use of self-service soft drink dispensers enabling them to mix different flavours. Results allowed Coca Cola to launch in 2017 a new product, the Cherry Sprite, that flavours Sprite with a cherry taste[100].

Apart from these solutions for Customer Profiling and Sentiment Analysis playing a key role in the industry, Artificial Intelligence is being adopted also for different Applications.
In Marketing Optimization, different solutions are finding adoption, ranging from the use of Natural Language Processing to provide guidance on how to improve marketing messages, to vending machines promoting products and flavours based on their location.
In Demand Forecasting, Artificial Intelligence is used to forecast demand of food and beverage products with better accuracy. In Pricing Recommendation, Artificial Intelligence is used to analyze massive amounts of current and historical data and optimize decisions about prices and promotions, achieving the highest possible rates of customer retention and adoption.

Level 3 - AI in Enterprise-Oriented Processes

To conclude, Level 3 is the one with the highest diffusion of Artificial Intelligence initiatives, underlying the great opportunity for the Food & Beverage industry to use Artificial Intelligence technology to improve internal activities. Three sub-Levels are particularly interested: Supply Chain Management, Operations, Human Resources.

Supply Chain Management

Artificial Intelligence is creating huge opportunities in the Food & Beverage industry for Supply Chain Management.
Particularly, several initiatives have been launched for Supply Chain Improvement, referring to projects in which Food & Beverage industries are collaborating with their Supply Chain Partners at advantage of the entire Supply Chain. Most of these solutions

are already Operative and leverage Intelligent Data Processing and Computer Vision to improve the whole Food & Beverage Supply Chain.
On the Supply Side, partners are farmers and growers providing raw materials to the Food & Beverage industry. Food & Beverage companies are collaborating with them to use Artificial Intelligence for understanding optimal growing conditions. For instance, the world's largest brewer AB InBev is one of the top global buyers of barleys. It is currently using Artificial Intelligence on Microsoft Azure to analyze data about barley crops, and help affiliate farmers in deciding the best moment to plant seeds or the quantity of fertilizer to use; at the same time, it is analyzing images of fields captured by drones, combined with other data about fields, to provide them suggestions to improve growing environments. Therefore, improvements in productivity, quality and efficiency are impacting not only farmers and growers, but AB InBev itself too [101].
On the Demand Side, partners involved in these collaborative projects are mainly retailers. Computer Vision and Intelligent Data Processing are used to ensure products are on the shelves of retail stores or to notify when inventories are running low, so that automatic reorders happen. For instance, Ab InBev is focusing not only on the Supply Side of its Supply Chain but also on the Demand Side, piloting a system for inventory visibility on the store shelves of its retail partner IGA Extra Beck in Montreal. Images of the store shelves are analyzed for out-of-stock detection, notifying the retail partner when needed. [102]

Artificial Intelligence initiatives have been launched for External and Internal Logistics too, even if all the analyzed projects are only Pilot projects and Project Proposals, marking their distance from the solutions for Customer Profiling and Sentiment Analysis that the industry has widely adopted in the recent years.
Initiatives for Externa Logistics mainly regards the use of Autonomous Vehicles. For snack deliveries to end consumers, mobile vending machines and robots using Artificial Intelligence to manoeuvre around people have been piloted in enclosed areas, such as university campuses and parks. For the transportation of products between warehouses of Food & Beverage companies and the ones of retailers, the possible use of autonomous trucks is under development.
When it comes to Internal Logistics, Nestlè has ushered in 2020 a warehouse using Autonomous Vehicles, collaborative Autonomous Robots and advanced sortation

systems to handle goods in the East Midlands, United Kingdom. Particularly, Autonomous Robots are able to take food & beverage products from storage and build mixed pallets, with cases of different dimensions and brands. [103]

Operations

At Level 3, Food & Beverage companies are also particularly interested in the adoption of Artificial Intelligence to support Operations. The main Applications of interest are Maintenance, Production Planning, Quality Control and Operations Automation, with initiatives at different stages of implementation.

In Maintenance, Artificial Intelligence initiatives typically based on Intelligent Data Processing are used to improve the ability of Food & Beverage companies to predict problems in equipment and failures before they occur. For instance, AB InBev and Kraft Heinz are using Artificial Intelligence in their plants to listen to motors in manufacturing lines and understand when they are going out of their baseline noise, or to analyze vibrations of machines to understand when a unit is going to fail. In both cases, value lies in the possibility to carry out repair or replace activities before any break of the system[104]. If shifting the focus to fault diagnosis, Coca Cola is piloting the use of Computer Vision and Augmented Reality solutions. Dedicated apps or smart glasses are used to support equipment's visual problem identification and to help technicians in diagnosis and solution of technical problems through augmented repair. This solution has been tested not only in manufacturing plants, but also for maintenance of vending machines at different locations[105]

In projects for Production Planning, Intelligent Data Processing solutions are applied to automatically adjust production parameters or for production schedule optimization in industrial plants. As a consequence, they can be applied both to optimize processes, such as beer filtration or potato peeling, adjusting specific parameters and optimizing the system to achieve optimal results, rather than to help companies in improving their production scheduling.

In Quality Control, Artificial Intelligence systems are used to control and maintain product quality, a challenge particularly relevant for the sector. Computer Vision is being adopted for visual inspection of products, enabling faster quality checks and higher accuracy of inspections and automation of quality controls. Alongside these intuitive solutions, new original initiatives are emerging: for instance, PepsiCo's Frito-Lay uses Machine Learning in its chips processing system, hitting them with a laser and listening to their sounds to assess their quality [106].

To conclude, when dealing with Operations Automation, the analyzed initiatives are not related to the factory, but they mainly refer to the automation of back-office operations through iRPA solutions, simplifying the execution of routine transactional activities. Only one project regarding the automation of production processes has been found, with Frito Lay using Computer Vision to automatically detect the amount and weight of processed potatoes along the production line. This solution led the company to considerable saving, compared with the previous weighing systems[107].

Human Resources

To conclude, Artificial Intelligence in the Food & Beverage is being applied also for Recruiting and HR Management. Despite this is not the main focus of this industry analysis, relevant solutions are briefly presented.

In Recruiting, different solutions have been adopted, including the use of Intelligent Data Processing to assess the suitability of employees based on results in cognitive and behavioural tests and Computer Vision to analyze body language in interviews. Chatbots and Voicebots are used to autonomously carry out job interviews, or their initial stages, while Natural Language Processing is used to extract keywords, skills and concepts from job applications to return the best job matches, or to filter job applications.

In HR Management, the use of Natural Language Processing regards Speech to Text for autonomous transcription of words in training videos for employees and Language Translation of the same words, Chatbots to help new employees in getting started in their jobs and Contents Recommendation Systems to suggest them new opportunities to increase their skills and experiences.

To conclude, Figure 71 summarizes the analyzed initiatives in the Food & Beverage sector, according to their Class of Solutions and the developed Qualitative Framework.

As shown, almost all the Classes of Solutions are adopted in the Food & Beverage Industry. At Level 1, most of the initiatives refer to the adoption of Chatbots and Virtual Assistants for Customer Interaction purposes, a combination widely spread also in other industries. At Level 2, the number of initiatives increases and mainly refers to the use of Natural Language Processing and Intelligent Data Processing solutions for Sentiment Analysis and Customer Profiling, to support New Product Development and Marketing Optimization. Level 3, the one with the highest number of initiatives, is particularly interesting for the use of Computer Vision and Intelligent Data Processing solutions to improve the Supply Chain, Autonomous Vehicles in External Logistics and solutions based on Intelligent Data Processing and Computer Vision for different Applications in Operations, using Artificial Intelligence to support the digital transformation of the factory.

Food & Beverage Industry		IDP	NLP	VA/CHATBOT	CV	iRPA	RECOMM.	AR	AV	IO	
AI-Enabled Product & Services		1		5	2		2				14,5%
AI in Customer-Oriented Processes		11	5		1		2				27,5%
AI in Enterprise-Oriented Processes	Supply Chain Management	5			3			1	4		58,0%
	Business & Financial Control	1									
	Operations	7			4	2					
	Corporate Security										
	Human Resources	2	6	3	1		1				
		39,1%	15,9%	11,6%	15,9%	2,9%	7,2%	1,4%	5,8%	0,0%	

Figure 71: Application-Class of Solutions Matrix applied to the Food & Beverage Industry

4.2 Manufacturing Industry

The Manufacturing Industry includes all those companies using manual labour or machines to transform raw materials, components or substances into new products. The realized products may be sold to other companies, as well as distributed through wholesalers to final consumers.

For the scope of the analysis, the Automotive Industry has been considered separately from the Manufacturing Industry: despite the adoption of similar Artificial Intelligence solutions in manufacturing plants, the Automotive Industry is highly focused on the development of Autonomous Vehicles and ADAS as additional on board functions. Similarly, the Food & Beverage industry has been separately considered to focus on food and beverages products, while technology firms like Sony and Samsung have been considered in the Technology industry.

Consequently, the companies analyzed in the Manufacturing Industry are 15 (3M; British American Tobacco; Caterpillar; Danaher; Deere & Company; Hyundai Motor; Hon Hai Precision; L'Oreal; LVMH Moët Hennessy Louis Vuitton; Midea Group; Mitsubishi (Electric); Philip Morris International; Procter & Gamble; Taiwan Semiconductor Manufacturing Company; Thermo Fisher Scientific), for a total number of 73 projects and an average of 4,8 projects per company. Two of the analyzed companies were found to be without active Artificial Intelligence initiatives.

As can be seen, the analyzed companies vary in terms of realized products and end market for their products. For instance, Caterpillar is the world leading manufacturer of heavy construction and mining equipment, engines and industrial turbines for a variety of sectors, above all Construction & Real Estate and Energy, Resources & Utility. Philip Morris International is the largest cigarette manufacturer in the world. Mitsubishi Electric Group, one of the main companies of Mitsubishi Motor, is a leader in manufacturing electronics and electrical equipment used for energy systems, industrial automation, electronic devices and home appliances. Procter & Gamble is a world leading manufacturer of consumer goods including beauty and grooming supplies, paper towels and toilet paper, razors and shaving creams, nappies, household care items, soap, detergents.

When considering the Status of the Project for Artificial Intelligence initiatives in the industry, results are shown in Figure 72.

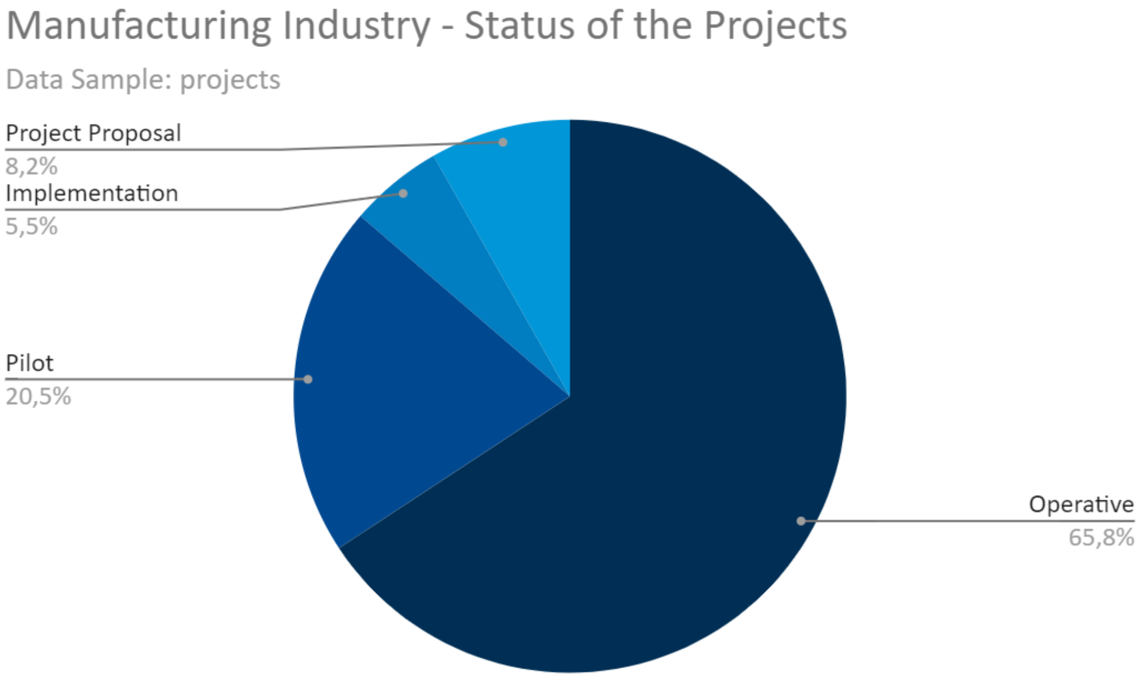

Figure 72: Status of the Projects in the Manufacturing Industry

As shown, the majority of the analyzed projects are already Operative, with the 65,8% of the cases. Initiatives in Implementation (5,5%), Pilot Projects (20,5%) and Project Proposals (8,2%) follow.

