

**dSEA**

**Università degli Studi di Padova**

DIPARTIMENTO DI SCIENZE ECONOMICHE E AZIENDALI "MARCO FANNO"

Master of Business Administration

# **The Impact of Artificial Intelligence on Strategic and Operational Decision Making**

**Supervisor**


Prof. Ivan De Noni

**Candidate**

Giacomo Coletto

2036091

Academic Year 2022–2023



**dSEA**

**Università degli Studi di Padova**

---

DIPARTIMENTO DI SCIENZE ECONOMICHE E AZIENDALI "MARCO FANNO"

Master of Business Administration

**The Impact of Artificial Intelligence on Strategic  
and Operational Decision Making**

**Supervisor**

Prof. Ivan De Noni

**Candidate**

Giacomo Coletto

2036091

Academic Year 2022–2023

## DECLARATION OF AUTHENTICITY

Dichiaro di aver preso visione del “Regolamento antiplagio” approvato dal Consiglio del Dipartimento di Scienze Economiche e Aziendali e, consapevole delle conseguenze derivanti da dichiarazioni mendaci, dichiaro che il presente lavoro non è già stato sottoposto, in tutto o in parte, per il conseguimento di un titolo accademico in altre Università italiane o straniere. Dichiaro inoltre che tutte le fonti utilizzate per la realizzazione del presente lavoro, inclusi i materiali digitali, sono state correttamente citate nel corpo del testo e nella sezione **Bibliography**.

I hereby declare that I have read and understood the “Anti-plagiarism rules and regulations” approved by the Council of the Department of Economics and Management and I am aware of the consequences of making false statements. I declare that this piece of work has not been previously submitted — either fully or partially — for fulfilling the requirements of an academic degree, whether in Italy or abroad. Furthermore, I declare that the references used for this work — including the digital materials — have been appropriately cited and acknowledged in the text and in the section **Bibliography**.

A handwritten signature in black ink, reading "Giacomo Coletto". The signature is written in a cursive style with a prominent flourish at the end.

*A mio padre,  
l'eroe che ha sempre creduto in me.  
Perché la sua saggezza e il suo supporto  
mi hanno portato dove non avrei mai immaginato.*

## ABSTRACT

Effective decision making lies at the core of organizational success. In the era of digital transformation, businesses are increasingly adopting data-driven approaches to gain a competitive advantage. According to existing literature, Artificial Intelligence (AI) represents a significant advancement in this area, with the ability to analyze large volumes of data, identify patterns, make accurate predictions, and provide decision support to organizations. This study aims to explore the impact of AI technologies on different levels of organizational decision making. By separating these decisions into strategic and operational according to their properties, the study provides a more comprehensive understanding of the feasibility, current adoption rates, and barriers hindering AI implementation in organizational decision making.

**Keywords:** decision making, artificial intelligence, strategy, bibliometric analysis

## ACKNOWLEDGMENTS

I want to express my gratitude to the following people who have supported me throughout the elaboration of this Master's thesis.

First and foremost, I am deeply grateful to my supervisor, Prof. Ivan De Noni, for his guidance, availability, and constant encouragement. I am thankful for his refining of my original thesis proposal, for introducing me to bibliometric analysis, for providing feedback on the research structure and interview questions, and for motivating me with the potential of a publication.

I want to express my recognition to Prof. Sébastien Douaillat for his expertise and teaching methods in Business Intelligence, as they inspired me to pursue this research interest. I am also thankful to him for providing me with two managers for my interviews.

I extend my heartfelt gratitude to the managers who volunteered their time and shared their experiences and insights for this research. Their contributions were fundamental in shaping the findings and conclusions of this thesis.

My thankfulness goes also to my colleagues at Accenture, for welcoming me in their team and understanding the importance of this research project for me. I am particularly grateful to Andrea, who gave me the chance to prove my abilities and connected me to AI experts.

I want to also thank my family for their support and motivation throughout this journey, and in particular Nicola and Martin for connecting me to different managers.

Lastly, but also most importantly, I owe my biggest thank you to Fabiola. She has always been on my side, listening to my ideas and appreciating what I was writing even before I did. Helping me countless times, being my constant motivator, and wanting me to succeed, are just few of the things that remind me every day of what it means to truly love someone.

# CONTENTS

<b>Abstract</b>	<b>iv</b>
<b>Acknowledgments</b>	<b>v</b>
<b>Contents</b>	<b>vi</b>
<b>List of Figures</b>	<b>viii</b>
<b>List of Tables</b>	<b>ix</b>
<b>List of Acronyms</b>	<b>x</b>
<b>1 Introduction</b>	<b>1</b>
1.1 Decision making . . . . .	2
1.1.1 Link with organizational structure . . . . .	2
1.1.2 Data-driven decision making . . . . .	5
1.1.3 Research interest . . . . .	6
1.2 Details of research process . . . . .	7
<b>2 Literature review</b>	<b>11</b>
2.1 Introduction to bibliometric analysis . . . . .	12
2.1.1 Principles and techniques . . . . .	13
2.1.2 Document selection process . . . . .	14
2.2 Aim and scope of literature review . . . . .	16
2.3 Broad search analysis . . . . .	17
2.3.1 Trending keywords . . . . .	19
2.3.2 Thematic evolution . . . . .	21
2.4 Introduction to Artificial Intelligence . . . . .	25
2.4.1 Evolution of the term . . . . .	25
2.4.2 Focus of current research . . . . .	26
2.4.3 Trends of current business adoption . . . . .	27
2.5 Narrow search analysis . . . . .	29
2.5.1 Co-citation analysis . . . . .	30

2.5.2	Bibliographic coupling . . . . .	44
2.6	Discussion and gap analysis . . . . .	56
2.7	Research question . . . . .	59
<b>3</b>	<b>Real cases of AI employment in decision making</b>	<b>61</b>
3.1	Research methodology . . . . .	61
3.1.1	Interview questions . . . . .	62
3.1.2	Companies selection . . . . .	63
3.1.3	List of selected companies . . . . .	64
3.2	Interviews results . . . . .	67
3.2.1	Current AI employment . . . . .	67
3.2.2	Drivers of AI employment . . . . .	71
3.2.3	Barriers to AI employment . . . . .	76
3.2.4	Considerations on AI employment . . . . .	82
3.3	Analysis and discussion of results . . . . .	88
3.3.1	Current AI employment . . . . .	88
3.3.2	Drivers of AI employment . . . . .	89
3.3.3	Barriers to AI employment . . . . .	90
3.3.4	Considerations on AI employment . . . . .	91
3.3.5	Answer to research question . . . . .	92
3.4	Conclusions . . . . .	93
	<b>Bibliography</b>	<b>95</b>



# LIST OF FIGURES

1.1	Mintzberg’s five basic parts of organizations . . . . .	3
1.2	Mintzberg’s structure of the decision process . . . . .	4
1.3	Structure of the theoretical part of the research process . . . . .	9
1.4	Structure of the empirical part of the research process . . . . .	10
2.1	Document selection process in bibliometric analysis . . . . .	14
2.2	Trending topics in broad search documents . . . . .	20
2.3	Thematic evolution in broad search documents . . . . .	22
2.4	Thematic networks in broad search documents in 2020–2023 . . . . .	24
2.5	Keyword co-occurrence of AI publications in 1998–2017 . . . . .	27
2.6	S&P 500 companies mentioning AI in their earning calls in 2013–2023 . . . . .	28
2.7	Most mentioned keywords in earning calls in Q1 2023 compared to Q4 2022 . . . . .	29
2.8	Co-citation analysis of narrow search references . . . . .	31
2.9	Framework to use Big Data in a decision process . . . . .	34
2.10	Explanation of object recognition through the LIME technique . . . . .	38
2.11	Bibliographic coupling of narrow search documents . . . . .	45
2.12	Jitter plot of the most important features in predicting back orders . . . . .	50
2.13	Architecture of a Monopoly-playing Deep Learning agent using both a rule-based and an autonomous learning approach . . . . .	57
3.1	Sankey diagram of companies invited to the interview . . . . .	65

## LIST OF TABLES

2.1	Elements of scientific research networks in bibliometric analysis . . . . .	14
2.2	Properties of search results for broad and narrow searches . . . . .	18
2.3	Properties of clusters in co-citation analysis . . . . .	32
2.4	Possible audience of eXplainable Artificial Intelligence . . . . .	38
2.5	Properties of clusters in bibliographic coupling . . . . .	46
2.6	Benefits and drawbacks of AI adoption in a business context . . . . .	52
2.7	Application areas, drivers, outcomes and technologies of AI in B2B marketing	54
2.8	Properties of strategic and operational decisions . . . . .	59
3.1	Properties of interviewed companies . . . . .	66
3.2	Current AI employment in decision making . . . . .	71
3.3	Drivers of AI employment in decision making . . . . .	75
3.4	Barriers to AI employment in decision making . . . . .	81
3.5	Considerations on AI employment in decision making . . . . .	87

## LIST OF ACRONYMS

<b>AI</b>	Artificial Intelligence . . . . .	1
<b>ANN</b>	Artificial Neural Networks . . . . .	27
<b>B2B</b>	Business to Business . . . . .	54
<b>BD</b>	Big Data . . . . .	33
<b>BI</b>	Business Intelligence . . . . .	19
<b>CFO</b>	Chief Financial Officer . . . . .	73
<b>COO</b>	Chief Operating Officer . . . . .	65
<b>CoT</b>	Chain of Thought . . . . .	85
<b>CRM</b>	Customer Relationship Management . . . . .	68
<b>CV</b>	Computer Vision . . . . .	27
<b>DL</b>	Deep Learning . . . . .	29
<b>DSS</b>	Decision Support Systems . . . . .	19
<b>ES</b>	Expert Systems . . . . .	19
<b>FOSS</b>	Free Open-Source Software . . . . .	12
<b>GAAM</b>	Gartner Analytics Ascendancy Model . . . . .	76
<b>GAN</b>	Generative Adversarial Network . . . . .	49
<b>GenAI</b>	Generative Artificial Intelligence . . . . .	27
<b>HRM</b>	Human Resource Management . . . . .	46
<b>IEEE</b>	Institute of Electrical and Electronics Engineers . . . . .	35
<b>LIME</b>	Local Interpretable Model-agnostic Explanations . . . . .	37
<b>LLM</b>	Large Language Models . . . . .	25
<b>LSTM</b>	Long Short-Term Memory . . . . .	49
<b>MCDM</b>	Multi-Criteria Decision Making . . . . .	35
<b>ML</b>	Machine Learning . . . . .	26
<b>N/A</b>	Not Applicable . . . . .	62
<b>NLP</b>	Natural Language Processing . . . . .	12
<b>OLAP</b>	Online Analytical Processing . . . . .	33
<b>RDB</b>	Relational Database . . . . .	19
<b>ROI</b>	Return on Investment . . . . .	53

## LIST OF ACRONYMS

---

<b>RQ</b>	Research Questions . . . . .	8
<b>SCM</b>	Supply Chain Management . . . . .	21
<b>SHAP</b>	SHapley Additive exPlanations . . . . .	49
<b>SMEs</b>	Small and Medium Enterprises . . . . .	53
<b>SQL</b>	Structured Query Language . . . . .	19
<b>XAI</b>	eXplainable Artificial Intelligence . . . . .	37

# 1

## INTRODUCTION

**T**HE research area of this master's thesis explores one of the topics that lead me towards this academic path. Five years ago, I was fascinated by how leaders of large corporations could take decisions that impacted not only their companies, but also society. Today, I want to express my gratitude for this learning journey by presenting my original contribution to research.

On a broad level, this thesis is an experimental research on the impact of Artificial Intelligence (AI) on the decisions of companies. The research process includes a comprehensive literature review and an empirical analysis based on primary qualitative data. Bibliometric analysis was chosen as the literature review methodology, and semi-structured interviews with manufacturing companies as the data collection methodology.

The research outline is presented as follows. We start our research by introducing the importance of decision making in organizations, and then by illustrating the details of the research design. Unlike other studies, we decide to not formulate a research question at the beginning of the research, opting instead for a more exploratory approach.

The second chapter begins with the introduction of bibliometric analysis, a powerful literature review methodology. During our literature review, we conduct two bibliometric analyses. The first analysis aims to understand the long-term developments of the study field, and is performed on a wide list of publications on *data-driven decision making*. From this analysis, we identify a recent research area of interest, which is the impact of *artificial intelligence in decision making*. We conduct a second bibliometric analysis to find a gap in the literature. This is achieved by understanding the current thematic organization of the research field, and the themes that influenced it. We find that the literature does not clearly identify what kind of business decisions can benefit the most from AI. We elaborate from the analyzed literature a framework to categorize business decisions into strategic and operational according to their properties. We then formulate a research question that aims

to answer if AI is more suited to strategic or more operational decisions.

To answer this research question, we derive interview questions on the current employment of AI in companies, on the main drivers behind these adoptions, on the main barriers hindering new use cases, and on possible actions to overcome them. We plan to collect data by hosting interviews with executives and data-savvy senior managers of large manufacturing companies. Answers to interview questions are organized, analyzed and discussed according to their topic. The insights collected allow us to answer to our research question and provide recommendations to scholars and AI practitioners. We conclude our research acknowledging its limitations and suggesting possible avenues for future research.

### 1.1 DECISION MAKING

**T**AKING decisions and being responsible for their outcomes is a crucial aspect of a manager's role. Managers must have the ability to make informed choices and be accountable for the consequences that arise from them. However, human decision makers are susceptible to their own biases and prejudices, whether consciously or unconsciously. Potentially, this can lead to suboptimal decisions and negative consequences for a company's performance (Leyer & Schneider, 2021).


Due to its role in organizations, its complex and challenging nature, combined with humans' bounded rationality, the process of decision-making has long been a topic of interest (Mintzberg, 1979; Simon, 1957).

In this section we introduce the concept of decision making in an organizational context, and our research interest in data-driven decision making. We elaborate from the research of Henry Mintzberg, as in one of his most renowned books he links the structure of organizations with different kinds of decisions.

#### 1.1.1 LINK WITH ORGANIZATIONAL STRUCTURE

In his book *The Structuring of Organizations: A Synthesis of the Research*, Mintzberg (1979) identifies five common components of organizational structures. These components vary in size and importance according to the nature of the organization (e.g. corporation vs. family business) figure 1.1.

Three of these five components are connected by a *single line of formal authority*:



**Figure 1.1:** Mintzberg's five basic parts of organizations


**Source:** Adapted from Mintzberg (1979)

- **Strategic apex:** leaders of the organization
- **Middle line:** managers of lower levels
- **Operating core:** workers directly involved with production or services

This central line of formal authority also represents the flows of control, operations, information and decisions within an organization. Operating performances are collected at every level and aggregated for structured reporting in an upward flow, providing what Mintzberg calls a Management Information System. However, this system cannot by its nature include at the right time all the information needed to take decisions (Rajagopal et al., 2022).

According to (Mintzberg et al., 1976), a decision is a process that spans from the perception of a need to the commitment to the action. As seen in figure 1.2, this process is divided in three phases, that involve different steps and feedback loops:

1. **Identification:** the problem or opportunity is recognized, and then a diagnosis of the problem is made
2. **Development:** there is the search for ready-made solutions, or the design of tailor-made solutions
3. **Selection:** the choice of the solution is made either through personal judgment, analysis of alternatives, or bargaining in case of group decisions. If needed, authorization to proceed is granted at this stage



**Figure 1.2:** Mintzberg's structure of the decision process

**Source:** Adapted from Mintzberg et al. (1976)

Empirical analyses illustrated that decisions differ according to their organizational level (Martin, 1956). In particular,

“At each successively lower level, the decisions were more frequent, of shorter duration, and less elastic, ambiguous, and abstract; solutions tended to be more predetermined; the significance of events and interrelationships was more clear; in general, lower-level decision making was more structured.”