To provide a general overview about the adoption of Artificial Intelligence in the industry, the Qualitative Framework has been applied. Results are shown in Figure 73.

Level 1	Service	23,3%
	Product	21,9%
	Customer Interaction	4,1%
Level 2	Customer Profiling	1,4%
	Demand Forecasting	4,1%
	Marketing Optimization	4,1%
	New Product Development	5,5%
	Price Recommendation	1,4%
	Sentiment Analysis	2,7%
Level 3	External Logistics	1,4%
	Internal Logistics	1,4%
	Supply Chain Improvement	0,0%
	Regulatory Compliance	1,4%
	Risk Management	0,0%
	Employees Support	2,7%
	Engineering Asset Management	0,0%
	Maintenance	4,1%
	Operations Automation	4,1%
	Production Planning	5,5%
	Quality Control	6,8%
	Cybersecurity/Fraud Detection	0,0%
	Safety	0,0%
	Security	0,0%
	HR Management	0,0%
	Recruiting	4,1%

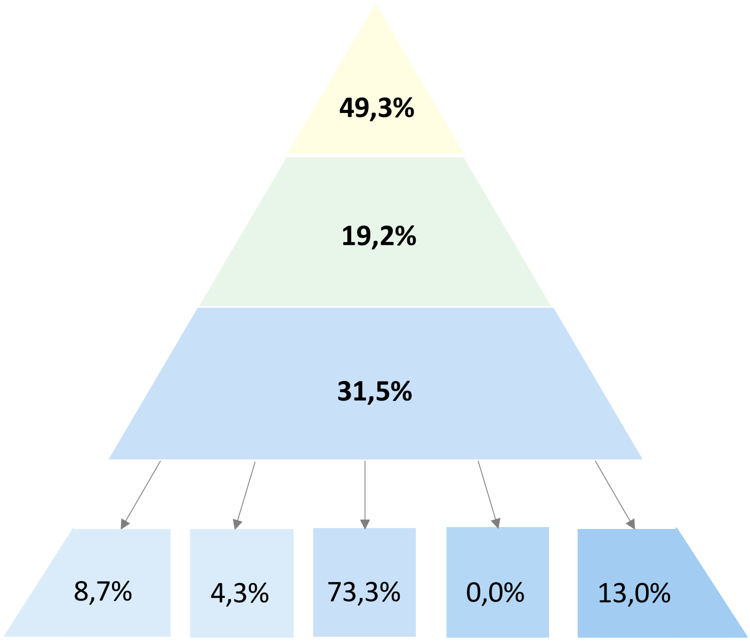

Figure 73: Qualitative Framework applied in the Manufacturing Industry

As shown, most of the projects regard Level 1 (AI Enabled Products & Services) and Level 3 (AI in Enterprise-Oriented Processes), with respectively the 49,3% and 31,5% of the initiatives. On one side, the diffusion of solutions at Level 1 suggests that the Manufacturing Industry is widely using Artificial Intelligence as enabler of the functionalities of Products and Services sold by the firms. On the other side, the majority of projects at Level 3 are related to Operations, suggesting the key role that Artificial Intelligence can play in the digital transformation of the factory, with several industrial applications of this technology possible. To conclude, despite projects mainly regard Level 1 and 3 of analysis, initiatives have been found at Level 2 (AI in Customer-Oriented Processes) too, supporting Marketing Optimization, Demand Forecasting and New Product Development.

In the following sections, each Level will be briefly discussed, to provide more information about the way Manufacturing companies are adopting Artificial Intelligence.

Level 1 - AI - Enabled Products & Services

Artificial Intelligence is showing great applicability at Level 1 in the Manufacturing Industry, with the technology as an enabler of the Products and Services sold by Manufacturing companies.

As regards Products, they are intended for both B2B and B2C markets, with several initiatives already Operative. In a B2B scenario, products range from elevators using Artificial Intelligence to optimize car assignments for each car of the elevator, to general-purpose inverters using Artificial Intelligence for failure diagnostics, to autonomous tractors and smart cameras using Computer Vision to distinguish between healthy and unhealthy crops. In a B2C scenario, Virtual Assistants are being progressively integrated into smart household appliances that people use in their everyday life. For instance, the appliance manufacturer Midea has incorporated Baidu's Xiadu Virtual Assistant into 15 categories of products such as heaters, rice cookers, air conditioners, and it is planning to introduce it into more and more product categories, like microwaves, lights, refrigerators.[108] Alongside this intuitive integration of Virtual Assistants in products, new interesting solutions based on different techniques are emerging. Among the others, they include: Intelligent Data Processing in toothbrushes, to understand and suggest areas requiring additional brushing, and air conditioners using Artificial Intelligence to optimize their performances; Computer Vision systems in refrigerators, to recognize foods and provide suggestions or health tips, and in washing machines, to recognize clothes and recommend washing methods to the user.

On the other side, the presence of several Services underlines how Artificial Intelligence is finding adoption in the industry not only in products, but also in the services that manufacturing companies offer to their clients. Also in this case, projects are at different implementation stages, ranging from Project Proposals to several already Operative initiatives.

The majority of the analyzed Services regards a B2B market, with manufacturing companies selling software solutions based on Intelligent Data Processing and Computer Vision. For instance, Caterpillar is not only manufacturing marine engines, but they have also launched a subscription-based Artificial Intelligence platform to support client shipping companies. Artificial Intelligence is used to reduce maintenance costs, monitoring the health of equipment and onboard systems to reduce the vessel's downtime, and to prove insights about operations.[109] In another project, an intelligent asset management system is using Machine Learning to analyze data collected by onboard sensors to allow the optimization of hull cleaning frequency, considering cleaning costs and improvements in performance led by different levels of cleanliness.[110] Or again, Deere & Company, the manufacturer of heavy equipment for the agriculture industry, offers to its customer a catalogue of high-tech data powered services. Artificial Intelligence algorithms are used to analyze field data and provide guidance to individual farmers, enabling them to get more and better crops from less resources[111]. Moving to a B2C scenario, Procter & Gamble and L'Oréal are providing to their clients solutions typically based on Computer Vision, with skin advisor services using image processing to analyze users skin and evaluate its condition, and apps using Computer Vision to let users virtually try on different haircuts or shades of lipsticks through their smartphones.

To conclude, some companies are implementing Artificial Intelligence solutions for Customer Interaction, with Chatbots used to let customers discover the full line of products of the company or to provide suggestions and information about specific items.

Level 2 - AI in Customer-Oriented Processes

At Level 2, New Product Development is an area in which manufacturing companies have implemented several solutions, exploiting the potential of Artificial Intelligence to bring new products to the market. Solutions are so different that the identification of a dominant scenario is not possible. For instance, Deere & Company is collecting data from the utilization of its machines and analyzing them through Artificial Intelligence to understand possible machine improvements and to guide the development of new products.[112] Meanwhile, 3M is testing the use of Machine Learning in supporting the new product development process, analyzing images to understand how changes in shape, size and orientation may have a positive impact on abrasiveness and durability.[113]

Alongside use of Artificial Intelligence for New Product Development, manufacturing companies are using the technology for other Applications at Level 2. Solutions range from Machine Learning in Customer Profiling to analyze data on browsing behaviours and identify customers segments, to Virtual Assistants used as influencers on digital media for Marketing Optimization, to Natural Language Processing solutions for Sentiment Analysis. To conclude, Intelligent Data Processing solutions can be used in Demand Forecasting to generate more reliable forecasts.

Level 3 - AI in Enterprise-Oriented Processes

Artificial Intelligence in the Manufacturing industry is becoming widespread at Level 3 too, especially in the Operations sub-Level. This points out the key role that Artificial Intelligence can have in the digital transformation of manufacturing plants, to concretize the concept of Smart Factory through several industrial applications of this technology.

For instance, in Quality Control several solutions based on Computer Vision are already operative inside the industry, automating the inspection of products and components and increasing the accuracy of controls if compared with traditional human inspection. In these solutions, Artificial Intelligence is used to control the processing quality of machines by autonomously detecting surface defects in products and components, improving the fault detection rate, and so product quality, and giving workers the chance to work on more value-added tasks.

In Production Planning, Artificial Intelligence is applied to optimize production and improve the process efficiency, with Intelligent Data Processing solutions to automatically adjust production parameters, cutting the time needed to set up optimal parameters manually, and to optimize production scheduling. New interesting initiatives are emerging too, in which Artificial Intelligence is used to analyze human motions of workers to improve the production process: for instance, Mitsubishi Electric is using Artificial Intelligence to analyse the human motions of assembly-line workers (measured through sensors or recognized in videos) and detect unnecessary, omitted or mistaken motions of workers, helping to upgrade and improve operations to increase work efficiency.[114]

For Operations Automation, companies are exploring solutions based on Autonomous Robots to automate operations in production plants, with projects still under development or at a pilot stage. For instance, Mitsubishi Electric has developed cooperative robots using Artificial Intelligence to learn and imitate the actions of skilled workers, and it is going to test them in its production sites-[115] Or again, the electronic manufacturer Hon Hai Precision's plant in Shenzhen, China, is completely automated and does not require human labour, with Artificial Intelligence used to run and maintain equipment.[116]

Artificial Intelligence is relevant in Maintenance too: manufacturing companies are typically using it to analyze massive amounts of sensor data for anomaly detection, improving the productivity of factories and production plants by alerting in case of possible failures.
Consequently, adoption of Artificial Intelligence solutions in the Manufacturing industry is particularly significant for Operations.

Besides these solutions, Artificial Intelligence can play a role also in the other sub-Levels. For instance, applications have been found in Internal Logistics, with Project Proposals for using Autonomous Vehicles in warehouses, and External Logistics, with Artificial Intelligence used to optimize the transportation network of Manufacturing companies.

To conclude, Figure 74 summarizes the analyzed initiatives in the Manufacturing sector, according to their Class of Solutions and the developed Qualitative Framework.

As shown, the most diffused initiatives regard Level 1, with the use of Intelligent Data Processing and Computer Vision as enablers of Products and Services sold to customers, and Level 3. Here the Operations sub-Level is dominant, with solutions based on Intelligent Data Processing and Computer Vision for Quality Control, Maintenance and Production Planning, as well as projects in early stages yet for the adoption of Autonomous Robots for Operations Automation.

Manufacturing Industry		IDP	NLP	VA/CHATBOT	CV	iRPA	RECOMM.	AR	AV	IO	
AI-Enabled Product & Services		14		4	8			1	6	3	49,3%
AI in Customer-Oriented Processes		9	1	1	2		1				19,2%
AI in Enterprise-Oriented Processes	Supply Chain Management	1							1		31,5%
	Business & Financial Control		1								
	Operations	5	2	1	6			3			
	Corporate Security										
	Human Resources	1	1	1							
		41,1%	6,8%	9,6%	21,9%	0,0%	1,4%	5,5%	9,6%	4,1%	

Figure 74: Application-Class of Solutions Matrix applied to the Manufacturing Industry

4.3 Retail Industry

The Retail Industry includes all those companies that sell finished products directly to the ultimate consumers, purchasing goods for personal use or consumption.

Retailers are not creating products, but they are purchasing them from wholesalers, or directly from manufacturers, and resell them in small quantities to end users at a mark-up. Different types of retail store exist: physical stores include, among the others, small groceries, supermarket chains, department stores; on the other side, online stores are used for E-Commerce, selling products to consumers through e-commerce websites. Consequently, also the customer experience is different based on the store, because of differences in terms of shopping process and purchasing experience. While in physical stores customers can have a physical contact with the products, in E-Commerce they see and buy them in online stores.

In both the cases, Artificial Intelligence technology can help companies operating in the Retail Industry to improve their businesses and offer a better shopping experience to customers.

The analysis of the industry has been carried out by considering the 12 Retail companies in the database (Alibaba; Amazon; Costco Wholesale; Home Depot; Lowe's; LVMH Moët Hennessy Louis Vuitton; Midea Group; Walgreens Boots Alliance; Seven & I Holdings; Target; Walmart), for a total number of 58 projects and an average of 4,8 projects per company.

When considering the Status of the Project for Artificial Intelligence initiatives in the Retail industry, results are shown in Figure 75.

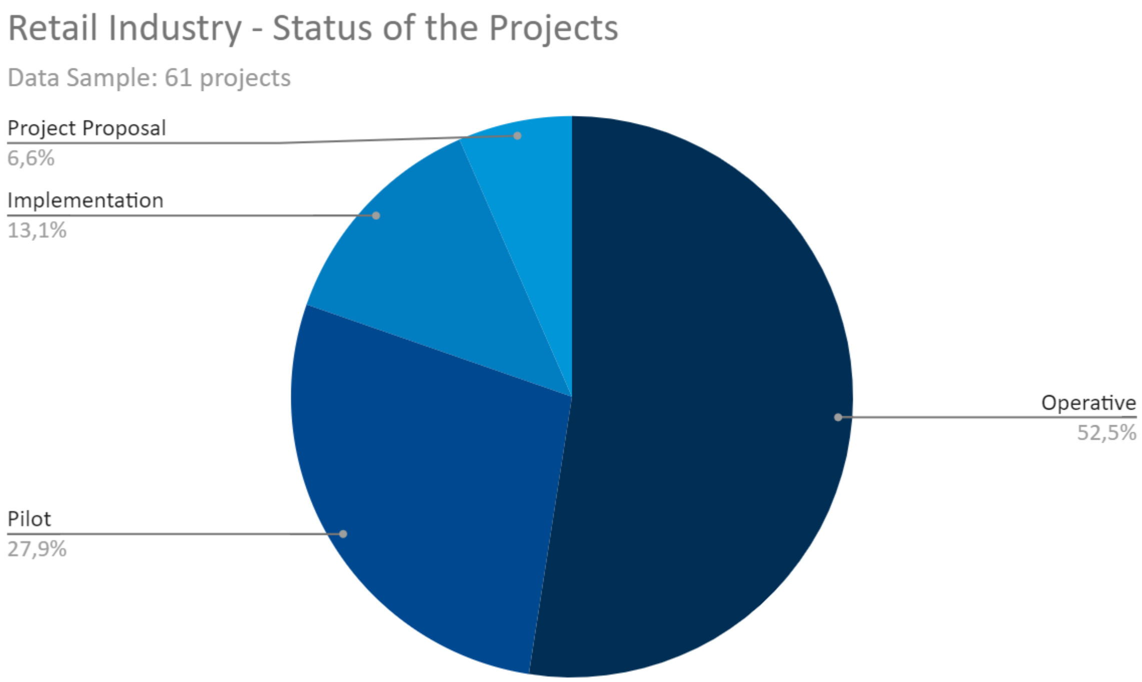

Figure 75: Status of the Projects in the Retail Industry

As shown, the 65,6% of Artificial Intelligence initiatives in the industry are Operative or in Implementation, suggesting positive outcomes of the implemented solutions. On the other side, the 34,5% of the initiatives are Pilots or Project Proposals, pointing out how several Retailers are now starting to explore the use of Artificial Intelligence at their advantage. To provide a general overview about the adoption of Artificial Intelligence in the Retail industry, the developed Qualitative Framework has been applied. Results are shown in Figure 76.

Level 1	Service	8,3%
	Product	0,0%
	Customer Interaction	20,0%
Level 2	Customer Profiling	5,0%
	Demand Forecasting	6,7%
	Marketing Optimization	18,3%
	New Product Development	0,0%
	Price Recommendation	3,3%
	Sentiment Analysis	0,0%
Level 3	External Logistics	8,3%
	Internal Logistics	6,7%
	Supply Chain Improvement	1,7%
	Regulatory Compliance	0,0%
	Risk Management	0,0%
	Employees Support	1,7%
	Engineering Asset Management	0,0%
	Maintenance	0,0%
	Operations Automation	13,3%
	Production Planning	0,0%
	Quality Control	1,7%
	Cybersecurity/Fraud Detection	1,7%
	Safety	0,0%
	Security	1,7%
	HR Management	1,7%
	Recruiting	0,0%

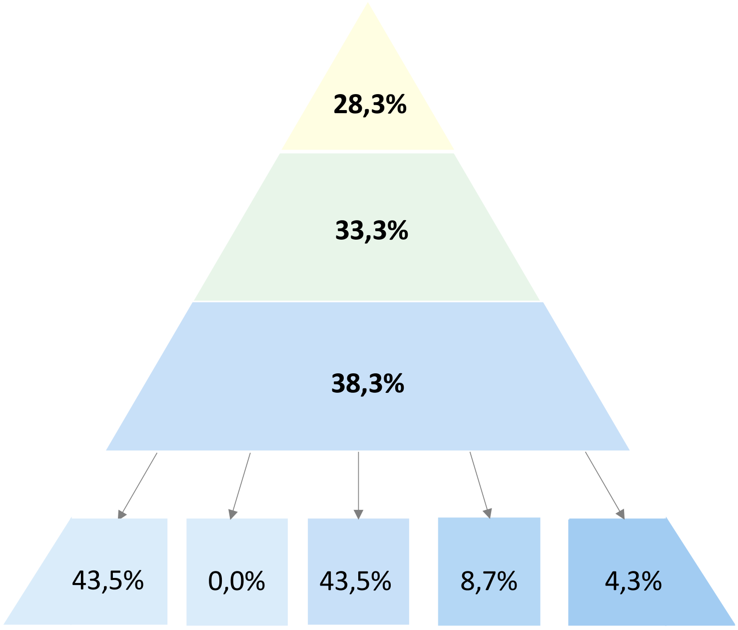

Figure 76: Qualitative Framework applied to the Retail Industry

As shown, companies in the Retail sector are adopting Artificial Intelligence solutions at all the levels, with quite similar diffusion rates. The high diffusion of initiatives at Level 1 (AI Enabled Products and Services) and 2 (Customer-Oriented Processes) can be explained by the fact that Retailers are in direct contact with the ultimate consumers, representing the last step of the Supply Chain before end users. As a consequence, the client is often at the center of the development of Artificial Intelligence Projects in the Retail Industry.

At Level 1 (28,3% of the projects) Artificial Intelligence is largely used by Retailers for Customer Interaction, particularly relevant in this sector for the closeness with consumers. Similarly, at Level 2 (33,3% of the projects) data about the customers and their behaviours are largely used for Marketing Optimization, Customer Profiling and to improve the accuracy of Demand Forecasting. Particularly, the use of Purchasing Recommendation Systems is relevant in online stores.

To conclude, Level 3 (AI in Enterprise-Oriented Processes) is the one with the highest diffusion of projects, underlying how Artificial Intelligence is also being internally applied by Retailers for Supply Chain Management (for Internal and External Logistics) and Operations Automation in physical stores.

In the following sections, each Level will be briefly discussed, to provide more information about the way Retailers are adopting Artificial Intelligence solutions.

Level 1 - AI - Enabled Products & Services

Level 1 includes the 28,3% of initiatives in the Retail sector. While companies in this industry are not making products with Artificial Intelligence capabilities because of the definition of Retail itself, solutions for Customer Interaction are finding widespread adoption in the industry. Particularly, this single Application is the one with the highest number of initiatives in the sector. Soon after, Services enabled by Artificial Intelligence are spreading too.

In Customer Interaction, the great majority of solutions regard the interaction with customers in online stores. Particularly, solutions mainly refer to the use of customer service Chatbots and Virtual Assistants in the online channels of retailers. These are mature solutions that have found widespread adoption in other industries too, but are even more relevant in the Retail sector because of the direct contact with consumers. Thanks to Chatbots and Voicebots powered by Artificial Intelligence, retailers can improve the customer engagement and increase sales in their e-commerce platforms. Handling spoken or written queries, they are used to answer FAQs, provide information about the status of orders, and support users in finding the right products when provided with text or voice description.

Alongside these well-known solutions, new initiatives often at an embryonic stage are emerging, such as the use of Computer Vision to support visual search of products, allowing users to upload pictures and supporting them in finding what they are looking for or similar items. Despite the focus on online stores, solutions can be found in physical stores too. For instance, the international chain of convenience stores Seven-Eleven has launched an experimental outlet in Tokyo in which customers can enter the outlet and make payments using Biometric Recognition. [117]

Meanwhile, new Services for consumers supported by Artificial Intelligence are emerging. For instance, several retailers are integrating voice order services in the Virtual Assistants of Google and Amazon, allowing consumers to make orders through smart home devices or their smartphones. Or, for example, Walmart launched in a pilot a shopping service based on Artificial Intelligence at a subscription cost of 50$ a month. Combining E-commerce with a personal shopper Chatbot, users can inform the chatbot about a specific product they need, and the Chatbot answers and delivers the product for the same or the next day[118].

Level 2 - AI in Customer-Oriented Processes

Level 2 includes the 33,3% of projects in the Retail Industry. At this level, data about customers and their behaviours are used for Marketing Optimization, Demand Forecasting, Customer Profiling and Price Recommendation.

Particularly, Marketing Optimization is the Application with the highest number of initiatives at this level, because of the importance for retailers to attract customers and entice them to purchase products. Solutions typically regards the use of Recommendation Systems, that all the companies have introduced, or are introducing, in their online stores. Using Artificial Intelligence, Recommendation Systems are used to show in online stores products based on the past behaviour and preferences of the users, enabling personalization of the shopping experience. The best-known example is the recommendation engine of Amazon, using data about the individual preferences of users and their shopping history to recommend products of interest for the customer.[119]

Artificial Intelligence is being used in the Retail industry also for Demand Forecasting, Customer Profiling and Price Recommendation. In Demand Forecasting, Intelligent Data Processing solutions are used to predict demand for both physical and online stores: Artificial Intelligence algorithms analyze historical sales data and information about upcoming sporting events, holidays, and weather conditions to predict demand for products, so that companies can refurnish their stocks. For instance, Walmart is using Artificial Intelligence to weekly forecast the demand of more than 500 million store-item combinations, so that when customers enter stores, they will find the desired product on the shelves.[120]

On the other side, Intelligent Data Processing solutions are adopted by retailers for Customer Profiling, to analyze customer data and classify them in clusters and provide them with specific offers. For instance, Walgreens Boots Alliance, a global leader in retail

pharmacy, has launched a pilot using Artificial Intelligence to analyze customer behaviours and purchase data to understand the preferences of clients.[121]
To conclude, Retailers are using Intelligent Data Processing solutions also for Price Recommendation, leveraging Artificial Intelligence to competitively price products with a better understanding of demand elasticity.

Level 3 - AI in Enterprise-Oriented Processes
Artificial Intelligence is showing great applicability in the Retail Industry at Level 3, with solutions adopted for Supply Chain Management and Operations in physical stores. This is the level with the highest number of initiatives in the industry.

Supply Chain Management
As regards Supply Chain Management, Internal and External Logistics are very relevant activities for a retailer, especially when it comes to online stores. These two Applications represent the 15% of projects in the Retail Industry, underlying the role that Artificial Intelligence can have in supporting logistic activities of retailers.