This reflection of the type of work, which is less repetitive and thus less formalized higher in the hierarchy.

Mintzberg (1979) argues that there is no standard classification of decisions, but rather a set of conceptual criteria for which they can be separated, such as structure (programmed or not), frequency (routine or one-off), functional area (product, investment, hiring), process (intuitive or analytical) and importance (impactful or not). The author illustrates one of the most common classifications based on importance, where decisions are divided in strategic, administrative and operating. This framework derives from military operations, where decisions are often organized in strategic, tactical and operational (Anthony, 1965).

According to Mintzberg, business decisions can be:



- **Strategic:** happen exceptionally, have a significant impact on organizations, involve a broad perspective (e.g. investments, acquisitions)
- **Administrative:** happen routinely as part of unstructured processes, have a medium impact, (e.g. coordination, planning and budgeting)
- **Operative:** happen routinely as part of structured processes, have a limited impact, have predetermined phases with little diagnosis and ready-made solutions (e.g. worker starting a machine, librarian searching for a reference)

The author highlights how strategic importance is relative to each business case, mentioning that pricing can be crucial for a large construction contract, but not as crucial for a small restaurant. Considering that decisions may present different features is a necessary first step towards better decision making.

### 1.1.2 DATA-DRIVEN DECISION MAKING

Around fifty to seventy years ago, business decision making heavily relied on human judgment. Professionals relied significantly on their intuitions, which had been established over years of experience and limited data in their particular disciplines. Whether it was determining optimal inventory levels or approving financial investments, relying on experience and gut instinct was the primary way to differentiate between good or bad, high or low, risky or safe business decisions (Siegel et al., 2020).

However, our intuitions are not ideal decision making instruments due to cognitive biases resulting from evolution. Early hunter-gatherers developed a system of reasoning based on simple heuristics, enabling quick, almost unconscious decisions in dangerous situations. Unfortunately, this does not always lead to optimal or accurate outcomes (Siegel et al., 2020).

Being in an interconnected world, today many companies employ a data-driven approach to make operational decisions more efficiently (Siegel et al., 2020). In a data-driven workflow, human judgment remains as the central processor, but elaborates summarized data as a new input. Despite the promises, this approach has limitations, such as not leveraging all the data available or needed, or bias in data or its aggregation. Ronald Coase explained the tendency of economists and managers to get the results they expect: “If you torture the data long enough, it will confess”.

Even if automation of business decisions may sound appealing at first, Rajagopal et al. (2022) acknowledge the limitations of such an approach. First, business decisions

are taken based on many more factors than just structured data. Long-term objectives, business strategies, organizational values, and competitive dynamics are all examples of knowledge that is confined to our minds and transmitted through non-traditional, non-digital channels. This information, although vital for business decisions, is inaccessible to information systems.

For instance, even hypothesizing an intelligent system able to decide that the optimal decision is to decrease the inventory level in a warehouse, company executives may decide to increase inventory levels everywhere to enhance the customer experience, even if it affects profitability in the long term. The extensive knowledge possessed by individuals regarding tactics, ethics, and economic circumstances may cause their decisions to diverge from the objective rationality of intelligent information systems. To address this, these systems can be utilized to generate possible alternatives for individuals to choose from, based on additional information Rajagopal et al. (2022).

### 1.1.3 RESEARCH INTEREST

Despite the potential benefits that prescriptive systems like these can bring to companies, we do not hear about recommender systems for managers as much as we do for e-commerce customers. Exploring the scientific literature on data-driven decision making can shed light on recent developments in the field, potentially explaining the implementation process of such systems.

However, the lack of scientific research on the field (900 documents found in Scopus for *prescriptive analytics* compared to 32 000 of *predictive analytics*) and the amount of references to Gartner shed some doubts on whether this concept originates from a scientific background or as the client offering of a private company (Maoz, 2013).

Therefore, we decide to conduct a broad and exploratory literature review on *data-driven decision making*, as described in section 1.2.

### 1.2 DETAILS OF RESEARCH PROCESS

**D**ESIGNING an effective research methodology is crucial for the accuracy and reliability of the results. In this section, we illustrate the process that we plan for this research. The thesis follows a structured research process, consisting of an initial brainstorming on a topic of interest, a comprehensive literature review, the formulation of a research question, the collection of primary qualitative data, and the discussion of the results obtained. Figures 1.3 and 1.4 illustrate this process as two flowcharts, respectively for the theoretical and empirical part.

After an initial brainstorming, my supervisor and I decided to perform an exploratory literature analysis on the topic of *data-driven decision making*. The bibliometric approach was chosen as the most suitable methodology for this kind of analysis, as it can illustrate the main themes of a scientific research field by analyzing the metadata of a broad amount of documents.

At this stage, no research question was formulated, as a clear and novel research proposal would be made after an extensive analysis of the literature. We started our research process with the intent of discovering the impact of technologies in decision making, and only then choosing a recent theme to analyze.

By selecting a wide publication time span we warranted a large and variegated literature corpus to analyze through the bibliometric approach. We name this process *broad search*. Two types of keyword-based analyses were conducted to understand the evolution of the broad research field. The analysis of trending keywords identified historical research trends, while the thematic evolution analysis displayed how these themes have been studied together or diverged over the years.

Based on these analyses, we decided to focus the research on the development of *artificial intelligence in decision making* over the past three years, by performing a second bibliometric analysis, that we name *narrow search*. Before doing so, we briefly introduce the concept of AI, including an historical evolution of the term and the modern focus of the research field. For this analysis, we exported a second list of documents along with their reference lists.

Two types of citation-based analyses were then conducted to understand the current organization of the narrow research field. The co-citation analysis grouped the references of the exported documents based on the documents that cited them. These clusters represent the past research themes that have influenced the current literature. On the other hand, the

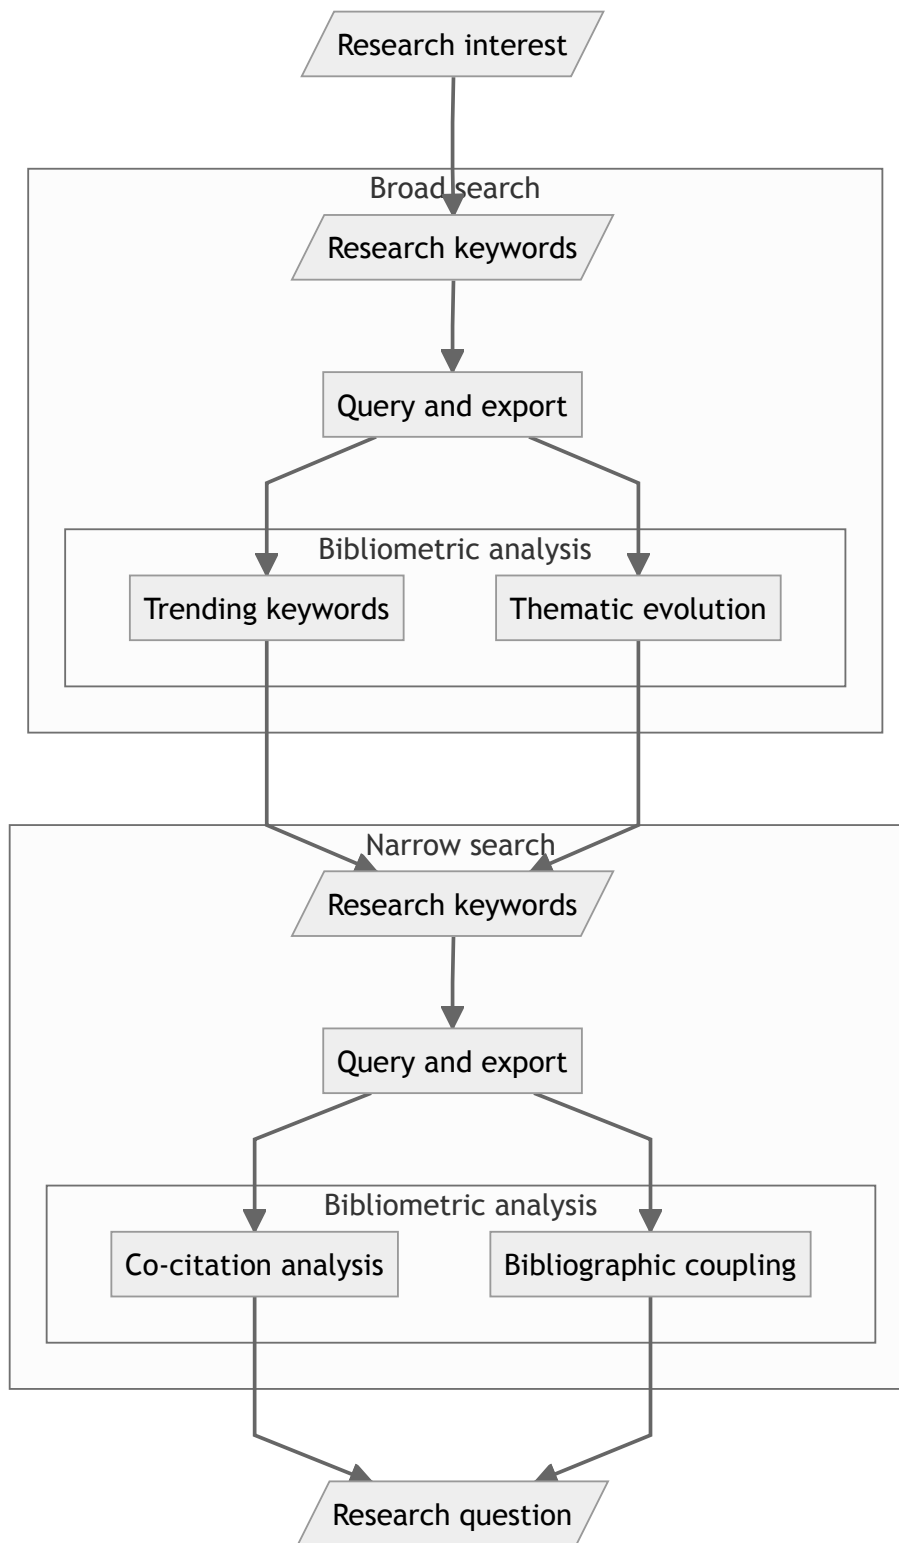
bibliographic coupling analysis grouped the exported documents based on their references, providing insight on how the current literature is organized.

While many articles stated that AI can support decision making in organizations, few explored this topic in depth, and to the best of our knowledge no one explained what kind of business decisions can benefit the most from AI. After identifying this important gap in the literature, we defined the main objective of the research and formulated our research question by taking into account the properties of different business decisions that emerged from the literature. Our Research Questions (RQ) asks if AI is better suited for more strategic or more operational decisions.

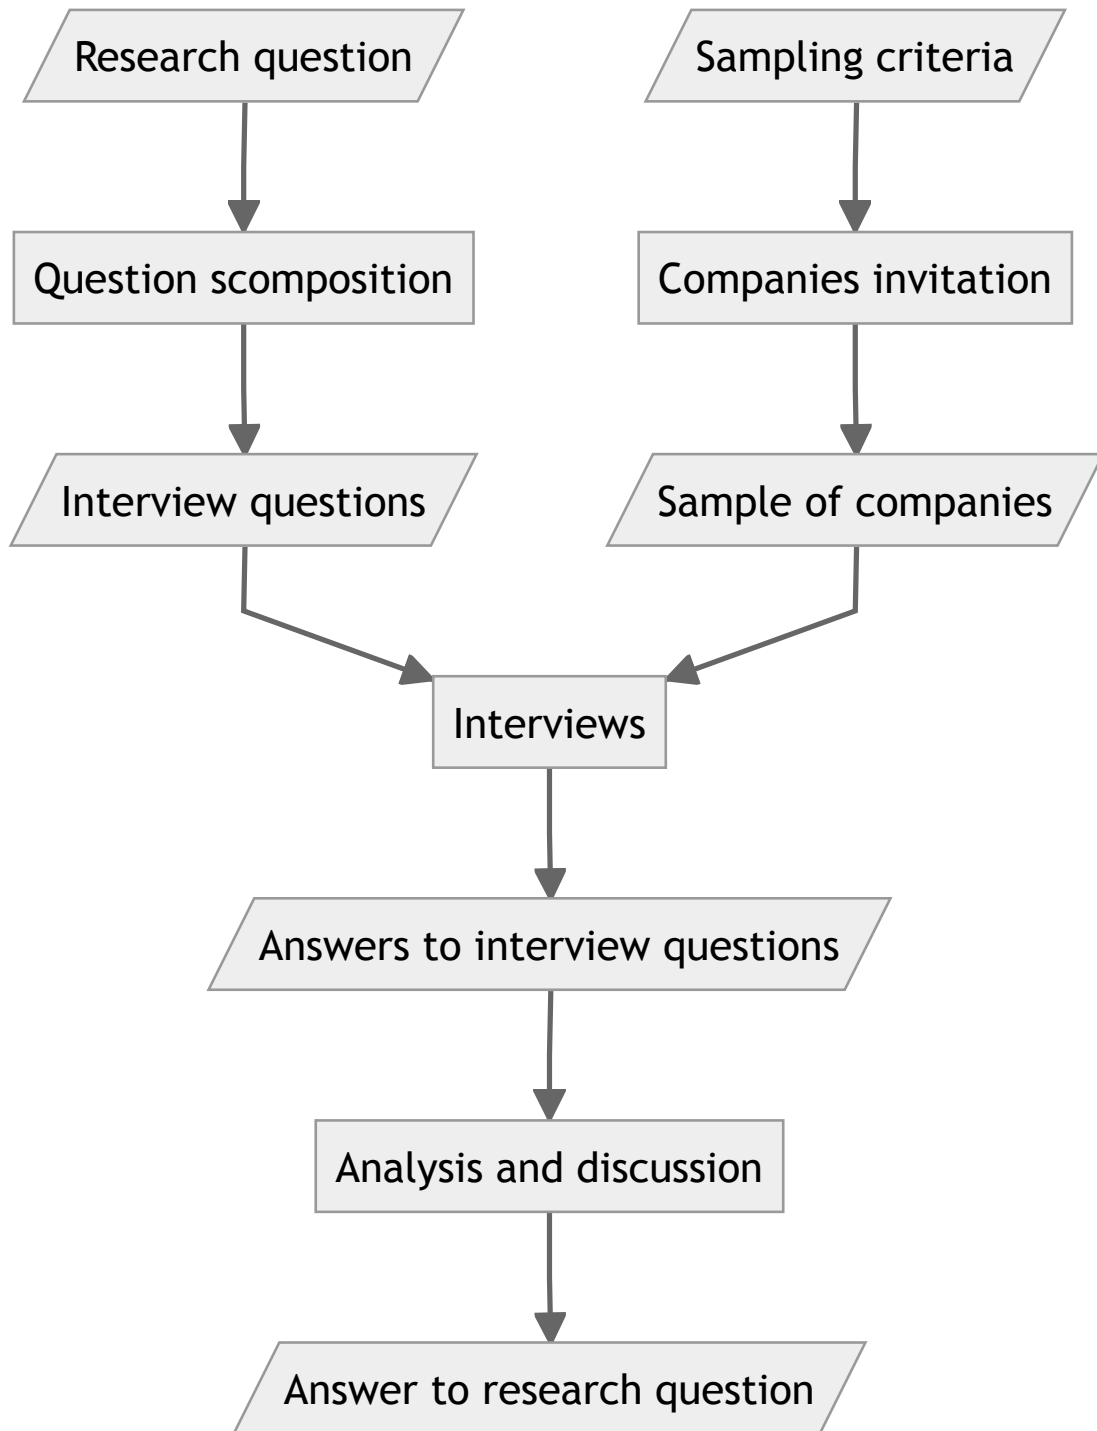
The empirical part of the research aims to understand the impact that AI is currently having on strategic and operational decisions in an organizational environment. We want to answer the RQ by collecting primary qualitative data from companies that are employing AI in their decisional processes. We aim to understand for what decisions they use AI, why they decided to employ AI, what is preventing them from employing AI to other kind of decisions, and whether these barriers can be overcome.