The use of Artificial Intelligence in External Logistics mainly refers to the use of Intelligent Data Processing algorithms for optimization of the shipping process and Autonomous Vehicles. On one side, Optimization algorithms are used to optimize delivery routes of home deliveries and to assign delivery trips to vehicles, considering a multitude of variables such as time slots, the quantity of orders in them, available vehicles and their type, distances, possible delays due to weather or traffic. For instance, Walmart uses an Artificial Intelligence algorithm to determine the fastest and cost-efficient delivery routes of vehicles, using data like the distance between the store and the shipping address, the traffic and general conditions of the road.[122]
On the other side, projects related to Autonomous Vehicles are Pilots or just Project Proposals, and typically require collaboration with third parties to internally adopt their solutions. Terrestrial or aerial self-driving vehicles, such as cargo vans, drones, or smaller solutions, can be used for transportation between fulfilment centers or for deliveries to final consumers. For instance, Amazon is testing a small autonomous vehicle to autonomously make package deliveries in a small neighborhood in Washington’s Snohomish County.[123]

In Internal Logistics, different solutions are under exploration by retailers, ranging from Intelligent Data Processing solutions to sequence the picking and packaging of items and Computer Vision, to Autonomous Robots to pack items and Autonomous Vehicles.
For instance, Amazon is testing the use of Computer Vision to eliminate the need for a handheld scanner: when workers pick up an item, they slide it under a scanner mounted nearby able to recognize it, and place it in a bin. The system is able to recognise the location of the item also for future references, without the need to scan the bin again, as well as to analyze hand movements to understand when a person is placing an item inside a shelf slot.[124]

Operations

The other relevant sub-Level is Operations, particularly with Operations Automation: several solutions, typically in a pilot stage, are currently being tested by retailers to automate basic activities and tasks in physical stores. Many of these solutions regard the deployment of Computer Vision technology.
The most diffused use case refers to the use of Computer Vision to automate the replenishment of shelves, essential for a retailer since they should be kept always full. Cameras can be used to detect when holes appear in shelves, sending real-time alerts to staff so that items can be restocked.
For instance, Walmart launched a Pilot project in some of its stores using Artificial Intelligence-enabled cameras to control availability of products on the shelves and alert about the need to restock them. The solution is also used to understand if products are on the shelves from too long, determining for example how bananas are ripe from their colour, and alerting workers in case of replacement needed.[125]
Other solutions include the use of Computer Vision and cameras to make sure shopping carts are always available, to open more registers in case of long queues, to check for abandoned shopping carts. More complex solutions. always at a pilot stage, regard the use of autonomous robots for autonomous floor cleaning or to check inventory, prices and misplaced items, as well as Amazon Go Grocery, the first physical store without registers or checkouts. Here Artificial Intelligence is used to recognize customers and the products they take, keeping track of them in a virtual cart, in a "Just Walk Out" shopping experience.[126]

Corporate Security

Even if the number of initiatives at this Sub-Levels is significantly lower, some retailers are adopting Artificial Intelligence for Security in physical stores, or for Cybersecurity and Fraud Detection, in this case more related to online stores.

For instance, Walmart uses Artificial Intelligence in thousands of its stores: videos from surveillance cameras are analyzed for real-time detection of thefts at registers and self-checkout kiosks. If a customer places an item in a bag without having it previously scanned, the system recognizes it and automatically informs store associates.[127]

In another project, Home Depot is using Intelligent Data Processing to identify anomalies and frauds in online transactions.[128]

To conclude, Figure 77 summarizes the analyzed initiatives in the Retail sector, according to their Class of Solutions and the developed Qualitative Framework.

As shown, Virtual Assistants and Chatbots are widely used at Level 1, for Customer Interaction or as enablers of the Services of the company.

Level 2 is characterized by the adoption of Recommendation Systems for Marketing Optimization and Intelligent Data Processing algorithms for Customer Profiling and Demand Forecasting.

To conclude, at Level 3 the maturity of solutions is lower and mainly regards the adoption of Computer Vision solutions for Operations Automation. Meanwhile, External and Internal Logistics are taking advantage of solutions mainly based on Intelligent Data Processing and Autonomous Vehicles.

Retail Industry		IDP	NLP	VA/CHATBOT	CV	iRPA	RECOMM.	AR	AV	IO	
AI-Enabled Product & Services			1	9	5		2				27,9%
AI in Customer-Oriented Processes		6	2		1		11				32,8%
AI in Enterprise-Oriented Processes	Supply Chain Management	5			1			1	4		39,3%
	Business & Financial Control										
	Operations	1		2	6			1			
	Corporate Security	1			1						
	Human Resources					1					
		21,3%	4,9%	18,0%	23,0%	1,6%	21,3%	3,3%	6,6%	0,0%	

ùFigure 77: Application-Class of Solutions Matrix applied to the Retail Industry

4.4 Banking & Finance Industry

The Banking & Finance Industry includes both Commercial Banks, providing general services like provision of savings and transactional accounts, mortgages, personal loans, debit cards, and credit cards to the general public, and Investment Banks, financial services companies or banks' divisions mainly focused on advisory investments for their clients, that could be people, businesses or even the government.

The companies analyzed in the Banking & Finance Industry are 20 (Agricultural Bank of China; Bank of America; Bank of China; BNP Paribas; China Construction Bank; China Merchants Bank; Citigroup; Goldman Sachs Group; HSBC Holdings; ICBC; JPMorgan Chase; Mitsubishi UFJ Financial; Morgan Stanley; Ping An Insurance Group; Postal Savings Bank Of China (PSBC); RBC; Santander; Sumitomo Mitsui Financial Group; TD Bank Group; Wells Fargo), for a total number of 115 projects and an average of 5,8 projects per company.

When considering the Status of the Project for Artificial Intelligence initiatives in the Banking & Finance Industry, results are shown in Figure 78.

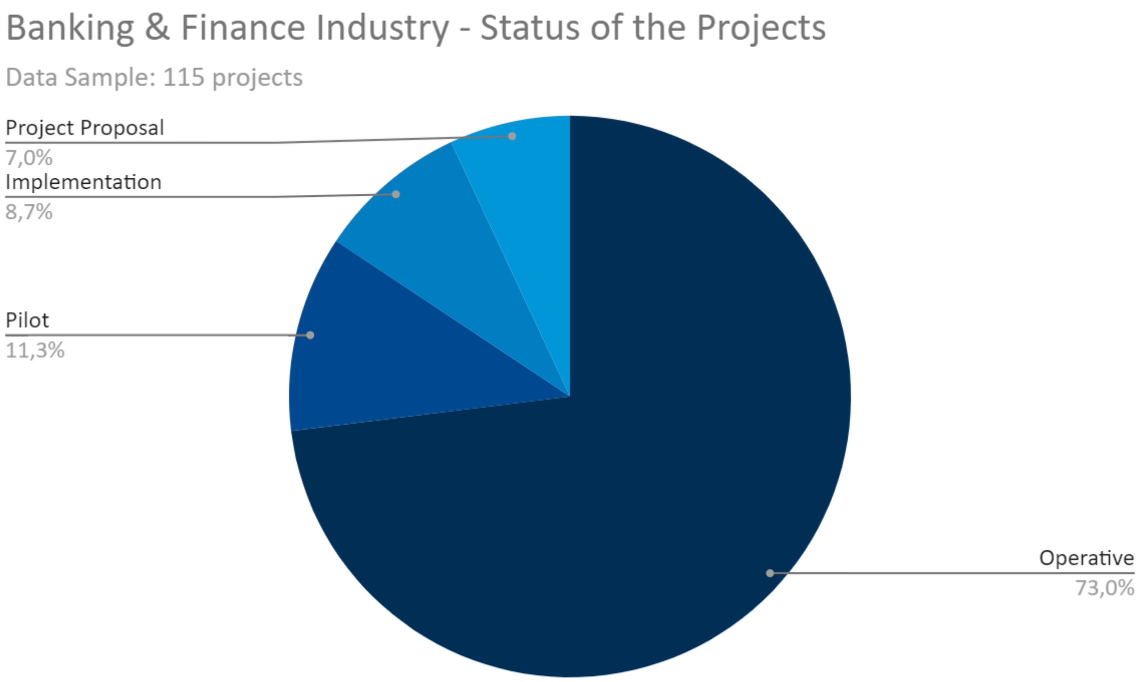

Figure 78: Status of the Projects in the Banking & Finance Industry

As shown, the majority of solutions are already Operative (73% of the projects in the industry) or in Implementation (8,7%), while the remaining 18,3% of solutions are Pilot projects (11,3%) and Project Proposals (7%). This underlines how, generally speaking, the vast majority of Artificial Intelligence solutions in the Banking & Finance sector have a high maturity, with solutions widely used in everyday processes. The results are confirmed by comparing these results with the average ones shown in Figure 15, suggesting that solutions applied in this industry are consolidated and have found widespread adoption in the recent years. Meanwhile, the several Pilots and Project Proposals point out how some companies are still testing Artificial Intelligence solutions or exploring new innovative ways to exploit the technology.

To provide a better understanding of the Industry, the developed Qualitative Framework has been applied. Results are shown in Figure 79.

Level 1	Service	18,3%
	Product	0,0%
	Customer Interaction	23,5%
Level 2	Customer Profiling	0,9%
	Demand Forecasting	1,7%
	Marketing Optimization	3,5%
	New Product Development	0,0%
	Price Recommendation	0,0%
	Sentiment Analysis	4,3%
Level 3	External Logistics	0,0%
	Internal Logistics	0,0%
	Supply Chain Improvement	0,9%
	Regulatory Compliance	1,7%
	Risk Management	5,2%
	Employees Support	5,2%
	Engineering Asset Management	0,0%
	Maintenance	1,7%
	Operations Automation	18,3%
	Production Planning	0,0%
	Quality Control	0,0%
	Cybersecurity/Fraud Detection	13,0%
	Safety	0,0%
	Security	0,0%
	HR Management	0,0%
	Recruiting	1,7%

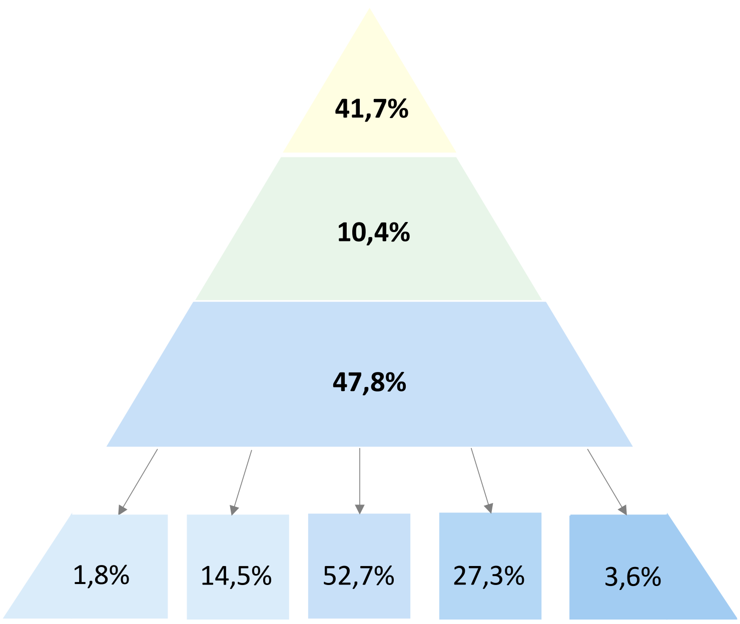

Figure 79: Qualitative Framework applied to the Banking & Finance Industry

As shown, most of the initiatives in the Banking and Finance Industry are related to Level 1 (AI Enabled Products & Services) and 3 (AI in Enterprise-Oriented Processes), with respectively the 41,7% and 47,8% of the projects in the industry.

At Level 1, Artificial Intelligence is at the basis of the Services that banks offer to their clients, or it is used for Customer Interaction. On the other side, at Level 3 banks are using this technology to make processes more effective and efficient, and especially for

Automations Operations and Cybersecurity/Fraud Detection, with respectively the 52,7% and 27,3% of projects at Level 3.