To provide more insights on the introduction of AI inside these companies, we planned to mostly interview companies whose core business is not AI or data-centered, but have introduced this technology only in later stage of their existence. The answers received during the interviews were organized by topic, to better compare the results. Similarly, these answers were analyzed and discussed by topic.

The final part of the research summarizes the main findings, acknowledges the limitations and suggests directions for further research.



**Figure 1.3:** Structure of the theoretical part of the research process



**Figure 1.4:** Structure of the empirical part of the research process

# 2

## LITERATURE REVIEW

**I**N this chapter we will conduct the literature review process illustrated in figure 1.3. After a brief introduction to the purpose of literature review and the possible methodologies to conduct it, section 2.1 will explain why bibliometric analysis was chosen, what are its principles and how its document selection process is structured. Section 2.2 will define the goal of the bibliometric analyses performed for this thesis., which spans multiple sections. Section 2.3 will cover the details of the broad bibliometric analysis, focusing on *data-driven decision making*. In this section it will emerge the concept of AI, which we will introduce in section 2.4. Section 2.5 will instead cover the details of the narrow bibliometric analysis focusing on *artificial intelligence in decision making*. Section 2.6 will discuss the results of the two analyses, and in section 2.7 a research question based on the findings will be presented.

Literature review is the process of examining, evaluating, and synthesizing existing literature or previous research on a particular topic. It is a crucial step in the research process because it provides background information, identifies knowledge gaps in previous research, and highlight unanswered research question. It can be carried out through qualitative, quantitative or combined methodologies, such as systematic literature review, meta-analysis and bibliometric analysis (Donthu et al., 2021).

Systematic literature review is usually carried out manually and is thus better suited for confined research areas. It relies on qualitative techniques to summarize and synthesize the findings of existing literature on a research topic.

Meta-analysis is a series of quantitative techniques used to summarize empirical evidence by analyzing the strength of effects and relationships among variables in a wide number of homogeneous studies in the field. It throws light on mixed empirical findings and boundary conditions of the studies.

Bibliometric analysis is used to gain a holistic view of a research field. Through this

analysis, scholars can understand in a short amount of time the intellectual structure (themes, authors, connections) of the field and its evolution, measure the impact of research and discover its emerging trends.

Advancements in Natural Language Processing (NLP), a series of technologies that make machines understand, interpret, and generate human language, enabled the creation of new tools to help researchers with literature review, ranging from brainstorming to finding and summarizing work around a research question, to visualizing the relationships between specific documents (Quinn, 2023). However, ethical concerns about algorithm choice and transparency are emerging, as these algorithms influence the results displayed, and can potentially influence what type of articles will be written to accommodate these metrics (Gadd, 2020).

Each of these literature review methodologies offers unique advantages and is better suited for a specific type and scope of literature review.

### 2.1 INTRODUCTION TO BIBLIOMETRIC ANALYSIS

**B**IBLIOMETRIC analysis has been chosen as the primary methodology for the goal of this thesis. Its versatility can be used to illustrate the evolution of trends in data-driven decision making, to group publications in clusters, and to visualize the relationship between themes covered. This methodology has been adopted to perform large literature reviews in the decision-making and AI field (H. Chen et al., 2012; Loureiro et al., 2021; Pietronudo et al., 2022; Raza et al., 2023; Tang & Liao, 2021).

This is achieved by using software to analyze a vast amount of data from scientific bibliographic databases such as Scopus (<https://www.scopus.com/>) or Web of Science (<https://www.webofscience.com/>). The increasing adoption of this methodology can be attributed to the literature coverage of these databases, the availability of Free Open-Source Software (FOSS) bibliometric software like *Bibliometrix* (Aria & Cuccurullo, 2017) or *VOSviewer* (Van Eck & Waltman, 2010) and the growth of scientific research itself (Donthu et al., 2021).



### 2.1.1 PRINCIPLES AND TECHNIQUES

An important distinction must be made in the literature review terminology. The term *references* indicates older publications that have been cited by the current document, while *citations* denotes that the current document has been mentioned by newer publications. Bibliometric analysis is based on assumptions that are either performance-related or network-related:

- The number of publications of an author represents his productivity (publication score)
- The number of citations of an author or publication represents his influence (global citation score)
- The number of citations of an author or publication within a review corpus represents influence in the current discipline (local citation score)
- Publications that are frequently cited together are similar in their theme (co-citation analysis)
- Publications that share their references are similar in their content (bibliographic coupling)
- Words that frequently appear together have a thematic relationship (co-word analysis)
- Authors or institutions that frequently collaborate together represent a research cluster (co-authorship analysis)
- Publications that are frequently cited among highly cited publications are of high-quality, even if they have a low numbers of citations (PageRank score)

Based on these principles, the most significant bibliometric techniques are performance analysis and scientific mapping. The first accounts for the contributions of research constituents based on metrics such as the number of publications or citations, whereas the latter focuses on the relationships between research constituents, such as the use of similar keywords.

Despite the insights that a bibliometric analysis can provide, this methodology should not be used alone to assess the overall quality of publications, as it only relies on pure quantitative metrics. For example, an article may be highly cited for negative reasons, and researchers may try to artificially boost their citation score to appear more influential. A higher publication frequency of an author does not imply that his documents offer a valuable contribution to his research field (The Open University, 2023). Tools like Scite

(2018) can help to determine the citation type (related work, comparison, using the work, extending the work) from the citation statement and context, allowing for a more meaningful representation of the citations received (Valenzuela et al., 2015), even if not free from bias (Gadd, 2020).


By combining different analysis techniques, one can overcome their individual limitations, and gain more impactful insights. For instance, niche or newer papers are often less cited than their mainstream or older counterparts, but nevertheless they can gain visibility through bibliographic coupling or their PageRank score. This would not be possible through citation or co-citation analysis, as a low number of citation limits the co-citation potential (Donthu et al., 2021). Co-authorship analysis can shed light on a regional or interest cluster, and allow studying collaboration pattern between authors. By observing how thematic or social clusters changes over time, one can understand how a research field manifests and develops.

**Table 2.1:** Elements of scientific research networks in bibliometric analysis

ELEMENT	MEANING
Node	Entity (e.g. document or keyword)
Node color	Thematic cluster
Node distance	Centrality degree
Node size	Number of occurrences of an entity
Link	Co-occurrence of two entities
Link thickness	Number of co-occurrences of two entities

### 2.1.2 DOCUMENT SELECTION PROCESS

According to Donthu et al. (2021), there are four steps to perform a bibliometric analysis.



**Figure 2.1:** Document selection process in bibliometric analysis

**Source:** Adapted from Donthu et al. (2021)

The first step of the analysis should define the aims and scope of the study (e.g. future of research, thematic evolution), according to the degree of focus on the structure and retrospection that the scholar wants to achieve. The definitions should be broad enough to warrant bibliometric analysis, and yet focused enough to remain in the dedicated research field. This step must occur before the selection of analysis techniques and the gathering of bibliometric data, otherwise the analysis will be limited and will lose credibility.

The second step consists in choosing the best analysis techniques to meet the aims of the first step. For instance, a co-citation analysis (section 2.5.1) may be better suited for a review of the past, while a bibliographic coupling (section 2.5.2) would shed light on the current state of the art of the discipline. A co-word analysis of notable words in the implications and future research directions of full texts is more adapt to understand upcoming developments of the study field, while a keyword co-occurrence analysis of papers published in a long time span (section 2.3.2) allows understanding how the main themes evolved over time.

The third step involves the gathering of the data required for the selected analysis techniques. This is accomplished by choosing a bibliographic database based on the adequacy of its coverage, defining the search terms and filters based on the scope of the first step, and then exporting a comprehensive list of publications with the required attributes (e.g. authors, title, year, keywords, affiliations). Search terms may be found by brainstorming, interacting with experts, or from a preliminary reading of relevant publications.

The last step consists in running the performance analysis and scientific mapping, and then reporting the findings. Performance evaluation should summarize the performance of prolific research constituents using publication, citation and mixed measures. Scientific mapping instead should give an overview of the intellectual structure through network analysis and visualization. Scholars should not limit themselves to the bibliometric summary, but should discuss the findings and their implications, engaging with relevant trends.

## 2.2 AIM AND SCOPE OF LITERATURE REVIEW

THE literature review of this thesis aims to understand the fundamentals of data-driven decision making. The bibliometric analysis offers a set of tools to understand where to start our research in this research field. While performance metrics can point out to influential publications on the topic, scientific mapping can help by displaying the relationships between documents, topics and authors (see table 2.1).

To accomplish this, two bibliographic searches are carried out, a broad search of the *data-driven decision making* theme, spanning across many decades, and a more targeted search on one of the main themes that covered by the state of the art, limited to the last three years. Their scope should be broad enough to warrant the use of bibliometric analysis, which according to Donthu et al. (2021) is around high hundreds or few thousands. The main indicators of these research corpora are summarized in table 2.2. A similar methodology has been adopted by Toorajipour et al. (2021), where a first pilot search allowed a general grasp of the literature and identify criteria for the inclusion or exclusion of literature in a second search.

By running a keyword co-occurrence analysis on papers published over many years (broad search), and then segmenting the time span, it is possible to understand the thematic evolution of the research, thus viewing what research areas were born, gained momentum and eventually transformed throughout the period (see figure 2.3).

By narrowing the research to the latest development in the field (narrow search), it is possible to target the latest developments and discoveries. This operation should also reduce the diversity within the review corpus. It could be useful to see what publications influenced the most the review corpus, and if these references can be organized in research clusters. This can be accomplished by running a co-citation analysis on the references of the exported articles (see figure 2.8).

To understand how the current literature is organized, we can clusterize it by shared references (see figure 2.11).

## 2.3 BROAD SEARCH ANALYSIS

THE bibliographic database that was chosen for this literature review is Scopus (Elsevier, 2004). Scopus is one of the largest curated abstract and citation databases, and covers scientific journals, conference proceedings, and books (Baas et al., 2020). Similarly to Web of Science, it allows to search scientific literature, visualize references and citations, apply a variety of filters and export the results in different formats. However, at the time of writing Scopus allow for a higher limit of exported results compared to Web of Science (20 000 vs. 1 000).

Given the comprehensiveness of the database, a simple search for *data-driven decision making* returned too many results, most of them not useful for the search purposes as strictly related to other subject area, especially *Computer Science*. By limiting the research to the *Business, Management and Accounting* subject area, the filter should exclude more technical papers, less related to decision making in a business context. Other filter that were applied included limiting the document type to journal articles, to ensure a better comparability of search results, and limiting the document language to English, the main language for publications in this field (Donthu et al., 2021). The final query for the broad search is the following:

```
data AND driven AND decision AND making
AND ( LIMIT-TO ( DOCTYPE , "ar" ) )
AND ( LIMIT-TO ( SUBJAREA , "BUSI" ) )
AND ( LIMIT-TO ( LANGUAGE , "English" ) )
AND ( LIMIT-TO ( SRCTYPE , "j" ) )
```

This lead to 19 867 articles in March 2023, which is just below the maximum number of exportable results in Scopus. The main metadata of the search results are shown in the first column of table 2.2.

Before running any analysis, it is important to consider that by choosing different programs based on their strengths can help overcome the limitations of specific software. *Bibliometrix* offers, through its graphical user interface named *biblioshiny*, a wide number of performance indicators and scientific mapping visualizations, other than document filtering and text editing capabilities (e.g. keywords to remove or consider as synonyms). On the other hand, *VOSviewer* excels in scientific mapping visualization such as co-citation analysis and bibliographic coupling.

The first part of the analysis is carried out through *Bibliometrix*, as it offers a thematic evolution analysis via *biblioshiny*. Unlike *VOSviewer*, which can be downloaded and

**Table 2.2:** Properties of search results for broad and narrow searches

	BROAD SEARCH	NARROW SEARCH
Objective	Thematic evolution	Thematic clustering
Bibliographic database	Scopus	Scopus
Bibliometric software	<i>Bibliometrix</i>	Bibliometrix, <i>VOSviewer</i>
Time span	1982–2023	2020–2023
Average age (years)	5.15	1.56
Documents	19 867	934
Citations received	590 249	9 368
References cited	1 299 755	52 339

launched from a standalone application, *Bibliometrix* comes in the form of package for the *R* programming language, so *R* must be installed beforehand in the operating system. From an *R* console of choice, the commands to install, run *Bibliometrix* and open its graphical interface are:

```
# install required packages
install.packages(c("Bibliometrix", "shiny"))

# launch bibliometrix interface with a file upload limit of 1GB
library(Bibliometrix); biblioshiny(maxUploadSize=1000)
```

The upload limit has been increased to allow the analysis of large files, e.g. when exporting the full list of reference of the search results.

After accessing *biblioshiny* and uploading the file downloaded from Scopus, *Bibliometrix* converts the file content into a bibliometric dataframe, and presents a table with the completeness of bibliographic metadata. This check informs the user of possible issues with the dataset that could compromise the application or veracity of future analysis. From now on, it is possible to apply further filtering to the data collection and run analysis on the documents.

### 2.3.1 TRENDING KEYWORDS

A powerful tool that *Bibliometrix* offers to understand the evolution of a research field is the chart that can be generated from author keywords in *Documents > Words > Trend Topics*. Figure 2.2 gives an idea of which topics are trending in the dataset time span. It appears as a horizontal box plot that selects the most frequent keyword for every year, and then determine how much that keyword has been used in the previous and following years.

As seen in the chart, the oldest articles of the corpus coincide with the commercial adoption of Relational Database (RDB) to store data, and the standardization of Structured Query Language (SQL) in 1987 (International Organization for Standardization, 1987). The idea of adopting a single database shared between multiple decision makers, that could reduce data entry costs and provide centralized management of data integrity, caused a rapid adoption of this technology (Jarke, 1986).

During the 80s and 90s the main research on data-driven decision making was focusing on Decision Support Systems (DSS) and Expert Systems (ES). The first mention of AI in the corpus appears in Jarke and Radermacher (1988), where the authors stress the importance of model management (knowledge bases, evolutionary principles, mathematical models) to achieve higher level of cognitive competence in DSS.

The chart confirms the findings of Duan et al. (2019) about a terminology shift from *Expert Systems* to *Knowledge-based Systems* in the 90s, due to the bad reputation of early ES, as they failed to deliver the early promises of AI. This terminology shift however does not reflect an evolution of the underlying technology, which was left unchanged.

A similar change in the label due to bad reputation can be seen in Simon (1997), where the first generation of *Management Information Systems*, that were designed to “inundate managers’ desks with irrelevant reports that they had neither time nor desire to study” without knowing what information would be relevant to managers’ decisions, were replaced by “new buzzwords like Management Decision Aids”.

The term *multi-criteria decision making* refers to a sub-discipline of operations research that evaluates multiple conflicting criteria in decision making (Hong & Vogel, 1991). Trade-off rules can be learned by machines and then be applied automatically. Thus, they can be categorized as a subset of AI tools (Sharma et al., 2022).

According to Arnott and Pervan (2005), *Executive Information Systems* are data-oriented DSS that provide reporting about the nature of an organization to management. Despite the executive title, they are used by all levels of management. The authors considers Business

## 2.3 BROAD SEARCH ANALYSIS




Figure 2.2: Trending topics in broad search documents



Intelligence (BI) a contemporary term for Executive Information Systems.

*Product attributes* have been a central theme for many years due to its impact in operations and marketing decision making, other than consumers' preferences, perceived quality and satisfaction (Michalek et al., 2005).

The implementation of *Information Systems* cover a wide variety of topics, such as simulations, computer-aided scenario analysis (Fabianova et al., 2021).

The term *Big Data* was not coined until the early 2000s when it became clear that the amount of data being generated and stored was growing exponentially (Ylijoki & Porras, 2018). According to the author, Big Data Analytics can improve the decisional processes by providing better quality decisions. The agility of an organization increases with the speed of decisions which, combined with quality decisions, causes a positive impact on organizational performances.

A term that gained traction in the last period is sustainability, considered as a central feature of new business models (Dhir et al., 2023), circular economies (Riggs et al., 2023) and smart cities (Ju et al., 2018).


AI (re)emerged in the last years as a leading topic, thanks to advancements in learning algorithms, availability of data to train them and their successful application in business contexts, especially Supply Chain Management (SCM) (Sharma et al., 2022).

### 2.3.2 THEMATIC EVOLUTION

The chart seen in figure 2.3 has been generated from authors' keywords and offers a more detailed thematic evolution. Here the corpus time span is divided in time slices to better understand the main research areas of the literature in specific years. In particular, the cutting points set for this chart allows studying the evolution of data-driven decision making over different decades. The chart can be generated by navigating the *Bibliometrix* menu to *Conceptual Structure > Network Approach > Thematic Evolution*. The algorithm creates a network of keywords for every time slice, assigns each keyword to a cluster in each network, and then plots a Sankey diagram that shows how keywords moved from a cluster in a previous time slice to another cluster in the following time slice.

These clusters were also very far and separate from each other, indicating that these themes were addressed in separate publications, with little to no overlapping. However, this can also be attributed to the low amount of documents belonging to this time slice. The periods 1991–2001 and 2001–2010 show instead the origin of new research areas, albeit small, and an increase in the number of links between areas. Then, the periods

## 2.3 BROAD SEARCH ANALYSIS



**Figure 2.3:** Thematic evolution in broad search documents

2011–2019 and 2020–2023 show a progressive convergence of the various disciplines into four highly interconnected clusters: big data, innovation, artificial intelligence and customer engagement. The network in figure 2.4 illustrates better the main keywords that represent these clusters. In particular, the most used keywords in each cluster are:

- **Red cluster:** innovation, performance, knowledge management, dynamic capabilities, smes
- **Green cluster:** artificial intelligence, machine learning, social media, data-driven decision making
- **Blue cluster:** big data, digital transformation, covid-19, supply chain, industry 4.0, sustainability
- **Purple cluster:** customer engagement, purchase intention, online shopping, perceived value, loyalty

The purple cluster that appears in figure 2.4 has too few nodes to appear in the Sankey diagram in figure 2.3.

From the analysis of the last time slice, one can see that AI plays a central role in modern data-driven decision making. It makes sense to explore the impact of these technologies in the research field with the help of a narrower query on recent publications.

The review corpus used for the previous analysis was well suited to illustrate the variety of themes on data-driven decision making that arose in the last 40 years. However, it would be too broad to analyze the impact of AI in recent years. Therefore, there is the need for a new, more restricted query, that considers both the new keyword and time slice. We report this new query in section 2.5, after a brief introduction on the meaning of AI and its recent interest in business adoption.

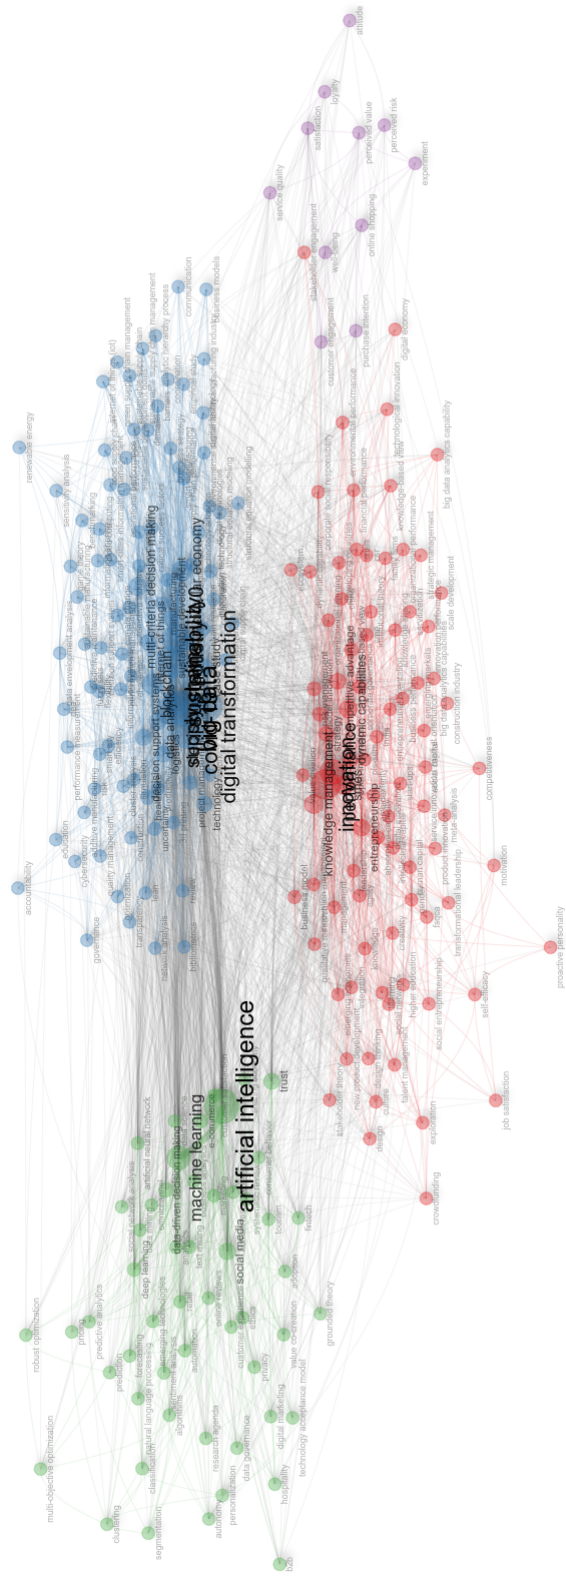


Figure 2.4: Thematic networks in broad search documents in 2020–2023

# 2.4 INTRODUCTION TO ARTIFICIAL INTELLIGENCE

ARTIFICIAL Intelligence has emerged as a promising technology to support, and potentially even replace, human managers in the process of taking decisions. In this section, we aim to provide an accessible introduction on the term, covering the evolution of its original meaning and the current focus of the research field. As with many terms that make their way from specialist academic domains into common usage, the term Artificial Intelligence (AI) is fraught with misinterpretation (Jordan, 2019).

AI has been a study field for more than sixty years, but has been influenced by different disciplines in thousands of years of human history (Russell & Norvig, 2020). As a consequence, its meaning has evolved during this period of time, and the technology has experienced summers, periods of hopes and promises, and winters, periods of disillusionment (Duan et al., 2019).

## 2.4.1 EVOLUTION OF THE TERM

The term *Artificial Intelligence* was originally coined in 1955 by Dartmouth professor John McCarthy to distinguish his research agenda from that of the MIT professor Norbert Wiener (Jordan, 2019). In 1947, Wiener invented the term *cybernetics* to describe his own vision of intelligent systems, which was intimately related to operations research, statistics, pattern recognition, information theory, and control theory. McCarthy, on the other hand, highlighted the links to logic and the exciting goal of establishing a human-level intelligence in software and hardware. Wiener research domains were frequently inspired by human or animal behavior and were allegedly focused on low-level signals and decisions (easy to imitate). McCarthy's AI was supposed to focus on something different: humans' high-level or cognitive ability to reason and conceive abstractly (difficult to replicate). High-level thinking and cognition, however, remain elusive more than sixty years later, despite claims from corporate researchers (Marcus, 2023; Tiku, 2022).

Oxford professor of Philosophy and Ethics of Information Floridi (2023) highlights this difference by explaining how of chatbots based on Large Language Models (LLM), such as *ChatGPT* (OpenAI, 2021), work:

“They do not think, reason or understand, and they are not a step towards any sci-fi AI. They have nothing to do with the cognitive processes present in the animal world, and in the human mind, to manage semantic contents successfully.

However, with the staggering growth of available data, quantity and speed of calculation, and ever-better algorithms, they can do statistically (working on the formal structure), what we do semantically (working on the meaning of texts).”

The advancements commonly referred to as AI primarily originated from engineering fields associated with fundamental pattern recognition, motion control, and statistics. An example is Machine Learning (ML), an algorithmic discipline that incorporates concepts from statistics, computer science, and numerous other fields to develop algorithms that process data and generate predictions. It is reasonable to conclude that nowadays Wiener’s intellectual purpose has prevailed, even if in McCarthy’s words (Jordan, 2019).

The confusion of these terms is exacerbated by the most celebrated achievements of ML, which have occurred in areas associated to human capabilities like game-playing, robotics, speech recognition, computer vision, chat interaction and art generation (Jordan, 2019). On top of this, the use of terms such as *neural networks* and *learning*, may feed the assumption that AI’s ultimate goal is imitating humans.

However, mere imitation has rarely been the focus of AI researches, as Russell and Norvig (2020) suggest in their analogy with history of flight:

“Planes are tested by how well they fly, not by comparing them to birds.

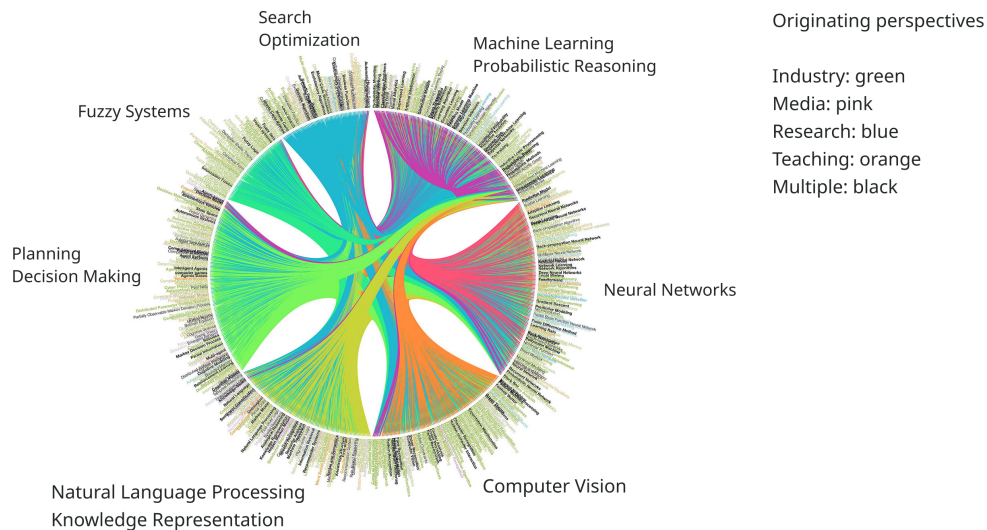
The quest for artificial flight succeeded when engineers and inventors stopped imitating birds and started using wind tunnels and learning about aerodynamics. Aeronautical engineering texts do not define the goal of their field as making machines that fly so exactly like pigeons that they can fool other pigeons.”

### 2.4.2 FOCUS OF CURRENT RESEARCH

Russell and Norvig (2020), authors of one of the most popular AI textbook in the world, define AI as the study and construction of agents that do the right thing.

The concept of *rational agent* finds wide adoption in numerous research fields, including economics, where a decision maker maximizes utility or some measure of social welfare. Rationality refers to making decisions that are deemed as most sensible given the available knowledge. By providing these definitions, the authors emphasize that an intelligent agent is not merely an entity that can exhibit human-like behaviors, but any system that can sense its environment, process information, and act upon it to achieve its objectives.

Figure 2.5 is the result of a keyword co-occurrence analysis of more than 600 000 AI-related publications issued between 1998 and 2017. This study has been conducted by Elsevier (2018), one of the largest academic publishers, with the aim of defining and understanding the AI field and its reach. The methodology they used is explained by Siebert et al. (2018).




**Figure 2.5:** Keyword co-occurrence of AI publications in 1998–2017

**Source:** Adapted from Elsevier (2018)


We observe how the literature on AI in this period focused on Machine Learning (ML), Artificial Neural Networks (ANN), Computer Vision (CV), Natural Language Processing (NLP), Probabilistic Reasoning, Knowledge Representation, Decision Making, Fuzzy Systems, Search and Optimization. We will refer to these terms as *AI technologies* and *AI use cases*.

### 2.4.3 TRENDS OF CURRENT BUSINESS ADOPTION

Given the popularity of AI in the current year, FactSet analyzed the quarterly earnings conference calls of S&P 500 companies over a 10 years period to find mentions of AI (Butters, 2023). Figure 2.6a shows that 110 companies mentioned AI, 46 of which also mentioned Generative Artificial Intelligence (GenAI). This figure is significantly higher than the 5-year average of 57 and the 10-year average of 34. Figure 2.6b shows that the communication services and information technology sectors have the highest percentage of companies citing AI in their earnings calls.



(a): Number of S&P 500 companies mentioning AI



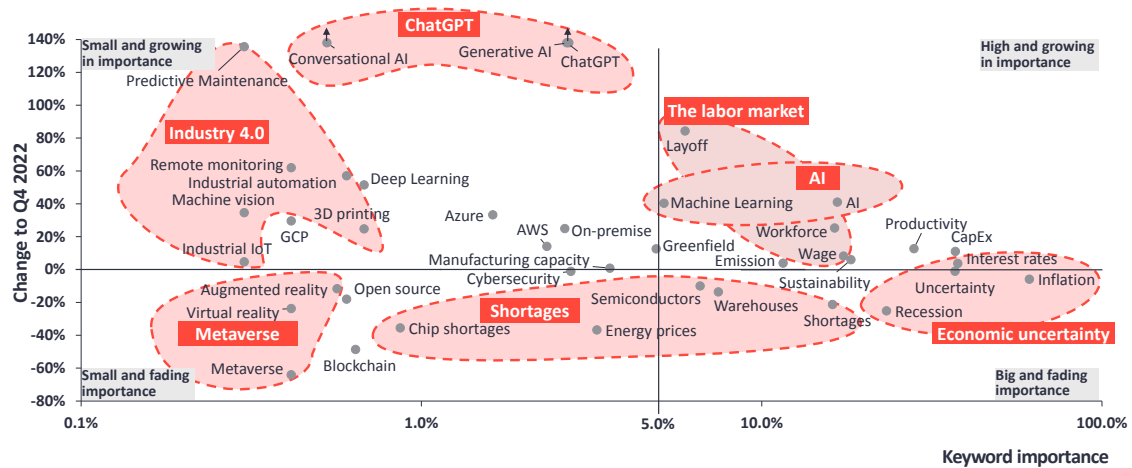
(b): Percentage of companies by sector mentioning AI

**Figure 2.6:** S&P 500 companies mentioning AI in their earning calls in 2013–2023

**Source:** Adapted from Butters (2023)



Similarly to Butters (2023), IoT Analytics (2023) conducted a keyword analysis on the Q1 2023 and Q4 2022 earnings calls of about 3 000 companies listed in the U.S. to analyze changes in the most mentioned keywords.