Level 2 (AI in Customer-Oriented Processes) includes only the 10,4% of projects in the industry: this suggests that banks are limitedly interested in applications at this level, but for their purposes the adoption of Artificial Intelligence is much more relevant at Level 1 and 3. Generally, it can be noticed how 70% of the initiatives can be referred to only 4 Applications: Services and Customer Interaction at Level 1, Operations Automation and Cybersecurity/Fraud Detection at Level 3.

Level 1 - AI Enabled Products & Services

As previously mentioned, Level 1 includes the 41,7% of projects in the industry, with the 23,5% of total projects for Customer Interaction and the 18,3% in Services.

Solutions for Customer Interaction mainly refers to the use of Chatbots powered by Artificial Intelligence for customer service. In the industry, Chatbots are used to support customers and to handle banking services, through a myriad of channels including SMS, mobile banking apps, WeChat, Facebook Messenger, Internet channels. Chatbots are used to provide information to customers, support them in managing and reviewing their bank accounts, fulfil requests, make recommendations and for several other purposes. It can be noticed that on one side this is a very mature typology of solutions, with almost all the projects that have become Operative in the recent years; on the other, 19 out of 20 of the analyzed companies are using Chatbots, pointing out how this kind of technology has become a standard in the industry.

Meanwhile, a second group of solutions related to the interaction between banks, and the services they offer, and customers, regards the use of Biometric Recognition. This has found widespread adoption in recent years in China: Chinese banks have extensively adopted Artificial Intelligence to allow their clients authentication, introducing facial recognition to withdraw cash and make payments at ATMs or in mobile phone banking applications. This use of Biometric Recognition eliminates the need for typing lengthy personal ID and security codes along the authentication procedure, allowing a simpler and much more secure identification of clients. For instance, the Agricultural Bank of China has launched in 2017 a withdrawal service based on facial recognition at over

20.000 branches across China; clients must verify their identity through Biometric Recognition at the ATM before withdrawing cash.[129]

As regards Services enabled by Artificial Intelligence, in the Banking & Finance industry they typically are investment advisory services, also known as robo-advisor services, extensively using Artificial Intelligence to offer investment recommendations to customers. These solutions can use Artificial Intelligence algorithms to forecast the asset price for a certain investment, rather than to tailor the portfolio of investments for a specific client by considering not only the expected return, but also the customer's risk appetite and his historical trading preferences. These kinds of solutions, as Chatbots for Customer Interaction, have generally spread in the industry in recent years.

For Instance, Goldman Sachs uses an Intelligent Data Processing Forecasting algorithm that analyzes various investments and provides the user with a prediction on the profitability of a certain investment. This analysis can reveal if a company is undervalued, and so possibly worth buying.[130] Or Again, China Merchants Bank launched in 2016 a robo-advisor using Artificial Intelligence to tailor the best investment portfolio based on the characteristics of the specific client, providing tailored recommendations. This algorithm starts from the choices of the client in terms of risk and return and then analyzes a multitude of global asset allocation or funds, to construct the best portfolio for the client.[131]

Level 2 - AI in Customer-Oriented Processes

Level 2 includes only the 10,4% of projects in the industry, with a limited number of initiatives related to Customer-Oriented Processes. However, projects related to Sentiment Analysis, Marketing Optimization, Demand Forecasting and Customer Profiling have been found.

For instance, in Sentiment Analysis Natural Language Processing is used to understand public opinions about the company or to identify issues to improve customer satisfaction. In one project, TD Bank Group is applying Natural Language Processing to transcriptions of calls to its call centers, to identify the mentioned topics and analyze the sentiment of the customer.[132]

Meanwhile, tailored solutions for the industry are emerging, that do not consider the sentiment of customers, but for instance the one of tweets and news articles to understand the opinion about firms to eventually suggest equity investments.
On the other side, in Marketing Optimization solutions based on Natural Language Generation can be used to generate more effective marketing contents, and Recommendation Systems can be used to suggest to clients services that could be interesting for them.

Level 3 - AI in Enterprise-Oriented Processes

Level 3 is the one with the highest diffusion of Artificial Intelligence Initiatives in the Banking & Finance industry. Specifically, the 52,7% of the projects at Level 3 regards the Operations sub-Level, the 27,3% Corporate Security and the 14,5% Business and Financial Control.

Operations

At this sub-Level, the most relevant Applications are Operations Automation and Employee Support, representing respectively the 18,3% and 5,2% of total projects in the Banking Industry.

Operations Automations includes a variety of solutions to automate processes or single tasks, ranging from Intelligent Data Processing and Natural Language Processing, to iRPA and Autonomous Robots.
Particularly, iRPA systems are widely diffused in the industry to automate a multitude of repetitive tasks. The adoption of these systems is extremely beneficial for the Banking & Finance industry, since it is generally considered one of the most data intensive sectors.
When it comes to Natural Language Processing, the industry requires both the analysis and the creation of immense quantities of documents: as a consequence, some banks are using Natural Language Processing to automatically extract relevant information from documents or to create reports. For instance, to correctly suggest investment to its customers, Goldman Sachs is using Information Retrieval to extract information from analysis and news reports.[133]

To conclude, some banks are piloting the use of humanoid Autonomous Robots in their branches, to welcome guests and teach them how to open accounts, cracking jokes, relaying credit card details and more.

As regards Employee Support, banks have introduced solutions based on Natural Language Processing and enterprise Virtual Assistants/Chatbots to support the employees in the execution of their day-by-day activities. For instance, Morgan Stanley is testing in several departments an internal Virtual Assistant, leveraging Artificial intelligence to provide employees answers to questions like the market value of a certain company.[134] On the other side, Natural Language Processing can be used by employees to filter documents when searching for internal information.

Corporate Security

At this sub-Level, all the projects refer to the Cybersecurity/Fraud Detection Application. Cybersecurity and Fraud Detection are vitally important in the Banking & Finance sector, to protect customer assets and keep their trust and to prevent financial losses. Particularly, Artificial Intelligence technology is widely used to support these activities, with the use of Intelligent Data Processing solutions for anomaly detection.

For instance, in Fraud Detection these algorithms are used to identify possible frauds, regarding for example credit cards or online transactions, and for anti money-laundering purposes. In Cybersecurity, Intelligent Data Processing algorithms are used to identify cyber-attacks targeting the information systems of the bank.

Business & Financial Control

At this sub-Level, the most relevant application is Risk Management, covering the 5,2% of projects in the industry. Banks are using Artificial Intelligence, with solutions based on Intelligent Data Processing, to analyse data related to clients and companies asking for loans to evaluate the risk exposure of the business. For instance, HSBC is using Artificial Intelligence to analyze supply chain data of the different client companies, to better understand which companies may have liquidity problems during supply chain disruption or during economic turmoil.[135]

To conclude, Figure 80 summarizes the analyzed initiatives in the Banking & Finance industry, according to their Class of Solutions and the developed Qualitative Framework.

Banking & Finance		IDP	NLP	VA/CHATBOT	CV	iRPA	RECOMM.	AR	AV	IO	
AI-Enabled Product & Services		13	1	19	9		5	1			41,7%
AI in Customer-Oriented Processes		3	7				2				10,4%
AI in Enterprise-Oriented Processes	Supply Chain Management	1									47,8%
	Business & Financial Control	7	1								
	Operations	6	10	2		8		3			
	Corporate Security	14			1						
	Human Resources		1	1							
		38,3%	17,4%	19,1%	8,7%	7,0%	6,1%	3,5%	0,0%	0,0%	

Figure 80: Application-Class of Solutions Matrix applied to the Banking & Finance Industry

As shown, most of the solutions in the Banking & Finance industry are based on Intelligent Data Processing, Natural Language Processing and Virtual Assistant/Chatbot, mainly at Level 1 and 3.

At Level 1, Intelligent Data Processing is at the basis of investment advisor services, while Chatbots and Voicebots are widely diffused for Customer Interaction.

At Level 3, Intelligent Data Processing algorithms are applied especially in Cybersecurity/Fraud Detection, while Intelligent Automation and Natural Language Processing solutions are typically applied for Operations Automation.

To conclude, the Banking & Finance Industry has widely adopted Artificial Intelligence solutions, to improve processes and offer a higher value to their customers.

5. Conclusion

Despite Artificial Intelligence has been an area of research for a long time, with the birth of the discipline dating back to 1950 with Alan Turing's seminal paper "Computing Machinery and Intelligence", practical applications in businesses have emerged only in recent years, generating great interest by public and private companies.

The primary aim of this thesis work was to answer a precise research question: *What is the present state about the international adoption of Artificial Intelligence solutions?*
Indeed, although the presence of an extensive body of literature related to Artificial Intelligence techniques, information about the adoption of Artificial Intelligence in enterprises is limited, and mainly referred to past years. However, Artificial Intelligence is a rapidly growing technology, and its diffusion is rapidly expanding. Therefore, its current state of adoption changes year after year, implying the necessity to provide a constantly updated picture about its evolution.

The results of the analysis show how Artificial Intelligence is rapidly spreading in a business context, with a variety of possible applications in the most varied sectors. Particularly, detailed results have been provided for each of the nine Classes of Solutions proposed by the Artificial Intelligence Observatory of Politecnico of Milan, such as the most relevant application fields and the involved sectors. Currently, solutions based on Intelligent Data Processing are the most diffused ones, constituting the 37,2% of the analyzed initiatives and pointing out the relevance of Artificial Intelligence algorithms in making sense of the impressive amount of data companies are overwhelmed with. Computer Vision (17,7%), Virtual Assistant/Chatbot (14%) and Natural Language Processing solutions (11,3%) are widely spreading in companies too. Particularly, Chatbots powered by Artificial Intelligence can be considered a mature technology that have found widespread adoption in almost all the sectors. Meanwhile, other solutions require further development before having a wider diffusion. Particularly, Autonomous Robots (4%) and Autonomous Vehicles (5,7%) are still at an embryonic stage, and further advancements are needed because of the complex Artificial Intelligence capabilities required.

Furthermore, the research highlights how the benefits on businesses of these solutions are so relevant that no company can ignore the ongoing Artificial Intelligence revolution and stay out of it. This becomes even more real by considering that, in many industries, companies have just started to explore the opportunities offered by this disruptive technology. Therefore, the future of Artificial Intelligence is expected to be very promising.

In addition to answering this first research question, the secondary aim of this thesis work is to provide a focus on how Artificial Intelligence is being adopted in specific industries, answering the research question *RQ2: How AI adoption changes from industry to industry and what is the potential contribution for the sector?*
Particularly, the Food & Beverage, Manufacturing, Retail and Banking & Finance sectors have been studied, supporting their analysis with the development of a specific Qualitative Framework.

Results show how in the Food & Beverage industry Artificial Intelligence is mainly used to gain a better understanding of customers, with the aim to bring new products to the market or provide tailored marketing solutions, and to improve the whole Food & Beverage Supply Chain. Moreover, the technology is finding widespread adoption at the factory level. In Manufacturing, Artificial Intelligence is mainly used as an enabler of the products and services sold by firms, and it is driving the digital transformation of the factory. In Retail, Artificial Intelligence is finding widespread adoption at different levels. The most relevant uses are in customer service, tailored purchasing recommendations in online stores, logistics activities and for the automation of operations in physical stores. To conclude, in the Banking & Finance Industry several solutions are now mature and consolidated in recent years. In this sector, Artificial Intelligence is typically at the basis of Chatbots to improve the customer experience, investment advisory services, and internal solutions for automation of the processes, cybersecurity and fraud detection.

To conclude, some limitations of this research work can be highlighted.
First, the sample of companies present in the database is limited. Despite the number of identified initiatives is high, the limited number of firms in the database does not allow a comparison between the investments of different industries, or an evaluation about the

effectiveness of national strategies for Artificial Intelligence development based on the projects' country of implementation.

Second, the selection of companies has been done by following the methodology described in Chapter 2, based on the Forbes Global 2000 list. As a consequence, all the analyzed companies are the largest public ones in the world, and the only perspective taken into account is the one big companies leaders in their sector, that have been the first to start investing in Artificial Intelligence .