**Figure 2.7:** Most mentioned keywords in earning calls in Q1 2023 compared to Q4 2022

**Source:** Adapted from IoT Analytics (2023)

Among these companies, less than 1% mention blockchain in their documents, with a 60% reduction over the last quarter. The companies mentioning AI instead are 20%, with a 40% increase over the last quarter. More specific terms such as ML and Deep Learning (DL) are mentioned respectively by 5% and 0.7% of the sample, with a 40% increment over the last period. The biggest variations compared to Q4 2022 appear in AI applications such as predictive maintenance, conversational AI and GenAI. Metaverse, blockchain and supply shortages are instead heavily fading of importance.

## 2.5 NARROW SEARCH ANALYSIS

FROM figures 2.3 and 2.4 we observe how AI has recently gained recognition among decision making scholars. This interest for AI also emerges in executive surveys (?? and ?? 2.6 and 2.7). To narrow down the literature focus on the role of AI in decision making, we restricted the search query by specifying different keywords and setting a time filter that limit the results to articles published in the last three years:

```
artificial AND intelligence AND decision AND making  
AND PUBYEAR > 2019  
AND ( LIMIT-TO ( DOCTYPE , "ar" ) )  
AND ( LIMIT-TO ( SUBJAREA , "BUSI" ) )  
AND ( LIMIT-TO ( LANGUAGE , "English" ) )  
AND ( LIMIT-TO ( SRCTYPE , "j" ) )
```

Compared to the first one, the query returned 934 results as of March 2023. The main attributes of the two review corpora are summarized in table 2.2.

### 2.5.1 CO-CITATION ANALYSIS

Before starting to analyze these 934 documents directly, it can be a good idea to understand what are the main research areas that influenced the current literature by analyzing the 52 339 references of the exported documents. To further reduce the computational complexity of the layout and clustering algorithms, only references that have been cited at least 3 times are considered. This reduces the group size from 52 339 to just 2 043. The publication date of these references ranges from 1907 to 2023, but half of them has been published in 2017 or later.

The co-citation network displayed in figure 2.8 can be generated in *VOSviewer* by navigating *Create... > Create a map based on bibliographic data > Read data from bibliographic database files > Scopus > Co-citation > Cited references*. For the exported corpus, the counting method did not impact the network layout, so the default full method was selected. No threshold or document number limit was selected. The size of each node represents its total link strength.

*VOSviewer* identified six clusters, and assigned each reference to one of them. Results are reported in table 2.3. A manual overview of the most cited papers per clusters confirmed the bibliometric principle that publications cited together multiple times are thematically similar. It is important to mention that nearby clusters are more related than cluster far apart from each other. In this case, the red cluster (fuzzy sets) is more thematically similar to the purple one (uncertainty) than the others, and the remaining clusters (ML, business cases, decision making, analytics) are more related to each other than the red or purple ones.

The clusters identified in this co-citation analysis appear similar to the ones identified in the keyword co-occurrence analysis performed by Elsevier (2018) with a far larger analysis corpus. Differences can be explained because of the search filters applied (more




Figure 2.8: Co-citation analysis of narrow search references

business-centric in this analysis) and the methodology adopted.

**Table 2.3:** Properties of clusters in co-citation analysis

ID	COLOR	SIZE	RESEARCH AREA	EXAMPLE
1	Red	530	Fuzzy sets	Rodríguez et al. (2012)
2	Green	461	Machine Learning	Jordan and Mitchell (2015)
3	Blue	441	Business applications of AI	Duan et al. (2019)
4	Yellow	307	Behaviors, bias and aversions	Dietvorst et al. (2014)
5	Purple	160	Uncertainty, belief functions	Xiao (2020)
6	Cyan	144	Data Analytics, Supply Chain	Dubey et al. (2020)

By analyzing the most influential references for each cluster, one can get valuable insights into the research area, and understand how it relates with the application of AI to decision making. Instead of reading through every reference, which can be time-consuming and overwhelming, focusing on the first five references can provide a good overview of the cluster's content.

#### FUZZY SETS

In the red cluster, L. A. Zadeh (1965) appears as the most cited reference, and this is reflected on its node size in figure 2.8. The author developed the theory about *fuzzy sets* to address the challenge of providing precise inputs to intelligent systems. While in classical logic an entity is either a member or a non-member of a particular set, fuzzy logic operates on the premise that an entity can possess a varying degree of membership within a (fuzzy) set.

Fuzzy logic addresses the issue of *vagueness* in the mapping from symbolic terms to real-world scenarios, rather than focusing on *uncertainty* about the truth of well-defined propositions. Vagueness is a pervasive challenge when applying logic, probability, or standard mathematical models to reality.

For instance in the hiring process, classical logic dictates that a candidate is either qualified or unqualified for a job. However, using fuzzy logic, we can assign a degree of membership to a candidate's qualifications. A candidate that possesses most but not all the required qualifications, gets a moderate degree of membership assigned to their qualifications. This approach can help employers consider a wider range of candidates

instead of just the ones that possess all the qualifications.


Another influential work of the author explain how linguistic variables can be used to model and analyze imprecise information (L. Zadeh, 1975). A linguistic variable is a mathematical representation of a concept that cannot be expressed precisely, but can instead be described using words or natural language. The values of a linguistic variable are not fixed but can vary over a range of possibilities, depending on the context and the speaker's perception. For example, the variable *temperature* can be described using linguistic values such as *hot* and *cold*, depending on the context and the speaker's interpretation.

*Intuitionistic fuzzy sets* are a type of fuzzy set that extends the traditional notion of a fuzzy set by introducing an additional degree of uncertainty. They were developed by Atanassov (1986) as an attempt to capture the hesitancy or vagueness that arises in decision making situations.

Similarly, *hesitant fuzzy sets* were proposed by Rodríguez et al. (2012) as an extension of fuzzy sets where elements have multiple degrees of membership called hesitant degrees, each of which represents a different level of confidence or preference, allowing for more precise representation of uncertain or ambiguous information. For example, you may be unsure whether a supplier is *very good* or just *good*, but you are confident that it is at least *good*. In hesitant fuzzy sets, you would represent this uncertainty by assigning multiple hesitant degrees of membership to the element, such as *0.7 very good* and *0.9 good*.

Tang and Liao (2021) conducted a bibliometric analysis on the literature about large-scale group decision making and its connection with Big Data (BD). Their analysis show that BD techniques, such as data mining, can be utilized to extract public preferences and identify key points in public opinions. This provides an objective foundation for relevant departments and large-scale experts to make informed decisions and enhance their decision making capabilities. However, only few articles actually implemented BD in their frameworks, suggesting that it is challenging to make full use of it. The authors proposed a framework to implement BD in the decision process (figure 2.9).

Figure 2.9 illustrates this framework and its four phases: intelligence, design, selection and implementation. In the intelligence phase, data is collected from various sources to identify the issue. This data should be stored and organized in RDB. In the design phase, different data mining and machine learning techniques can be applied to analyze data and create a model, which can be used with Online Analytical Processing (OLAP) or predictive analysis to predict future outcomes. The selection phase involves evaluating and making decisions using methods such as fuzzy set and visualization tools like Gephi. Finally, in



**Figure 2.9:** Framework to use Big Data in a decision process

**Source:** Adapted from Tang and Liao (2021)

the implementation phase, the BD techniques can be used to monitor the results of the solution obtained from the previous steps.

The authors also illustrated the differences with small-scale group decision making, summarizing the main results from five perspectives: dimension reduction, weighting and aggregating decision information, preference modeling, consensus reaching, and social network analysis. Their analysis show that a considerable part of the literature focus on emergency problems such as earthquakes and fire accidents, as these situations involve several decision makers from different professional backgrounds. The authors claim that monitoring the flood of information in social media (such as Weibo and Twitter) during and after these emergencies allow for public emotion analysis and viewpoint mining.

Other important research area for large-scale group decision making are water management, energy management and supply chain management. According to the authors, future scholars should focus on issues in the operation research and management science fields, as they seem less represented by the literature.

### UNCERTAINTY, BELIEF FUNCTIONS

The purple cluster is strongly tied with the red one and covers topic of uncertainty, entropy, pignistic probability and belief functions.

Dempster (1967) published a seminal paper in the field of uncertainty reasoning, introducing what is now known as Dempster’s rule of combination. Glenn Shafer, who was a graduate student at the time, studied Dempster’s work and realized that it could be further developed by generalizing the notion of probability measures to belief functions. Shafer’s extension of Dempster’s work led to the development of the theory of evidence, which is also known as Dempster-Shafer theory. In his book “A Mathematical Theory of Evidence”, Shafer (1976) acknowledges Dempster’s contributions and discusses how his work builds upon and extends Dempster’s theory. Specifically, Shafer shows that Dempster’s rule of

combination can be derived from a more general framework based on the theory of belief functions, which provides a more flexible and powerful way of reasoning with uncertain and imprecise information. Belief functions are a generalization of probability measures, allowing for the representation of partial knowledge and conflicting evidence. The book also presents a number of applications of the theory, including ES, decision making, and data fusion.

Murphy (2000) explores the use of belief functions in ES and proposes solutions to the normalization problems that arise when combining multiple evidence. The author points out that Dempster's rule increase the measure of belief in the dominant subset, and suggest that averaging the masses of decisional rules identifies combination problems, shows the distribution of belief, and preserves a record of unassigned belief (ignorance).

Deng (2020) discusses the topic of measuring uncertainty in evidence theory. Evidence theory is an extension of probability theory that is better equipped to handle uncertain and imprecise information. While there are several methods for measuring uncertainty in evidence theory using basic probability assignment, these methods are not without controversy, and there is still much debate on the ideal way to measure the uncertainty of basic probability assignments. Deng's article reviews existing uncertainty measures in evidence theory and introduces Deng entropy as a new method for measuring uncertainty.

Xiao (2020) introduced a fuzzy Multi-Criteria Decision Making (MCDM) method that integrates Dempster-Shafer theory with belief entropy to address the issue of uncertainty in MCDM. Each decision criteria is modeled as evidence and all alternatives compose the frame of discernment in the framework of Dempster-Shafer theory. This method considers both subjective and objective weighting of criteria to generate more appropriate basic probability assignments and uses Dempster's rule of combination to fuse multiple pieces of evidence into composite evidence. It can therefore model uncertainty and help decrease uncertainty caused by subjective human cognition to improve decision making.

### MACHINE LEARNING

The green cluster covers mainly the topics of ML and ANN. Although it is found on the left side of the co-citation network, close to the remaining clusters, it represents another technical cluster. The most influential sources are Nature, Science and a wide variety of journals from Institute of Electrical and Electronics Engineers (IEEE), but there are also management journals and seminal publications of management literature, such as March and Simon (1958).

Sutton and Barto (2018) provide a comprehensive introduction to the key concepts and algorithms in the field of reinforcement learning, on top of what they already wrote in the first edition of the book in 1998. The main topics covered are online learning algorithms, function approximation, and off-policy learning. They also discuss reinforcement learning's relationship to psychology and neuroscience, and include case studies on topics such as AlphaGo and IBM Watson's betting strategy. The authors suggest that the technology could have significant implications for a range of fields and industries. For example, they note that reinforcement learning could lead to advances in autonomous vehicles, robotics, and personalized medicine. However, they also acknowledge that there are potential risks associated with the technology, such as the possibility of job displacement or ethical concerns related to the use of reinforcement learning in areas such as finance or law enforcement.

Mnih et al. (2015) describe how reinforcement learning, which is deeply rooted in psychological and neuroscientific perspectives on animal behavior, can be used to optimize an agent's control of an environment. However, agents must derive efficient representations of the environment from high-dimensional sensory inputs and use these to generalize experience to new situations. The article is impactful because it introduced an artificial agent named *deep Q-network* that can learn successful policies directly from high-dimensional sensory inputs using end-to-end reinforcement learning. The authors tested the agent on *Atari 2600* games, and surpassed all previous algorithms performances, obtaining a level comparable to a professional human player in 49 games. The *Q* in *Q-network* refers to the *quality* function  $Q(s, a)$ , that represents the discounted future reward (i.e. the best possible score at the end of game) when action *a* is performed in state *s*, assuming an optimal continuation from that point on (Matiisen, 2015).

Jordan and Mitchell (2015) provide an overview of the current state of Machine Learning research and discuss future trends and potential applications. The authors note that ML has made significant progress in recent years and has become increasingly important in fields such as computer vision, natural language processing, and robotics. They discuss the role of Big Data in ML and the challenges involved in developing algorithms that can effectively learn from large datasets. The paper also covers topics such as Deep Learning, unsupervised learning, and reinforcement learning. Additionally, the authors discuss potential applications of ML in fields such as healthcare, transportation, and education, and consider the ethical and societal implications of these technologies.

Silver et al. (2016) describes the development of *AlphaGo*, a program that was able to defeat 5-0 the European champion Fan Hui in the game of Go. Few months after, in



March 2016, *AlphaGo* will win 4-1 against the world champion Lee Sedol. Although Go has simple rules, it is a highly complex game. In comparison to chess, Go has a bigger board that allows for more extensive gameplay and longer games. Additionally, there are more potential moves to consider per turn. For instance, there are 361 possible initial moves in Go compared to just 20 of chess. Up to this point, computer Go programs were unable to compete with high level players in a full-size 19x19 board due to the complexity of the game. The article outlines the various components of *AlphaGo*'s design, including the deep neural networks that were trained on large datasets of self-played Go games, and a Monte Carlo tree search algorithm that was used to select moves more efficiently during gameplay. The article also discusses the challenges involved in developing an AI program capable of playing Go at a world-class level, including the game's complexity and the difficulty of constructing a training dataset.

Barredo Arrieta et al. (2020) discuss the importance of developing explainable AI systems that are transparent and understandable to human users. The paper begins by introducing the concept of eXplainable Artificial Intelligence (XAI) and discussing its relevance in today's world where AI systems are being used in various fields such as healthcare, finance, and transportation. The authors argue that XAI is necessary to build trust in AI systems and to ensure that they are used responsibly. Table 2.4 illustrate the different explanation purposes of AI audience. The authors classify different ML frameworks according to their level of explainability. In particular, they identify two categories:

- **Transparent models:** linear regression, logistic regression, decision trees, k-nearest neighbors, rule-based learners, general additive models, Bayesian models
- **Post-hoc explainable models:** tree ensembles, multiple classifier systems, support vector machines, multi-layer neural networks, convolutional neural networks, recurrent neural networks

The models in the first group present human-readable variables, but in more complex models with a large amount of rules or variables it is necessary to decompose the model into readable chunks. Mathematical tools are overall needed to analyze the interaction between variables. For the models in the second group, none of the above is applicable. They need *post-hoc* analysis through model simplification, feature relevance estimation or visualization techniques in case of image classification.


In figure 2.10, the Local Interpretable Model-agnostic Explanations (LIME) technique has been applied to explain the object recognitions of a CV model. The model is prompted

**Table 2.4:** Possible audience of eXplainable Artificial Intelligence**Source:** Adapted from Barredo Arrieta et al. (2020)

AUDIENCE PROFILE	EXAMPLE	EXPLANATION PURPOSE
Domain expert	Doctors	Trust the model
Regulatory entities	Audit agencies	Certify model compliance
Executives	CEOs	Understand applications
Product owners	Data scientists	Ensure efficiency
Affected users	Job candidates	Verify fairness of decisions

to recognize the object contained in figure 2.10a. It returns a list of two objects along with its confidence level: *electric guitar* 0.32 and *acoustic guitar* 0.24. By applying the LIME technique, the model also returns the area of the original image that resemble an electric guitar and acoustic guitar (figures 2.10b and 2.10c).

However, the authors also note that there are technical and ethical challenges associated with the development of explainable systems, such as the trade-off between explainability and accuracy. Finally, the authors emphasize the importance of responsible AI development and the need for interdisciplinary collaborations to address the challenges associated with XAI.



(a): Original image      (b): *Electric guitar* explanation      (c): *Acoustic guitar* explanation

**Figure 2.10:** Explanation of object recognition through the LIME technique**Source:** Barredo Arrieta et al. (2020)

### DATA ANALYTICS, SUPPLY CHAIN

Similarly to the green one, the cyan cluster is concentrated on the left side of the co-citation network, but extends to the center with few nodes. It is thus expected to be strongly tied with business applications of AI and behavioral decisions.

Min (2010) explores various subfields of AI suitable for solving practical problems relevant to Supply Chain Management. For decades, this technology has demonstrated potential for enhancing productivity and decision making, due to its capacity to recognize patterns, understand business issues, and process information effectively. Despite this potential, AI applications to SCM have been limited. The author describes three categories of AI tools and list their application areas:

- **Agent-based systems:** demand planning and forecasting, customer relationship management, negotiation, order picking
- **Genetic algorithms:** network design
- **Expert systems:** inventory planning, make-or-buy decisions, supplier selection

According to the author, there are three reasons for which AI adoption is not so widespread. The first is that this technology relies heavily on computer software, which may lead to wrong decisions if badly programmed. AI cannot overcome this issue by itself since it does not have free will. The second reason refers to the difficulties in understanding the decisions of AI solutions. The third reason recalls that AI is better suited for specific and narrowly focused SCM issues, since it features knowledge acquisition bottlenecks that prevent handling uncertainty in cross-functional or cross-border SCM environment. Finally, the author suggests a list of AI research topics to advance the decision making process in SCM. The most important ones are the application of agent-based systems to supply chain integration and partnerships, the adoption of expert systems to assist in outsourcing decisions, the implementation of Machine Learning to overcome the existing issue of supplier selection.

Dubey et al. (2020) explore the use of BD analytics and AI in manufacturing organizations to improve operational performance, taking in account entrepreneurial orientation and environmental dynamism. In the context of this article, entrepreneurial orientation refers to an organization's tendency to be innovative, proactive, and risk-taking in pursuing new opportunities, such as investing in emerging technologies like BD analytics and AI. Environmental dynamism refers to the degree of change and unpredictability in a firm's external environment. The study found out that when the external environment is moderately dynamic, organizations with an entrepreneurial orientation are more likely to adopt

these technologies and see improvements in their operational performance. The article suggests that more research is needed to address the concerns of data quality, which can affect the accuracy and reliability of predictive analytics, and to explore the impact of other factors, such as organizational culture and leadership.

Toorajipour et al. (2021) conducted a systematic literature review on 64 articles about the use of AI in SCM and its potential benefits, similar to what Min (2010) did 11 years before. Indeed, the authors elaborate on Min's future research proposals to elaborate theirs. The prevailing AI techniques in SCM are ANN, agent-based systems, multiagent systems and fuzzy logic, compared to the expert systems used in the past. In their conclusions, the author stress another important concept to enhance the use of AI in SCM: it is important to have both appropriate AI-based software and well-defined SCM problems that can benefit from such software. Therefore, an effort is required from both AI researchers and SCM practitioners.

Belhadi et al. (2022) evaluate different AI techniques in building supply chain resilience, considering the adaptive capacity to deal with disruptive events and to swiftly regain its previous performance level. They found out that fuzzy logic programming, machine learning and agent-based systems are the most promising techniques. In particular, fuzzy logic can be used to address the incompleteness and ambiguity in data collected, while machine learning can leverage the large amount of data collected in supply chain operations to deliver accurate predictions. Agent-based systems can instead simulate reasonable actions under constraints.

Helo and Hao (2022) focused on how AI can be implemented in SCM and how it helps to improve operational performance, illustrating four AI applications in customer, production, quality and services management. According to previous AI research and cases, AI can be implemented in the following areas:

- **Learning systems:** can adjust behavior based on dynamically observed data
- **Situation-aware systems:** can detect and understand the prevailing conditions, and adjust behavior according to modes and situations
- **Autonomous decision making systems:** can execute decisions in contrast with traditional DSS
- **Analysis systems:** can process streaming images, video, audio and non-structured text type of data

As a future research suggestion, the authors suggest to focus on the cultural and organizational factors influencing the adoption of AI.

### BUSINESS APPLICATIONS OF ARTIFICIAL INTELLIGENCE

The blue cluster represents the application of AI in different organizational contexts.

Davis (1989) researched and developed two psychometric scales to measure the impact of perceived usefulness and ease of use to the actual usage of four computer software used to edit electronic mail, computer programs, business charts and bitmap images. The author found out that usefulness is more correlated to actual usage and user acceptance than ease of use, as applications are adopted primarily because of the functions they perform, and only secondarily for how easy or hard it is to get the system to work. The impact of this research finds application in software development, where according to the author “designers [...] have tended to overemphasize ease of use and overlooked usefulness”.

H. Chen et al. (2012) discusses the evolution of Business Intelligence (BI) and Analytics research and its impact on different businesses cases. They highlight the emergence of web-based and mobile-based analytics and the mining of unstructured user-generated content. The authors suggest that existing BI educational courses should be revised, as now more than ever the subject requires an interdisciplinary approach, spacing from statistical skills and familiarity with business departments, to specific domain knowledge for BI application. These courses would address the predicted shortages of data-savvy managers and business professionals with strong analytical skills.

Huang and Rust (2018) explores the impact of AI on service tasks and jobs, and provides a theory on how firms can decide between humans and machines for accomplishing those tasks. The authors elaborate a framework made of four intelligence levels, that represents the progressive job replacement by AI. This progress is represented as follows:

1. AI takes over mechanical tasks
2. AI takes over analytical tasks
3. AI takes over intuitive tasks
4. AI takes over empathetic tasks

According to the authors, since AI applications are adopted first for cost reduction reasons, firms that follow a cost leadership theory will use them sooner than quality leaders, which will rely more on human skills. In addition, services that benefit from a stronger human interaction will be more difficult for AI to replace. To remain competitive in the AI era, the authors suggest that learning programs should emphasize creative thinking and intuition in interpreting data or making decisions, rather than training students to overlook data and analysis skills, which will be replaced sooner by AI.

Duan et al. (2019) identified the challenges and opportunities from the applications of AI for decision making through a systematic literature review of publications covering the AI topic from the International Journal of Information Management. Historically, AI has been acting both as a supporter and a replacement of human decision makers, and the progress in enabling AI to do increasingly complicated tasks that involve cognitive abilities, feeling emotion, and driving processes that were previously thought impossible. The authors highlight the peak of rule-based systems in the 90s, known as Expert Systems (ES) globally, but often referred to as *Knowledge-based Systems* within the business context. It is argued that the evolution of ML and the introduction of multilayered Artificial Neural Networks (ANN), whose results are not trivial to explain and codify into human knowledge, lead to a comeback of the AI term. Using AI for decision making, either in assisting or replacing humans, has been one of the most important applications in AI history. Twelve research propositions are provided for future researchers to address challenges and opportunities of applying AI in decision making, spacing from theoretical developing of the concept to actual implementation of AI solutions and their interaction with humans.

### BEHAVIORS, BIAS AND AVERSIONS

The yellow cluster covers the topic of behavioral decision making, bias and algorithmic aversion.

Lawler and Elliot (1996) investigate the impact of an ES used as a decision aid in a job evaluation system. The results suggest that the ES improved accuracy and reduced complexity for users, but did not significantly affect their confidence or satisfaction with the decisions taken. Existing literature showed that users were less motivated to use the decision

aid in circumstances where the decision consequences were little, favoring its usage when consequences are significant. However, the study also indicates that subjects became less satisfied with using an ES as task complexity increased. In fact, the ES took more time to complete a complex task than humans, contrary to authors' theoretical expectations. Already at the time, computer applications were sometimes considered as a threat to employee autonomy or job security.

Highhouse (2008) discusses the implicit beliefs that hinder the adoption of decision aids in employees selection, and how understanding these beliefs can lead to successful interventions. People assume that the complex characteristics of each candidate can be best assessed by a sensitive, equally complex human being, and this leads to two main beliefs. The first one considers possible to achieve perfect accuracy in predicting job performance. This leads to resistance towards analytical selection methods, because peoples' selection is seen as certain and not subject to error. A second belief argues that experience improves the prediction of human behavior. This results in over-reliance on intuition and reluctance to use selection decision aids, as it may undermine one's credibility. To address this belief, interventions can focus on structuring expert intuition and mechanically combining it with other decision aids.

Dietvorst et al. (2014) explore the concept of algorithm aversion, where individuals prefer human forecasters over evidence-based algorithms, despite the latter being more precise in predicting future outcomes. Through five studies, it is demonstrated that people have lower confidence in algorithms and are less inclined to select them over a less competent human forecaster, even when the algorithm has a better performance record. The reluctance to rely on algorithms can be expensive, given that numerous decisions necessitate a forecast and algorithms are typically superior to humans in forecasting accuracy. However, the authors do not provide in the paper any advice on how to address algorithm aversion.

Logg et al. (2019) argue that nowadays people rely on different recommendation systems for their daily life, showing more appreciation than aversion to algorithms, unlike previously stated by Dietvorst et al. (2014). The authors conducted a series of experiments to prove that people would rely on algorithms if given the choice. One of the experiments consisted in asking people to forecast the rank of a song on the Billboard, before and after receiving advice that was said coming from an algorithm. The results confirmed this hypothesis, and also that experts were less open to taking advices. In future research directions, the authors point out to the *black box* problem, where the user does not understand how the algorithm took a decision. Researchers should find a way to make users more conscious on

their reliance on algorithms, to better elaborate their advices and avoid the risk of being manipulated.

Glikson and Woolley (2020) present a literature review on human trust in AI, and propose a framework to study the relationship between humans and AI across its embodiment and capabilities. This framework considers three embodiment levels of AI:

1. **Robotic:** physical presence, like a humanoid robot
2. **Virtual:** visible digital presence, like a chatbot or a 2D assistant
3. **Embedded:** abstract representation, like an algorithm

Existing literature suggests that a tangible and capable AI generates more emotional trust, even if erroneous robots are sometimes liked more than flawless ones. The authors argue that AI will not entirely replace human employees, however a strong skill adjustment is needed to remain competitive. A human-centered integration of AI in organizations will help future leaders to manage teams made of humans, machines or both.

### 2.5.2 BIBLIOGRAPHIC COUPLING

Until now, the aim of the literature review has been understanding what influenced the research on AI and decision making of the last three years, by analyzing 2 043 references of these documents.

It emerged that both technical areas such as fuzzy sets and ML algorithms, psychological areas such as beliefs and aversion, and business areas such as AI applications and analytics heavily influenced the current literature. Now, the aim of the literature review shifts to understanding how the current literature is organized, and if there is a similarity with the research areas of the past.

The bibliographic coupling network displayed in figure 2.11 can be generated in *VOSviewer* by navigating *Create... > Create a map based on bibliographic data > Read data from bibliographic database files > Scopus > Bibliographic coupling > Documents*. Also in this case the counting method did not alter the network layout significantly, so the default full method was selected. Again, only the largest set of connected items was chosen to be shown, and the size of each node has been set to represent its total link strength.

*VOSviewer* identified eight clusters, and assigned each document to one of them. Results are reported in table 2.5. In this case, MCDM (green) and fuzzy sets (orange) clusters are clearly separated from the others. The AI in operations (blue) and AI in decision making (purple) ones are closely tied to each other, while explainable AI (red), robotics and



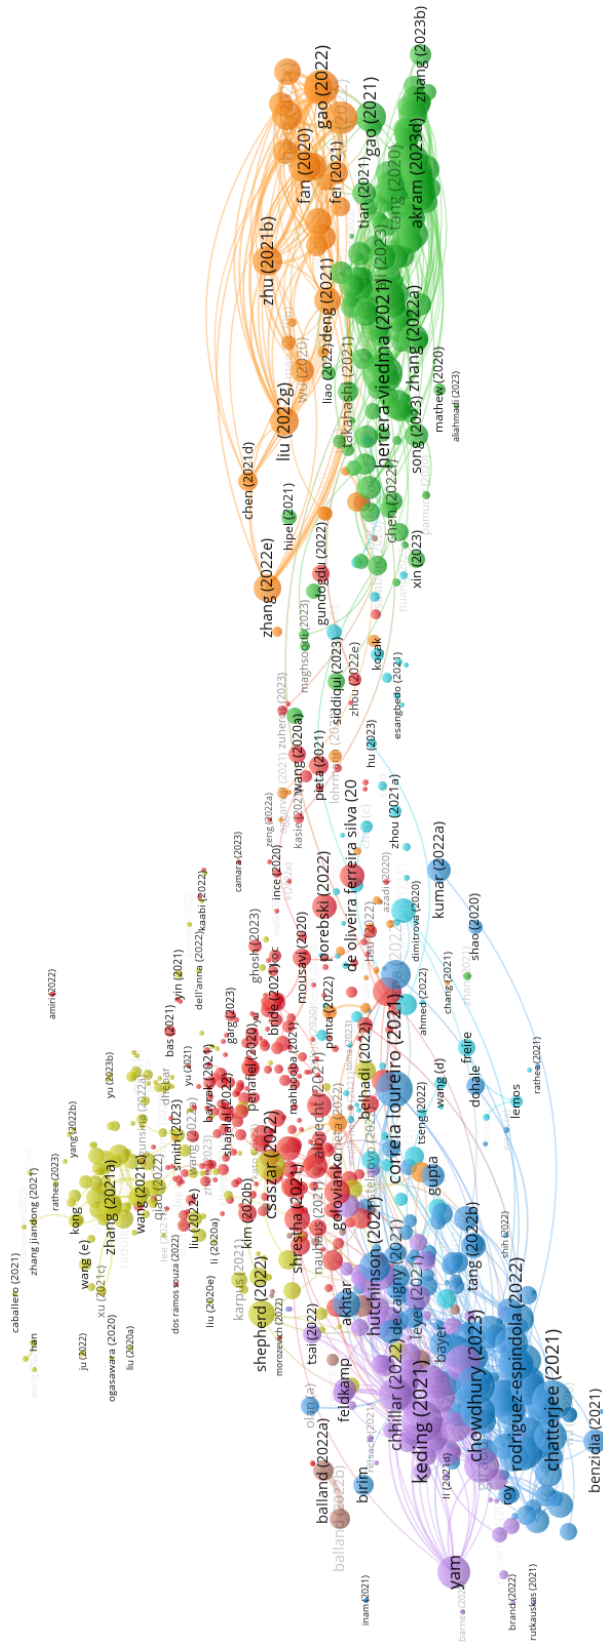


Figure 2.11: Bibliographic coupling of narrow search documents

automation (yellow), AI in HRM (cyan) and AI in innovation (brown), are more loosely tied.

These clusters show a partial thematic overlap with the ones from the co-citation analysis (table 2.3). In particular, fuzzy sets remain separate from the others, as these publications tend to be cited and cite references within the discipline. The most relevant sources for these themes are published by IEEE. Similarly to belief functions in the co-citation analysis (figure 2.8), the MCDM research area is separate from the others, but closely related to fuzzy sets. Technical papers on ML algorithms are cited in documents covering ML (red cluster) and robotics (yellow cluster). The research on business applications of AI diverge according to specific departments (blue, purple, cyan and brown clusters) or industries (red and yellow clusters). Behavioral decision making and aversion towards AI influenced the literature on the applications of AI in decision making (purple cluster), Human Resource Management (HRM) (cyan cluster) and healthcare (red cluster). The research area on analytics has confluent into operations (blue cluster).