Lastly, the database has been created by using secondary sources, like managers' interviews, news and articles, information on corporate websites. The use of these sources may have limited the visibility of certain projects and Classes of Solutions, like iRPA, because of a lower press coverage if compared with other solutions. Meanwhile, when considering the status of the projects, the intercepted initiatives are just the ones made publicly available, and these typically are the ones at the most advanced stages.

To conclude, directions for further development in research are provided.

Firstly, it is essential to continue the mapping of Artificial Intelligence initiatives at a global level also in future years, so that it is possible to have an always updated picture about the adoption of this technology, as well as to analyze and comment on its adoption's evolution year after year.

Then, the number of Artificial Intelligence initiatives is expected to increase even more year after year. Consequently, the workload needed to provide a satisfactory overview of the international adoption of Artificial Intelligence solutions is expected to be higher and higher. For this reason, a possible suggestion is the launch of several research projects in which the focus is on bundles of similar industries. Results may be later integrated to obtain a comprehensive overview about the global diffusion of the technology.

Moreover, it could be interesting to introduce a higher level of detail for some Applications, to provide an even more detailed overview. For instance, Services and Operations Automation are the most diffused Applications at a global level, with a high number of projects. It can be useful, for instance, to distinguish between different operations, like automation of manufacturing processes and the one of back-office operations.

To conclude, the developed Qualitative Framework could be adopted also in future research, to analyze different sectors or the same ones with larger sample volumes, to provide more detailed overviews of the industries. Furthermore, the adoption of the model in future research may be relevant to improve it, adjusting the framework based on the future adoption of Artificial Intelligence solutions.

Bibliography

Accenture. (2018). *Chatbots are here to stay, So what are you wating for?* Accenture.

Adomavicius, G., & Tuzhilin, A. (2005). *Toward the Next Generation of Recommender.* IEEE Computer Society.

Agatonovic-Kustrin, S., & Beresford, R. (1999). *Basic concepts of artificial neural network (ANN) modeling.* Dunedin, New Zealand: School of Pharmacy, Uni6ersity of Otago.

Agostinelli, S., Marrella, A., & Mecella, M. (2020). *Towards Intelligent Robotic Process Automation for BPMers.* Sapienza Universita di Roma.

Ajah, I. A., & Nweke, H. F. (2019). *Big Data and Business Analytics: Trends, Platforms, Success Factors and Applications.* Abakaliki: Ebonyi State University.

Annor-Antwi, A., & Al-Dherasi, A. A. (2019). *Application of Artificial Intelligence in Forecasting: A Systematic Review.* Zhejiang: Zhejiang University of Science and Technology.

Berruti, F., Nixon, G., Taglioni, G., & Whiteman, R. (2017). *Intelligent process automation: The engine at the core of the next-generation operating model.*

Bimbraw, K. (2015). *Autonomous Cars: Past, Present and Future. A Review of the Developments in the Last Century, the Present Scenario and the Expected Future of Autonomous Vehicle Technology.* Scitepress.

Buchanan, B. G. (2006). *A (Very) Brief History of Artificial Intelligence .* American Association for Artificial Intelligence.

Burger, R., & Wheelock, C. (2015). *Computer Vision Report: Market Trends, M&A Transactions, Private Placement.* Tractica.

Burke, R., Felfernig, A., & Göker, M. H. (2011). *Recommender Systems: An Overview.* Association for the Advancement of Artificial.

Burnett, S. (2017). *Automating Content-Centric Processes with Artificial Intelligence (AI).* Everest Global.

Cohen, K. B. (2014). *Biomedical Natural Language Processing and Text Mining.* Elsevier.

Dale, R. (2016). *The return of the chatbots.* Cambridge University Press.

Grigorescu, S., Trasnea, B., & Macesanu, T. C. (2019). *A survey of deep learning techniques for autonomous driving.* Wiley.

Gruyer, D., Magnier, V., Hamdi, K., Claussmanna, L., Orfila, O., & Rakotonirainy, A. (2017). *Perception, information processing and modeling: Critical stages for autonomous driving applications.* Elsevier.

Gubbi, J., Buyya, R., Marusic, S., & Planiswami, M. (2012). *Internet of Things (IoT): A vision, architectural elements, and future directions.* Elsevier.

Guo, Y., Liu, Y., Oerlemans, A., Lao, S., Wu, S., & Lew, M. S. (2015). *Deep learning for visual understanding: A review.* Elsevier B.V.

Hamet, P., & Tremblay, J. (2017). *Artificial intelligence in medicine.* Elsevier Inc. .

Heung-yeung, & Xiao-dong. (2018). *From Eliza to XiaoIce: challenges and opportunities.* Frontiers of Information Technology & Electronic Engineering.

Hinds, P. J., Roberts, T. L., & Jones, H. (2011). *Whose Job Is It Anyway? A Study of Human-Robot Interaction in a Collaborative Task.* Lawrence Erlbaum Associates Inc. .

Hirschberg, J., & Manning, C. (2015). *Advances in natural language processing.* sciencemag.org.

Iwamura, Y., Shiomi, M., Kanda, T., Ishiguro, H., & Hagita, N. (2011). *Do Elderly People Prefer a Conversational Humanoid as a Shopping Assistant Partner in Supermarkets?* ATR Intelligent Robotics and Communication Laboratory.

Jannach, D., & Jugovac, M. (2019). *Measuring the Business Value of Recommender Systems.*

Jordan, M. I., & Michell, T. M. (2020). *Machine Learning: Trends, perspective and prospects .* sinecemag.org.

Joseph, A., Dutta, D., Gillard, A., & Seewooruttun, G. (2018). *Robotics and intelligent automation: Combining the power of human and machine.* Ernst & Young LLP.

Khurana, D., Koli, A., Khatter, K., & Singh, S. (2017). *Natural Language Processing: State of The Art, Current Trends and Challenges.* Faridabad: Manav Rachna International University.

Kragic, D., Gustafson, J., Karaoguz, H., Jensfelt, P., & Krug, R. (2018). *Interactive, Collaborative Robots: Challenges and Opportunities.*

Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2010). *ImageNet Classification with Deep Convolutional.*

Lane, N. D., Bhattacharya, S., Georgiev, P., Forlivesi, C., & Kawsar, F. (2015). *An Early Resource Characterization of Deep Learning on Wearables, Smartphones and Internet-of-Things Devices.* Acm.

LeCun, Y., Bengio, Y., & Hinton, G. (2015). *Deep Learning.* Macmillan Publisher.

Liddy, E. D. (2001). *Natural Language Processing.* Syracuse: Syracuse University.

Liu, W., Wang, Z., Liu, X., Zeng, N., Liu, Y., & Alsaasi, F. (2017). *A survey of deep neural network architectures and their applications.* Brunel University London, Uxbridge, Middlesex: elvister.

Lu, J., Wu, D., Mao, M., Wang, W., & Zhang, G. (2015). *Recommender system application developments: A survey.* Elsevier B.V. .

Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A., Veness, J., Bellemare, M. G., . . . King, H. (2015). *Human-level control through deep reinforcement.* Macmillan Publishers .

Mohammadi, M., Al-Fuqaha, A., Sorour, S., & Guizani, M. (2018). *Deep Learning for IoT Big Data and Streaming Analytics: A Survey.* IEEE Xplore.

Moujahid, A., Hina, M. D., Soukane, A., Ortalda, A., Tantaoui, M. E., Khadimi, A. E., & Ramdane-Cherif, A. (2020). *Machine Learning Techniques in ADAS: A Review.* IEEE Xplore.

Niemelä, M., Heikkilä, P., & Lammi, H. (2017). *A Social Service Robot in a Shopping Mall – Expectations of the Management, Retailers and Consumers.* VTT Technical Research Center of Finland .

Nimavat, K., & Champaneria, T. (2017). *Chatbots: An overview Types, Architecture, Tools and Future.* International Journal for Scientific Research & Development|.

Osservatorio Artificial Intelligence - Politecnico di Milano. (s.d.). *NLP Applications Analyses.* Politecnico di Milano.

Padhy, N., Mishra, P., & Panigrahi, R. (2012). *The Survey of Data Mining Applications And Feature Scope.* International Journal of Computer Science.

Perego, A., Tubaro, S., Migliarotta, G., Gatti, N., Piva, A., Restelli, M., . . . Broggi, M. (2019). *Artificial Intelligence: on your marks!* Politecnico di Milano - Dipartimento di Ingegneria Gestionale .

Perego, A., Tubaro, S., Miragliotta, G., Gatti, N., Piva, A., Restelli, M., . . . Broggi, M. (2020). *Articial Intelligence: Learn to fly!* Politecnico di Milano - Dipartimento di Ingegneria Gestionale.

Qiu, J., Wu, Q., Ding, G., Xu, Y., & Feng, S. (2016). *A survey of machine learning for big data.* EURASIP Journal on Advances in Signal Processing.

Reddy, N., Harichandana, U., Alekhya, T., & Rajesh. (2019). *A Study of Robotic Process Automation Among Artificial Intelligence.* Banglore: GITAM school of technology.

Righetti, A., & D'Aquino, D. (2018). *La Robotic Process Automation per ridurre i costi delle Operations.* NTT Data.

Rughi, A., & Ceriani, A. (2019). *Le prospettive future dell'Intelligent Automation, secondo le aziende italiane.* Deloitte.

Saha, H. N., Mandal, A., & Sinha, A. (2017). *Recent Trends in the Internet of Things.* IEEE .

Schmidhuber, J. (2014). *Deep Learning in Neural Networks: An Overview.* Galleria 2, 6928 Manno-Lucano Switzerland: University of Lugano & SUPSI.

Silwattananusarn, T., & Tuamsuk, K. (2012). *Data Mining and Its Applications for Knowledge Management : A Literature Review from 2007 to 2012.* International Journal of Data Mining & Knowledge Management Process.

Stieglitz, S., Brachten, F., & Kissmer, T. (2018). *Defining Bots in an Enterprise Context.*

Szeliski, R. (2010). *Computer Vision: Algorithms and Applications.* Springer.

Thompson, W. (2015). *Recommendation Systems: An Overview of System Types and Benefits – and Why You Should Implement With SAS®.* SAS.

Tsai, C.-W., Lai, C.-F., & Vasilakos, A. (2014). *Future Internet of Things: open issues and challenges.* Spinger.

Vassakis, K., Petrakis, E., & Kopanakis, I. (2018). *Big Data Analytics: Applications, Prospects and Challenges.* Spinger International Publishing.

Voulodimos, A., & Doulamis, N. (2018). *Deep Learning for Computer Vision: A Brief Review.* Hindawi.

Voulodimos, A., Doulamis, N., Doulamis, A., & Protopapadakis, E. (2018). *Deep Learning for Computer Vision: A Brief Review.* Hindawi.

Wiley, V., & Lucas, T. (2018). *Computer Vision and Image Processing: A Paper Review.* International Journal Of Artificial Intelegence Research.

Willcocks, L., Lacity, M., & Craig, A. (2017). *Robotic process automation: strategic transformation lever for global business services?* LSE Research.

Winfield, A., McDermid, J., Müller, V., Porter, Z., & Pipe, T. (2019). *Ethical Issues for Robotics and Autonomous Systems.*

Young, T., Hazarika, D., Poria, S., & Cambria, E. (2018). *Recent Trends in Deep Learning Based Natural Language Processing.*

Zanchin, B. C., Silva, R. A., Santos, M. M., & Linares, K. (2017). *On the Instrumentation and Classification of Autonomous Cars.* IEEE International.

Zhang, C. (2018). *Intelligent Process Automation in Audit.* China: Southwestern University of Finance and Economics.

Zhang, T., Li, Q., Zhang, C.-s., Liang, H.-w., Li, P., Wang, T.-m., . . . Wu, C. (2016). *Current trends in the development of intelligent unmanned autonomous systems.* Frontiers of Information Technology & Electronic Engineering.

Ziler, M., & Fergus, R. (2014). *Visualizing and Understanding.* Switzerland: Springer International Publishing.

Sitography

[1]https://www.verifiedmarketresearch.com/product/global-computer-vision-market-size-and-forecast-to-2025/

[2] https://www.verifiedmarketresearch.com/product/global-internet-of-things-iot-market-size-and-forecast-to-2026/

[3] https://www.forbes.com/sites/andreamurphy/2018/06/06/2018-global-2000-methodology-how-we-crunch-the-numbers/?sh=328d1ba41504

[4] https://www.bankinfosecurity.com/jpmorgan-chase-develops-early-warning-system-a-12855

[5] https://www.pymnts.com/news/digital-banking/2019/wells-fargo-aml-biometrics-ai/

[6] https://www.manulife.com/en/advanced-analytics.html

[7] https://gdmissionsystems.com/products/electronic-warfare/signaleye

[8] https://sg.asiatatler.com/life/hyundai-uses-technology-to-build-smart-sustainable-cars

[9] https://www.shell.com/energy-and-innovation/overcoming-technology-challenges/d

[10] https://www.globaltimes.cn/content/1194086.shtml

[11] https://www.insidetelecom.com/orange-utilizes-ai-to-run-network-operations-more-efficiently/

[12] https://blogs.nvidia.com/blog/2019/07/11/walmart-nvidia/; https://www.datanami.com/2019/03/22/how-walmart-uses-gpus-for-better-demand-forecasting/

[13] https://www.refinitiv.com/en/media-center/press-releases/2020/march/bank-of-china-tea ms-up-with-refinitiv-to-launch-new-foreign-exchange-price-prediction-app-deepfx; https://www.thetradenews.com/bank-china-unveils-ai-currency-price-prediction-app-eikon/

[14] https://www.businesswire.com/news/home/20190927005166/en/Ping-AI-Disease-Prediction-Paper-Published-EBioMedicine; https://kr-asia.com/reducing-contagion-ping-ans-ai-team-in-china-can-now-predict-flu-outbreaks-with-90-per-cent-accuracy; https://pharmaboardroom.com/interviews/geoff-kau-co-gm-cso-ping-an-smart-city-china/

[15] https://www.beckershospitalreview.com/artificial-intelligence/google-cloud-ai-harvard-release-covid-19-forecasting-model-6-details.html

[16] https://www.shell.com/energy-and-innovation/overcoming-technologychallenges/digital-innovation/artificial-intelligence.html; https://www.shell.us/businesscustomers/shellrechargeplus.html#iframe=L2Zvcm1zL2VuX3VzX3NoZWxsX3JlY2hhcmdlX3BsdXM ; https://apkpure.com/it/rechargeplus/com.shell.sitibv.smartercharge

[17] https://enterpriseiotinsights.com/20180925/channels/news/caterpillar-and-osisoft-debut-ai-platform

[18] https://www.automotivetestingtechnologyinternational.com/news/rd/volkswagen-demonstrates-quantum-computing-in-traffic-routing.html; https://www.dwavesys.com/sites/default/files/Dwave_VW_Case_Study.pdf

[19] https://www.forbes.com/sites/bernardmarr/2018/06/15/the-brilliant-ways-ups-uses-artificial-intelligence-machine-learning-and-big-data/#256a55c65e6d; https://www.technologyreview.com/2018/11/21/139000/how-ups-uses-ai-to-outsmart-bad-weather/; https://sustainability.ups.com/media/optimized-network_making-the-smart-logistics-network-a-reality.pdf

[20] https://enterprisersproject.com/article/2019/2/ai-how-bayer-crop-science-uses

[21] https://www-file.huawei.com/-/media/corporate/pdf/publications/ict/ict22-en.pdf; https://e.huawei.com/pl/publications/global/ictnewhorizonspodcasts/Stories/sinopec;

[22] https://digital.hbs.edu/platform-rctom/submission/machine-learning-to-save-millions-of-dollars-of-advertising-waste-at-pg/

[23] https://financialpost.com/sponsored/business-sponsored/how-artificial-intelligence-can-leverage-big-data-to-boost-trade

[24] https://www.telekom.com/en/company/topic-specials/beethoven-year-2020-special/details/artificial-intelligence-will-complete-beethoven-s-10th-symphony-587432

[25] https://all3dp.com/4/volkswagen-uses-ai-and-3d-printing-to-produce-lightweight-car-parts/; https://www.greencarreports.com/news/1123945_volkswagen-tests-ai-informed-design-process-to-cut-weight-and-flaunt-it

[26] https://www.forbes.com/sites/bernardmarr/2018/06/08/honda-gets-ready-for-the-4th-industrial-revolution-by-using-ai-big-data-and-robots/#6d7a1f3c4bc1

[27] https://www.maana.io/resources/shell-partners-with-maana-to-accelerate-digital-transformation/; https://www.maana.io/resources/shell-partners-with-maana-toaccelerate-digital-transformation/; https://www.bestpractice.ai/studies/shell_plans_to_automate_information_extraction_from_internal_documents_to_collate_answers_to_operational_problems_using_maanas_knowledge_platform

[28] https://expertsystem.com/it/resource/zurich-customer-interview-artificial-intelligence-for-underwriting/

[29] https://www.distilnfo.com/payer/2019/12/03/anthem-we-are-now-a-tech-company-first/

[30] https://press.aboutamazon.com/news-releases/news-release-details/aws-announces-five-new-machine-learning-services-reinvent-and/

[31] https://expertsystem.com/it/gruppo-generali-sceglie-lintelligenza-artificiale-di-expert-system-per-ottimizzare-processi-di-registrazione-e-richieste-di-rimborso/; https://www.youtube.com/watch?v=ax26GuDRS7E&feature=emb_logo

[32] https://blog.google/products/news/new-google-news-ai-meets-human-intelligence/

[33] https://www.geeksforgeeks.org/how-does-google-use-machine-learning/

[34] http://technoitworld.com/how-does-gmail-use-artificial-intelligence/; https://www.techspot.com/news/74533-new-ai-powered-gmail-feature-can-write-emails.html

[35] https://www.softwareone.com/it-it/blog/articles/2019/07/15/ai-in-the-workplace-microsoft-word-ideas

[36] https://www.marketingaiinstitute.com/blog/acrolinx-spotlight

[37] https://venturebeat.com/2019/12/31/6-ai-features-microsoft-added-to-office-in-2019/

[38] https://www.youtube.com/watch?v=aIFbQNE3KOI

[39] https://data-flair.training/blogs/how-google-uses-machine-learning/

[40] https://techwisegroup.com/weekly-tech-tips/easily-use-artificial-intelligence-in-word/

[41] https://www.digfingroup.com/xiaoi/; http://www.ccb.com/en/newinvestor/upload/20160513_1463122653/20160513145721469892.pdf https://www.bankingexchange.com/news-feed/item/7971-transforming-customer-banking-experiences-with-ai

[42] https://www.appsruntheworld.com/customers-database/customers/view/generali-france-france https://www.techcircle.in/2019/03/22/using-ai-rpa-to-ease-customer-pain-points-future-generali-life-insurance-s-byju-joseph; https://www.abtasty.com/resources/generalis-chatbot-streamline-customer-support/

[43] https://group.bnpparibas/en/news/banking-jobs-cto-artificial-intelligence-lab-bnp-paribas-cib

[44] https://www.ecommercetimes.com/story/85934.html; https://www.hdblog.it/2019/04/02/google-assistant-walmart-spesa-voce/

[45] https://www.accenture.com/gb-en/_acnmedia/Accenture/next-gen-5/insight-new-energy-consumer-2017/Accenture-NEC2017-May-the-Bots-POV.pdf

[46] https://www.citigroup.com/citi/news/2020/200206a.htm

[47] https://www.forbes.com/sites/bernardmarr/2019/04/05/the-fascinating-ways-pepsico-uses-artificial-intelligence-and-machine-learning-to-deliver-success/#147a85cb311e

[48] https://www.prnewswire.com/news-releases/ping-an-research-on-ai-applications-for-retinal-imaging-selected-at-top-international-medical-imaging-conference-isbi-301012716.html;

https://www.asmag.com/showpost/26667.aspx;
https://global.chinadaily.com.cn/a/202003/03/WS5e5e19c4a31012821727bfb2.html;
https://pharmaboardroom.com/interviews/geoff-kau-co-gm-cso-ping-an-smart-city-china/;
https://www.prnewswire.com/news-releases/ping-an-medical-technology-showcases-five-smart-medical-solutions-at-csce-300701997.html;

[49] https://www.geeksforgeeks.org/how-does-google-use-machine-learning/

[50] https://www.loreal.com/en/usa/articles/research-and-innovation/loreal-unveils-perso/

[51] https://www.press.bmwgroup.com/global/article/detail/T0298650EN/fast-efficient-reliable:-artificial-intelligence-in-bmw-group-production?language=en;
https://www.youtube.com/watch?v=icGLjUCshNE;

[52] https://www.intel.it/content/www/it/it/customer-spotlight/stories/midea-industrial-vision-customer-story.html;

[53] https://www.prnewswire.com/news-releases/tokio-marine-to-use-tractables-ai-to-analyse-auto-damage-across-japan-301054962.html;

[54] https://www.chinadaily.com.cn/regional/chongqing/liangjiang/2017-08/17/content_30770287.htm;
http://www.chinadaily.com.cn/regional/chongqing/liangjiang/2018-09/28/content_37003425.htm;

[55] https://www.engadget.com/2019-08-22-google-photos-face-grouping.html;

[56] https://www.bbc.com/news/technology-51309186;
https://www.forbes.com/sites/bernardmarr/2016/12/29/4-amazing-ways-facebook-uses-deep-learning-to-learn-everything-about-you/#1d52a5d4ccbf;

[57] https://iktva.sa/wp-content/uploads/2020/03/Industrial-Revolution-4.0-Digital-Transformation.pdf;

[58] https://www.shell.com/energy-and-innovation/overcoming-technology-challenges/digital-innovation/artificial-intelligence.html;
https://customers.microsoft.com/it-it/story/shell-mining-oil-gas-azure-databricks;
https://blogs.microsoft.com/ai/shell-iot-ai-safety-intelligent-tools/;

[59] https://www.wired.com/story/walmart-shoplifting-artificial-intelligence-everseen/;
https://techsnow.xyz/walmart-using-artificial-intelligence/;

[60] https://www.cmple.com/learn/smart-home-artificial-intelligence-nest-was-just-the-start

[61] https://www.ge.com/research/technology-domains/artificial-intelligence/computer-vision

[62] https://emerj.com/ai-sector-overviews/ai-at-disney-viacom-and-other-entertainment-giants/;
https://www.7wdata.be/big-data/how-disney-is-using-artificial-intelligence-to-figure-out-exactlyhowmuchyouenjoyitsfilms/;
https://www.youtube.com/watch?v=tfq1bCx9pnU;
https://relate.zendesk.com/articles/ai-is-what-dreams-are-made-of-ask-disney/;
https://www.cbc.ca/news/technology/disney-ai-real-time-tracking-fvae-1.4233063;

[63] https://www.forbes.com/sites/bernardmarr/2019/08/23/the-amazing-ways-youtube-uses-artificial-intelligence-and-machine-learning/#2a499c275852; https://www.weforum.org/agenda/2018/04/ai-is-now-youtube-s-biggest-weapon-against-the-spread-of-offensive-videos/;
https://www.theverge.com/2020/9/21/21448916/youtube-automated-moderation-ai-machine-learning-increased-errors-takedowns;

[64] https://emerj.com/ai-sector-overviews/machine-learning-big-box-retail/;

[65] https://www.uxconnections.com/how-loreal-uses-machine-learning-and-augmented-reality-to-disrupt-the-beauty-industry/;

[66] https://www.wired.com/story/walmart-shoplifting-artificial-intelligence-everseen/;
https://techsnow.xyz/walmart-using-artificial-intelligence/

[67] https://www.bbc.com/news/technology-51309186;
https://www.forbes.com/sites/bernardmarr/2016/12/29/4-amazing-ways-facebook-uses-deep-learning-to-learn-everything-about-you/#1d52a5d4ccbf;

[68] https://finadium.com/bnp-paribas-using-advanced-tech-to-automate-asset-servicing/

[69] https://emailtree.ai/news/emailtree-supplier-edf-group/

[70] https://www.nytimes.com/2019/10/22/business/mcdonalds-tech-artificial-intelligence-machine-learning-fast-food.html