**Table 2.5:** Properties of clusters in bibliographic coupling

ID	COLOR	SIZE	RESEARCH AREA	EXAMPLE
1	Red	264	Explainable AI, ML, healthcare	Shajalal et al. (2022)
2	Green	164	MCDM	Bączkiewicz et al. (2021)
3	Blue	155	AI in operations	Benzidia et al. (2021)
4	Yellow	135	Robotics, automation	Qiao et al. (2022)
5	Purple	88	AI in decision making	Keding and Meissner (2021)
6	Cyan	55	AI in HRM	Qamar et al. (2021)
7	Orange	48	Fuzzy sets	Xie et al. (2022)
8	Brown	25	AI in innovation	Krakowski et al. (2022)

With the co-citation analysis, it makes sense to analyze the most influential publications of each cluster to understand the state of the art. However, relying on the mere citation count in this case can be misleading, as the older documents of the co-citation analysis have had more time to receive citations than the more recent documents of the bibliographic coupling. For this purpose, we rely on VOSviewer *normalized citations* attribute, which corrects for this fact, and *total link strength*, which can shed light on thematically similar

publications with a lower citation score.

### FUZZY SETS

Zhou et al. (2021) explained the belief rules-base model for modern ES. Belief rules are able to model complex systems and use both quantitative data and qualitative knowledge to express fuzziness (vagueness), randomness and ignorance. By embedding expert knowledge into rules, the outcomes and reasoning process is fully interpretable. In listing practical applications of belief rules, the authors illustrate an ES used to decide between suppliers, one operating in machine fault diagnosis and one in clinical diagnosis. The authors argue that these programs can be adapted to represent complex systems and decisional situation. Further research in ES should focus on optimizing large scale use cases, integrating an interpretable feature extraction capability, whereas deep learning still represent a black box model.

Xie et al. (2022) propose a novel method to rank alternatives in intuitionistic fuzzy decision making. The authors explain how intuitionistic fuzzy sets introduced by Atanassov (1986) contain both positive and negative information, applicable for example to the efficacy and strength of side effects of drugs in a clinical trial. Eventually, the authors show how their model can be used to choose the best drug for a disease treatment and the best component supplier according to different criteria.

### MULTI-CRITERIA DECISION MAKING

Bączkiewicz et al., 2021 propose a novel recommender system based on the combination of five MCDM methods, demonstrated by choosing the most suitable smartphone. The system is designed to be used as a central component of a DSS for e-commerce websites, and consists of four main stages:

1. **Dataset preparation:** data is gathered from various platforms to identify a list of alternatives and relevant criteria for evaluating the product
2. **Preferences induction:** weights representing user preferences are calculated using an objective or subjective methods
3. **Evaluation of alternatives:** a set of alternatives is evaluated using MCDM methods
4. **Compromise ranking construction:** a recommendation for the most favorable product is provided to the user in the form of a compromise ranking

The final suggestion addresses the issue of inconsistent results that arise from the varying

assumptions of each MCDM method. According to the authors, future research may explore other ways to obtain a compromise solution.

Herrera-Viedma et al., 2021 conducted an extensive review of the trends and developments in fuzzy and linguistic decision making in environment characterized by uncertainty. In explaining the developments in fuzzy set theory, the authors mention influential publications already seen in the co-citation analysis, such as L. A. Zadeh (1965), Atanassov (1986) and Rodríguez et al. (2012). Then, the authors describe three decision making scenarios: MCDM, group consensus-driven decision making and multi-person MCDM. The simplest MCDM framework consists of three elements:

- A finite set of decision alternatives  $A = \{a_1, a_2, \dots, a_n\}, n \geq 2$
- A finite set of evaluation criteria  $C = \{c_1, c_2, \dots, c_q\}, q \geq 2$
- An importance weight for every criterion  $W = \{w_1, w_2, \dots, w_q\}$

Preference information is expressed in an evaluation matrix  $M = (x_{ij})_{i,j}$  where each assessment  $x_{ij}$  represents the evaluation given to alternative  $a_i$  in accordance to criterion  $c_j$ . The matrix is then combined with the weight vector  $W$  to generate a decision matrix and find the best alternative  $a$ .

$$M = \begin{bmatrix} x_{11} & x_{12} & \dots & x_{1q} \\ x_{21} & x_{22} & \dots & x_{2q} \\ \vdots & \vdots & \ddots & \vdots \\ x_{n1} & x_{n2} & \dots & x_{nq} \end{bmatrix} \quad W = [w_1 \quad w_2 \quad \dots \quad w_q]$$

Group decisions often lead to better or less biased solutions, but add complexity to the framework as there may be the need to discuss and reach a consensus. Additionally, it could happen that leaders' opinions may influence the ones of other participants, especially if opinions are communicated sequentially. Multi-person MCDM differs as participants make decisions independently and then provide their decision matrix, which are aggregated to obtain either a unified preference vector or a ranking of alternatives. The authors explain the challenges that effect these and more recent frameworks, such as large scale decision making, recommender systems, crowd decision making and data-driven decision aids. In particular, they call for a real-world validation of the proposed frameworks. The authors attribute this shortage to the lack of real dataset of decision making problems, and to the lack of standard evaluation methodologies to compare models objectively. In their concluding remarks, they state that AI can lead to smarter decision aid tools by enabling personalized

recommendations, extracting opinions from other data and embedding uncertainty and common sense in the model.

### ROBOTICS, AUTOMATION


Zhang et al. (2022) propose a model to recognize the intention of enemy targets in a military context. The model rely on multiple sensorial inputs such as target pictures and trajectory. Then, information is elaborated by a Long Short-Term Memory (LSTM) network to assess the probability of each intention (e.g. attack, scout or withdraw). In parallel, a Generative Adversarial Network (GAN) assess the probability of each target vehicle type (e.g. aircraft carrier, cruiser or patrol ship). Fuzzy reasoning rules are applied to provide the command center a detailed analysis, calculated through the intention probability, the vehicle type probability, their target distance from the observer, and the firing range of their type.

Qiao et al. (2022) examined state-of-the-art research on brain-inspired intelligent robots in areas such as visual perception, autonomous learning, decision making, and musculoskeletal control. Despite being in the experimental phase, considerable progress has been made in replicating biological structures and functions, such as the impact of emotions in decision making. Brain-inspired intelligent robots have the potential to revolutionize robotics by paving the way for the development of new types of robots with human-like intelligence and behavior. The authors highlight the main issues that researchers are facing with their proposed architecture. Computer vision is achieved either through ANN systems and biologically-inspired models. The former focus on extracting features from an image, the other on replicating the structure of the human visual pathway. Decision making is achieved as either reinforced learning (based on rewards and punishments) or Bayesian learning (based on Bayesian inference), and emotion is formulated in these framework through mathematical representations. Motion can be achieved from either model-based and model-free methods. The former are based on simple control theory or optimization, the latter on complex ML or ANN systems.

### EXPLAINABLE ARTIFICIAL INTELLIGENCE, MACHINE LEARNING, HEALTHCARE

Shajalal et al. (2022) ideated an eXplainable Artificial Intelligence (XAI) model to predict product back orders in an inventory management system. A back order is an order for a good that cannot be filled in time. The authors referred to existing explainability techniques such as SHapley Additive exPlanations (SHAP) and LIME to explain which model feature have the most impact on the model predictions and its individual decisions (figure 2.12).

SHAP values have no unit of measure and can assume any real number, however, they represent the marginal effect that the observed level of a feature for an order has on the final predicted back order probability for that observation (figure 2.12). Summing the SHAP values of each feature of a given observation yields the difference between the prediction of the model and logistic regression (Mazzanti, 2020). The model suggested by the authors indicates that inventory level, lead time, and performance in the past 12 months are the most important features when predicting back orders. Explainability techniques allows stakeholders to analyze why a certain product has a high probability of back order in the future, but can also be generalized to other use cases where ML predictions are widely adopted, such as customer behavior, credit worthiness and fraud detection. The authors also plan to develop in a future publication a prescriptive system that suggests the next action to the user.



**Figure 2.12:** Jitter plot of the most important features in predicting back orders

**Source:** Shajalal et al. (2022)

Johnson et al. (2022) proposed a three-stage framework for developing an AI-based decision support system to predict the survival rates of lung cancer patients after diagnosis, using a publicly available dataset. The first stage involved data preprocessing and target creation, while the second stage applied different algorithms with feature selection and hyperparameter tuning. The third stage used permutation importance to interpret the models and gain insights into the relationships among influential features. The authors propose how future research could apply their framework to other types of cancer and

include genetic variables that are considered in real world situations.

### ARTIFICIAL INTELLIGENCE IN DECISION MAKING

Keding and Meissner (2021) studied why and to what extent managers rely on AI advisory systems rather than human recommendations. In this context AI refers to machines performing the various cognitive functions usually associated with human intelligence. The authors highlight how intelligent ES were praised already in the 70s for their potential aid in decision making, but their adoption declined due to the systems' incapacity to deliver these promises. With the subsequent technological advancements and cost reductions, ML predictions become affordable and able to provide accurate forecasts. Recent studies emphasize that humans keep playing the key role of central processors and final authorities in strategic decision makers thanks to their unique sense-making skills. Senior managers are shifting their role from generating solutions to evaluating the ones proposed by machines. To test their hypothesis, the authors gathered 150 senior managers and asked them to evaluate R&D investment opportunities (such as buying the patent for a biofuel production process) that were said to come from either an AI-based advisory system or a team of humans. The results confirm the authors' hypothesis of algorithmic appreciation, indicating that managers tend to rely heavily on AI-based advisory systems when making strategic decisions. The use of AI-based recommendations also increases the likelihood of investment action and positively affects perceived decision quality. This adherence to AI-based recommendations originates from a higher level of trust and a more structured perception of the decision making process, compared to the human one. However, this also represents overconfidence in the machine capabilities, that may not be sufficient enough to perform such a strategic recommendation. The results of the experiment align with the ones from Logg et al. (2019), contrasting with the theory of algorithm aversion proposed by Dietvorst et al. (2014). In their further research proposals, the authors suggest to study how humans can overcome this decisional bias to consciously leverage machines.

Leyer and Schneider (2021) discuss how managers relate to AI in task delegation and augmentation. The authors conducted three experiments to understand how much managers are willing to delegate a strategic decision, how they react to these decision outcomes, and why they would or would not delegate such decisions. The results show that managers are less willing to delegate such decisions to AI than to human colleagues, and the ones that delegated to AI showed milder reactions to both positive and negative decision outcomes. A significant chunk of candidates was not willing to delegate a second time a similar decision

to AI, regardless of the outcome of the previous decision. However, more than 40% of managers who made a wrong decision on their own were willing to delegate to AI. The main reasons to delegate to AI were:

- AI’s superiority in reducing bias
- AI’s potential for workload reduction
- AI’s valuable insights for decision making

The main reasons not to delegate to AI were:

- Confidence in human capabilities
- Desire to remain in control
- Lack of trust in AI
- AI’s limited flexibility to adapt to the decision context

After highlighting the technical differences between humans decision makers, traditional software and AI software, the authors point out some potential benefits and drawbacks of AI adoption in a business context (see table 2.6).

**Table 2.6:** Benefits and drawbacks of AI adoption in a business context

**Source:** Adapted from Leyer and Schneider (2021)

ASPECT	BENEFITS	DRAWBACKS
Task automation	Replacement of human labor	Unlearning of activities or skills, fewer training opportunities for more complex cases, unclear accountability for outcomes
Machine architecture	Efficiency and accuracy, reduction in the working time needed	Cannot question complex algorithms about their predictions or suggestions
Machine reliability	Trust in the tool	Over dependence on the tool, loss of the ability to challenge the technology



### ARTIFICIAL INTELLIGENCE IN HUMAN RESOURCE MANAGEMENT

Qamar et al. (2021) conducted a systematic literature review to understand the state-of-the-art research on HRM applications of AI. The main theme is employee selection, while other topics such as human perception of service robots are covered sparsely. Based on their content analysis, the authors deduced that AI application facilitates or supports decision making in HRM. They divide problems in structured and unstructured. Issues of the first type are agreeable to mathematical models and feature standard solutions, such as predicting employee performance, while members of the latter group do not feature standard solutions and therefore require human judgment to be solved, e.g. selecting the employee benefits to offer. Different methodologies, such as ES, fuzzy sets and ML are used to tackle different kind of problems. The authors propose several research questions articulated in employee perception, workforce management, Return on Investment (ROI) evaluation and leaders' perspectives.

Lemos et al. (2022) applied a multi-criteria decision system to understand what are the most important drivers for AI adoption for Small and Medium Enterprises (SMEs). A panel of six experts was gathered to brainstorm about the drivers of SMEs adoption of AI. A total of 112 criteria were organized in five clusters, which are presented as follows in order of importance. Then, the experts were asked to fill weight matrices and the most impactful factors were calculated:

1. **Human Resources:** lack of practical knowledge, need for experienced professionals, team motivation and upskilling, management of expectations regarding AI applications
2. **Know-How and Knowledge:** research on similar cases of AI usage, incorrect information about AI, lack of clear benefits from using AI, business questions determined via rapid ideation, information sharing within the organization
3. **IT Infrastructure:** data organization, data quality, testing opportunities before investing, digitalization, adoption of existing platforms to reduced development cost
4. **Organizational policies:** funding programs, work in collaborative networks, use of existing standards, difficulty of evaluating the results obtained, development of turnkey AI projects
5. **Leadership:** leaders' knowledge about adaptation processes, commitment and motivation

Therefore, according to the authors, SMEs should prioritize these aspects to facilitate AI

adoption.

ARTIFICIAL INTELLIGENCE IN OPERATIONS

L. Chen et al. (2022) elaborate a framework to apply AI to Business to Business (B2B) marketing. They conduct a systematic literature review to find the most prolific technologies, drivers and outcomes (table 2.7).

**Table 2.7:** Application areas, drivers, outcomes and technologies of AI in B2B marketing

**Source:** Adapted from L. Chen et al. (2022)

APPLICATION	DRIVERS	OUTCOMES	TECHNOLOGIES
Decision support	Customer priority, pricing complexity	Target most important customers, improve customer retention, optimize pricing	DSS, ES, ML
Process automation	Human inefficiency, reliance on intuition, excess information for employees	Improve efficiency, reduce personnel costs	Agent-based systems, ML
Customized service provision	Excess information for customers, limitations of existing recommendations	Increase sales, improve recommendation accuracy	Recommendation systems

The authors identify the following barriers to AI adoption:

- **Lack of awareness, knowledge or motivation:** managers may be unaware of inefficiencies in existing processes, may not have the skills to spot them or may not be motivated to implement such a change
- **Resistance and concerns:** resistance is often caused by ethical concerns and fear of job replacement
- **Business complexity:** there is little data supporting whether AI technologies can solve large-scale problems with real-world complexity
- **Compliance:** privacy regulations may prevent companies from processing customer data freely

They also present an abstract implementation framework that starts from data about the environment and tasks, and ends with business outcomes mentioned in table 2.7. In the future research directions, the authors highlight how other marketing areas are less explored, and that the literature lacks a cost analysis of AI implementation and employee training, other than a benefit analysis with real business data.