[71] https://www.analyticsinsight.net/how-walmart-is-adopting-ai-technologies/

[72] https://www.prnewswire.com/news-releases/lufax-introduces-know-your-intention-300887949.html; https://picante.today/latest-news/2019/07/31/64517/lufax-vice-president-yunny-lee-envisi ons-a-self-improving-future-of-finance/

[73] https://www.bernardmarr.com/default.asp?contentID=1275; https://blog.hootsuite.com/how-the-youtube-algorithm-works/; https://modassicmarketing.com/google-rankbrain-ai-artificial-intelligence

[74] https://www.creativebenefitsinc.com/cvs-health-launches-anxiety-treatment-app/

[75] https://www.shrm.org/resourcesandtools/hr-topics/technology/pages/ibm-transforms-human-resources-ai.aspx

[76] https://www.forbes.com/sites/bernardmarr/2016/12/29/4-amazing-ways-facebook-uses-deep-learning-to-learn-everything-about-you/#1d52a5d4ccbf

[77] https://www.getfeedback.com/resources/cx/cx-takeaways-from-dreamforce-2019/

[78] https://digital.hbs.edu/platform-rctom/submission/machine-learning-at-disney-solving-happiness/

[79] https://www.volkswagenag.com/en/news/stories/2018/06/self-learning-factory.html#

[80] https://roboticsandautomationnews.com/2020/06/04/mitsubishi-electric-develops-co-ope rative-ai-for-human-machine-work/32787/

[81] https://www.automationworld.com/factory/robotics/article/21138274/bmw-outfits-robots-with-artificial-intelligence

[82] https://builtin.com/artificial-intelligence/ai-in-banking; https://finovate.com/pepper-power-hsbc-brings-robot-retail-banking-to-finovatefall/; https://www.cbc.ca/news/canada/toronto/hsbc-introduces-pepper-the-robot-promising-fun-efficiency-and-job-growth-1.5261812

[83] https://www.worldoil.com/news/2018/4/2/total-partnership-to-deploy-worlds-first-autonomous-offshore-robot

[84] https://www.sony.net/SonyInfo/sony_ai/

[85] https://research.samsung.com/robot; https://news.samsung.com/global/get-a-glimpse-of-the-next-generation-innovations-on-display-at-samsungs-technology-showcase; https://news.samsung.com/global/video-meet-the-samsung-bots-your-companions-of-the-future;

[86] https://aibusiness.com/document.asp?doc_id=761304

[87] https://transmetrics.eu/blog/best-artificial-intelligence-logistics-startups/

[88] https://pubs.spe.org/en/ogf/ogf-article-detail/?art=5052; https://www.marinetechnologynews.com/news/group-flatfish-543315; https://www.saipem.com/sites/default/files/2019-03/FlatFish%20Phase%203_it_09012019.pdf

[89] https://www.fastcompany.com/90246028/how-ai-is-helping-amazon-become-a-trillion-dollar-company

[90] https://www.theverge.com/2017/9/27/16373058/toyota-research-institute-self-driving-car-luminar; https://europe.autonews.com/automakers/lexus-prepares-introduce-level-2-autonomy; https://www.theverge.com/2019/10/24/20930227/toyota-self-driving-car-taxi-tokyo-test-2020-olympics;

[91] https://www.axora.com/insights/artificial-intelligence-five-ways-ai-is-disrupting-the-mining-industry/; https://www.cnbc.com/2018/06/25/from-wearables-to-autonomous-drills-heres-how-innovation-is-transfor.html

[92] https://www.globalbrandsmagazine.com/intel-and-tencent-youtu-team-on-artificial-intelligence-products-to-deliver-customer-insights/

[93] https://www.sensory.com/sensory-unveils-domain-specific-voice-assistant-for-smart-appliances/

[94] https://www.healthcarefinancenews.com/news/study-finds-racial-bias-optum-algorithm; https://quertle.com/healthcare-ai-gone-bad/; https://www.boozallen.com/c/insight/blog/ai-bias-in-healthcare.html;

[95] https://builtin.com/artificial-intelligence/ai-in-banking; https://finovate.com/pepper-power-hsbc-brings-robot-retail-banking-to-finovatefall/ ; https://www.cbc.ca/news/canada/toronto/hsbc-introduces-pepper-the-robot-promising-fun-efficiency-and-job-growth-1.5261812

[96] https://www.cnet.com/news/samsung-neon-project-finally-unveiled-humanoid-ai-chatbot-artificial-humans/; https://thenextweb.com/artificial-intelligence/2020/01/07/samsung-is-working-on-lifelike-ai-powered-avatars-to-fill-in-for-humans/;

[97] https://www.fooddive.com/news/nestles-personalized-nutrition-pilot-taps-ai-uses-consumer-dna/531560/; https://twitter.com/quicktake/status/1035010038560116741;

[98] http://thecloudbigdata.com/how-coca-cola-uses-big-data/;

[99] https://digiday.com/marketing/artificial-intelligence-influencing-unilevers-marketing/; https://www.campaignlive.co.uk/article/ai-helped-unilever-discover-breakfast-dessert/1522441?zd_source=hrt&zd_campaign=13370&zd_term=chiradeepbasumallick;

[100] https://www.swisslog.com/en-us/newsroom/news-press-releases-blog-posts/2019/11/nestle-distribution-center-of-the-future; https://www.youtube.com/watch?v=nY7xnMxvGgw&feature=emb_logo; https://www.supplychaindive.com/news/nestle-and-xpo-prepare-a-warehouse-of-the-future/525941/

[101] https://www.pensasystems.com/wp-content/uploads/2019/03/Pensa_OnePager_v3.pdf; https://www.beveragedaily.com/Article/2019/01/14/AB-InBev-trials-drones-advanced-computer-vision-and-AI-in-retail; https://www.therobotreport.com/pensa-systems-raises-10m-drone-inventory-ai/;

[102] https://www.forbes.com/sites/bernardmarr/2019/04/05/the-fascinating-ways-pepsico-uses-artificial-intelligence-and-machine-learning-to-deliver-success/#147a85cb311e;

[103] https://news.microsoft.com/transform/from-seed-to-sip-how-anheuser-busch-inbev-uses-ai-to-drive-growth/; https://thirdeyedata.io/anheuser-busch-inbev-brews-up-game-changing-business-solutions-with-microsoft-azure/; https://news.microsoft.com/transform/videos/cheers-anheuser-busch-

inbev-taps-data-for-even-better-beer/; https://customers.microsoft.com/en-us/story/ab-inbev-consumer-goods-azure;

[104] https://www.just-food.com/futures/downloads/FUTURES-Vol8_Machine-Dreams.pdf; https://www.ibm.com/watson/stories/coreml;

[105]https://bernardmarr.com/default.asp?contentID=1171#:~:text=Social%20data%20mining&text=The%20company%20has%20used%20AI,way%20to%20serve%20them%20advertisements.; https://www.bernardmarr.com/default.asp?contentID=1280

[106] https://www.forbes.com/sites/bernardmarr/2019/04/05/the-fascinating-ways-pepsico-uses-artificial-intelligence-and-machine-learning-to-deliver-success/?sh=78e9e81b311e

[107] https://www.bcg.com/it-it/publications/2018/interview-tassilo-festetics-anheuser-busch-inbev-artificial-intelligence-is-everywhere; https://bernardmarr.com/default.asp?contentID=1171#:~:text=Social%20data%20mining&text=The%20company%20has%20used%20AI,way%20to%20serve%20them%20advertisements

[108] https://www.btimesonline.com/articles/133144/20200603/baidu-working-with-midea-to-produce-ai-powered-smart-home-appliances.htm;

[109] https://enterpriseiotinsights.com/20180925/channels/news/caterpillar-and-osisoft-debut-ai-platform

[110] https://syncedreview.com/2019/07/06/ai-x-global-500-manufacturing-companies-robots-data-maintenance/

[111] https://spectrum.ieee.org/view-from-the-valley/robotics/artificial-intelligence/want-a-really-hard-machine-learning-problem-try-agriculture-say-john-deere-labs-leaders; https://www.forbes.com/sites/bernardmarr/2018/03/09/the-incredible-ways-john-deere-is-using-artificial-intelligence-to-transform-farming/#27293590330d

[112] https://digital.hbs.edu/platform-digit/submission/john-deere-using-ai-to-reimagine-farming/

[113] https://www.wired.com/wiredinsider/2019/11/machine-learning-manufacturing-giant-3m-innovates-products-old-new/

[114] https://www.mitsubishielectric.com/en/about/rd/research/highlights/automation/ai.html; https://www.biospace.com/article/releases/mitsubishi-electric-s-kotsumon-system-uses-ai-video-technology-to-analyze-production-line-workers-motions/; http://blog.999automation.com/2017/07/10/mitsubishi-compact-ai-algorithm-slashes-machine-setup-times-2/

[115] https://roboticsandautomationnews.com/2020/06/04/mitsubishi-electric-develops-co-operative-ai-for-human-machine-work/32787/

[116] https://www.ft.com/content/82fdd5fa-be73-11e9-9381-78bab8a70848; http://www.junlinghu.com/2019/05/21/ai-in-manufacturing-industrial-ai/

[117] https://www.japantimes.co.jp/news/2018/12/17/business/tech/seven-eleven-tries-store-tokyo-unmanned-payment-system/

[118] https://etailwest.wbresearch.com/jet-black-case-study-ty-u

[119] https://arekskuza.com/the-innovation-blog/amazon-and-artificial-intelligence-in-retail/

[120] https://blogs.nvidia.com/blog/2019/07/11/walmart-nvidia/

[121]https://customers.microsoft.com/pt-br/story/733091-walgreens-boots-alliance-pharmaceuticals-azure

[122] https://venturebeat.com/2020/05/11/how-walmart-uses-ai-to-enable-two-hour-express-delivery/

[123]https://www.supermarketnews.com/online-retail/amazon-tests-autonomous-vehicle-deliveries

[124] https://www.fastcompany.com/90246028/how-ai-is-helping-amazon-become-a-trillion-dollar-company

[125] https://corporate.walmart.com/newsroom/2019/04/25/walmarts-new-intelligent-retail-lab-shows-a-glimpse-into-the-future-of-retail-irl

[126] https://www.pocket-lint.com/phones/news/amazon/139650-what-is-amazon-go-where-is-it-and-how-does-it-work; https://towardsdatascience.com/how-the-amazon-go-store-works-a-deep-dive-3fde9d9939e9;

[127]https://www.wired.com/story/walmart-shoplifting-artificial-intelligence-everseen/; https://techsnow.xyz/walmart-using-artificial-intelligence/

[128] https://risnews.com/ris-exclusive-how-home-depot-employing-prescriptive-analytics

[129] https://www.chinadaily.com.cn/regional/chongqing/liangjiang/2017-08/17/content_30770287.htm; http://www.chinadaily.com.cn/regional/chongqing/liangjiang/2018-09/28/content_37003425.htm;

[130] https://emerj.com/ai-sector-overviews/artificial-intelligence-goldman-sachs/#:~:text=AI%2DEnabled%20Investment%20Trust,-Goldman%20Sachs%20tested&text=Goldman%20Sachs%20claims%20the%20software,based%20on%20historical%20investment%20data.

[131] http://www.chinabankingnews.com/2018/10/29/artificial-intelligence-will-transform-traditional-banking-china-merchants-bank/

[132] https://www.forbes.com/sites/tomdavenport/2019/07/16/how-td-ameritrade-uses-ai-to-hear-the-voice-of-the-customer/#4b9c166f6349

[133] https://emerj.com/ai-sector-overviews/artificial-intelligence-goldman-sachs/#:~:text=AI%2DEnabled%20Investment%20Trust,-Goldman%20Sachs%20tested&text=Goldman%20Sachs%20claims%20the%20software,based%20on%20historical%20investment%20data.

[134]https://www.bi-kring.nl/nieuws/data-science-archief/99-data-science-archief/1008-an-overview-of-morgan-stanley-s-surge-toward-data-quality

[135] https://financialpost.com/sponsored/business-sponsored/how-artificial-intelligence-can-leverage-big-data-to-boost-trade