Raza et al. (2023) perform a bibliometric analysis to discover the main research topics of ML in SCM, similarly to what Toorajipour et al. (2021) did with its systematic literature review. The authors report how just 10% of the analyzed articles focus on Big Data, suggesting that research has yet to focus on the use of ML tools for handling BD in SCM. The clusters they identified in their analysis focus on five themes similar to the ones proposed by Min (2010): supplier selection, sustainability, demand forecasting, inventory management and decision making. They noted that few authors account for the most influential works, and ML aspects are treated at a superficial level, as articles are more directed toward BD analytics in general and not the specific aspects of ML.

### ARTIFICIAL INTELLIGENCE IN INNOVATION

Krakowski et al. (2022) built on top of the Resource-Based View theory to examine the influence of AI on the competitive advantage of companies. The authors studied the outcomes of several chess games where AI engines replaced or helped humans, and eventually proved that human cognitive abilities still play a role in these situations. Highly skilled players would rather remain in legacy niches that are no longer competitive rather than lose against chess engines. The authors argue that domain-specific capabilities in business contexts are plentiful and much more interdependent than in simpler situation such as chess, implying that AI is likely to substitute some, but not all of business activities. According to the authors, companies could either invest in employees' upskilling to develop complementary abilities, or deploy them alongside AI experts to generate competitive advantage by combining their skills.

Rajagopal et al. (2022) conducted an exploratory study to understand the frontiers of AI on business decision making. In recent times, enthusiasm around the potential of AI technologies has been fomented by consulting and technology companies, and stakeholders have great expectations. The authors recall different aspects of strategic decision making, suggesting that AI has the potential to offload managers from analytical tasks, allowing them to imagine and explore new opportunities. They also propose a conceptual model for AI implementation, a process that depends heavily on the peculiarities of each case.

## 2.6 DISCUSSION AND GAP ANALYSIS

THE literature on data-driven decision making is prolific and has covered a wide variety of topics since its inception (figures 2.2 and 2.3). Many topics come back recurrently, as in the case of high expectations from AI (Duan et al., 2019; Russell & Norvig, 2020) and the fear of job substitution (Lawler & Elliot, 1996). The technological innovations and computing performance improvements of the last decades made possible to analyze large quantities of data and perform complex computations, increasing the research interest in field such as ML and BD Analysis. During the last years, people seem to have moved from algorithm aversion (Dietvorst et al., 2014) to algorithm appreciation (Keding & Meissner, 2021; Logg et al., 2019).

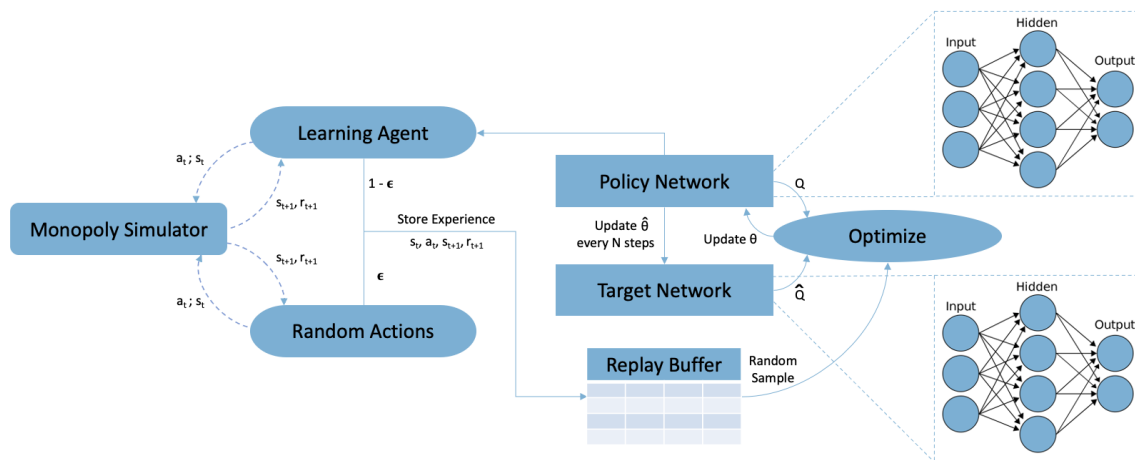
The literature on AI in organizational decision making seems split in two (figures 2.8 and 2.11): a more technical branch focuses on fields such as fuzzy sets and MCDM, while a more social branch focuses on applications of AI technologies such as ES and ML in different industries or functional areas. The various attempts to classify the topic of AI (Duan et al., 2019; Russell & Norvig, 2020; Schmitt, 2023) allowed to understand the shift from the rule-based approach, adopted in ES, to the example-based one, adopted in DL. However, no classification of AI has been standardized or is universally accepted.

Many technical articles in the literature introduced new ML models (Bączkiewicz et al., 2021; Dandolo et al., 2023) or discussed the application of existing models to available datasets (Nguyen et al., 2021; Zhang et al., 2022). Theoretical publications instead proposed adoption frameworks (Rajagopal et al., 2022), technology roadmaps (Davenport et al., 2019; Huang & Rust, 2018; Maoz, 2013) and hypothetical use cases of AI models (Keding & Meissner, 2021; Min, 2010). In decisional situations where datasets are not existing or available in tabular format, the authors adopted an unsupervised learning approach and developed a DL agent (Bonjour et al., 2022; Silver et al., 2016). In these cases, the real life situation is represented digitally through modeling and sometimes simplification. Then, the agent is trained via self-learning, and improves its performance at every iteration.

It is widely claimed that AI can support decision making (Bornet, 2022; Davenport et al., 2019; Schmitt, 2023; Weber, 2023; Ylijoki & Porras, 2018), but few articles acknowledge that not all decisions are alike (Bonjour et al., 2022; Edwards et al., 2000; Shrestha et al., 2021). Among other factors, organizational decisions may vary in complexity, frequency and importance (Mintzberg, 1979).

An enlightening example of applying state-of-the-art AI technologies to a complex

business decision making scenario featuring a long term objective, different decision types and relationships with other decision makers, is presented by Bonjour et al. (2022). By illustrating the modeling process of their Monopoly-playing DL agent (figure 2.13), the authors explained the challenges they faced and the shortcomings of existing models: some decisions are much more complex than others, and some are less frequent. Decisions of the first kind are heavily simplified in the model compared to real life, while decisions of the second kind are addressed outside of the DL model.



**Figure 2.13:** Architecture of a Monopoly-playing Deep Learning agent using both a rule-based and an autonomous learning approach

**Source:** Bonjour et al. (2022)

To take less frequent decisions, the authors opted for a rule-based approach, choosing what they considered as the best practice, resulting in a fixed policy of always buying a new property and accepting trades that lead to an increase of monopolies. Complex decisions, such as proposing the price to buy a property from another player, have been discretized to only consider three values (below market, above market and market level). This hybrid approach eventually led to a faster win rate convergence compared to both traditional DL and fixed-policy agents, simplifying at the same time the model complexity of the game. The authors state the following: “Evidently, instead of letting the agent explore the rare state-action pair it may be better suited if these are replaced by rule-based logic, especially if we know what actions might be good in the given state.” This process effectively shows that having a well-defined goal (March & Simon, 1958; Min, 2010), a structured model of the problem (Russell & Norvig, 2020; Siegel et al., 2020), and handling different types of decisions accordingly (Edwards et al., 2000; Mintzberg, 1979), are necessary to excel in

decision making.

Higher-level decision making, such as strategic business decisions (Mintzberg, 1979), is difficult to frame in a conceptual model. By its very definition (Merriam-Webster, 1831), a model is a simplification that only captures the most important aspects of a real-life situation. Even the concept of *digital twin*, a complex model that aims to replicate as many features as possible of the real-life situation for simulations, performance monitoring and predictions, presents limitations (Yan et al., 2022). These models are constrained to the variables and assumptions specified by the designer's bounded rationality (Simon, 1957) and cannot account for information that cannot be captured digitally (Mintzberg, 2015).

As any other technology, the application of ML to decision making presents benefits and drawbacks, both intrinsic and extrinsic. Among the most mentioned benefits, there are automatic feature extraction, high prediction and classification performance, bias avoidance, potential for workload reduction (Leyer & Schneider, 2021) and DL scalability to large datasets. The most common downsides instead include the need for training on large datasets, the presence of knowledge acquisition bottlenecks, difficulty in handling uncertainty and novelty, limited applicability to complex environments (Min, 2010), limited human understanding of machine decisions (Barredo Arrieta et al., 2020) and limited machine replacement of human intuition (Huang & Rust, 2018). In addition to the intrinsic properties of these technologies, there are external factors that need to be considered: employee training (Krakowski et al., 2022; Lemos et al., 2022), overconfidence in human capabilities (Kahneman, 2011; Leyer & Schneider, 2021; Sibony et al., 2010), the need for well-defined problems to solve (Toorajipour et al., 2021), the need to remain accountable for machine-based decisions, unlearning due to overreliance (Leyer & Schneider, 2021), the need for external explainability techniques (Shajalal et al., 2022).

These properties are quite different from the ones of ES and other rule-based systems, which have represented the leading AI approach for many decades (Duan et al., 2019; Jordan, 2019; Russell & Norvig, 2020). ES were praised for the explainability of their decisions, but failed to deliver their promises when applied to more complex situations, as it was difficult to code human knowledge into *if-then* rule systems, and computationally expensive to scale these programs to large issues (Lawler & Elliot, 1996).

Recent literature calls to examine what functional areas and decision making levels can benefit the most from DL and more generally from AI technologies (Shrestha et al., 2021). From an overview of the study cases analyzed in the literature, it emerges that ML strengths can be leveraged for frequent, data-intensive and well-structured situations, where decisions

are mostly analytical, and the possible choices are limited. Common example are fraud detection in financial transactions, demand forecast, preventive maintenance, customer classification and inventory optimization (Davenport et al., 2019; Russell & Norvig, 2020; Shajalal et al., 2022; Shrestha et al., 2021; Siegel et al., 2020).

However, there seems to be little applicability of these technologies to strategic decision making, where situations tend to be less frequent, less data-intensive and less structured (Pietronudo et al., 2022). According to (L. Chen et al., 2022), there is not enough evidence on whether AI technologies can solve problems with real-world complexity.

## 2.7 RESEARCH QUESTION

By acknowledging that not all organizational decisions are alike, and separating them according to their properties, we aim to obtain a more comprehensive and truthful picture of AI applicability to decision making. Table 2.8 illustrates a classification of organizational decisions, whose main principle is taking into account that day-to-day decisions are typically decoupled from the overarching strategies. This framework has been developed by expanding on the previous work of Bonjour et al. (2022), Edwards et al. (2000), Mintzberg (1979, 2015), Pietronudo et al. (2022), Russell and Norvig (2020), and Shrestha et al. (2021).

**Table 2.8:** Properties of strategic and operational decisions

	STRATEGIC DECISIONS	OPERATIONAL DECISIONS
Focus	Long-term	Short-term
Importance	High	Low
Frequency	Low	High
Complexity	High	Low
Structurability	Low	High
Dependency on data	Low	High

This classification must not be intended as a dichotomy, but rather as the extremes of a spectrum of some decisional properties. Strategic decisions typically have a long-lasting impact on organizations, while the same is not true for more operational decisions. The individual importance of each decision is higher at the strategic level, but the routine of operational decisions makes their impact also valuable, considering a lengthier time span.

Strategic decisions tend to be more complex as they are affected by multiple factors, some of which may be difficult to quantify or structure in a model (e.g. bargaining power of buyers). The range of possible solutions may be wide or undefined (e.g. what product to launch next, where to open a new production plant), and the selection process may be articulated in multiple stages. On the other hand, operational decisions are more standardized and documented. In general, we assume that importance is inversely correlated with the frequency of these decisions, and that complexity is inversely correlated with their structurability. While some decisions can be taken based solely on the available data, we consider most of them to be more operational. When it comes to more strategic decisions, other factors influence the decisional process, including culture, context, and values. Furthermore, the classification of each decision depends on the context of the organization, so no decision is inherently strategic or operational *per se*.

Considering this framework, we aim to answer the following research question:

**RQ 1.** Is AI more suited for operational rather than strategic decision making?



# 3

## REAL CASES OF AI EMPLOYMENT IN DECISION MAKING

**W**E aim to collect empirical evidence from large manufacturing organizations to address the research question proposed in section 2.7. In this chapter, we implement the empirical process illustrated in figure 1.4. Section 3.1 will introduce the methodology adopted to carry and analyze the interviews with companies. In particular, section 3.1.1 will list the interview questions derived from RQ 1, section 3.1.2 will explain the filter criteria of the interview sample, applied to ensure comparability. Section 3.1.3 will briefly introduce the interviewed companies. Section 3.2 will instead collect and organize the answers received during the interviews, which will be analyzed and discussed in section 3.3. In the same section, a clear answer to the research question will be provided. Section 3.4 will summarize the research and present the main findings. To conclude, it will acknowledge the limitations of the study and propose some future research directions.

### 3.1 RESEARCH METHODOLOGY

**P**RIMARY data to answer these research questions can be collected from surveys, interviews and experiments. We consider that having individual discussions with high-stake company representatives is the most authoritative way to discuss real AI applications in organizational decision making. This methodology offers more valuable insights than multiple choice questionnaires (Lamarre et al., 2023), and is based on existing use cases rather than hypothetical situations like controlled experiments (Keding & Meissner, 2021).

### 3.1.1 INTERVIEW QUESTIONS

We want to address RQ 1 by splitting it in four intermediate questions:

- RQ 1.** Is AI more suited for operational rather than strategic decision making?
- a. For what decisions are companies currently employing AI?
  - b. What are the main factors driving AI employment in decision making?
  - c. What are the main barriers hindering AI employment in decision making?
  - d. Can we overcome the barriers of AI employment in decision making?

This section will explain the interview questions related to each intermediate question.

To answer RQ 1a, we will individuate a list of decisions and activities for which interviewed companies are currently employing AI. The emphasis on *current* employment rather *possible* employments will avoid wishful thinking on possible applications, maintaining the focus on the present situation (Raza et al., 2023). Possible or planned future employments of AI are addressed later in a following set of questions. The framework in table 2.8 will be used to separate these decisions into strategic and operational, and collect them in table 3.2. If we cannot reasonably identify a strategic or operational use case within the company, we fill the cell with Not Applicable (N/A). This will allow the comparability of answers and highlight if AI is employed for more operational or strategic decisions. The following questions will be asked during the interviews (section 3.2.1):

- For what decisions are people in your company currently employing technologies associated to Artificial Intelligence, including Machine Learning?
- What roles do these technologies play? How are they being applied?
- How would you define these decisions in terms of frequency, complexity and importance for the company business?
- Is your company using AI solutions in decisions such as pricing, launching a new product or choosing a potential acquisition?

To answers RQ 1b, we will uncover the main reason that led to the adoption of AI solutions for the arose decisions. These reasons may be related to the properties of the issue (such as data-intensity) or of the technology (such as forecast accuracy), also compared to existing alternatives (Leyer & Schneider, 2021). The following questions will be asked during the interviews (section 3.2.2):

- Why did your company decide to employ AI for the decisions you said before?
- In these situations, what are the main advantages of the AI solution over existing or alternative solutions?

To answer RQ 1c, we will understand what are the main challenges that prevent companies from applying AI to other kind of decisions. These factors can be intrinsic in the technology or external in the company environment (Lemos et al., 2022), and make the adoption impossible or less convenient. These answers provide realistic boundaries that companies willing to adopt AI for decision making need to take into account. The following questions will be asked during the interviews (section 3.2.3):

- What is hindering your company from applying AI to other types of decisions?
- How do you take these decisions without relying on AI technologies?

To answer RQ 1d, we will collect opinions on possible improvements to AI employment to decision making (Helo & Hao, 2022), referred to both the technology itself and the company roadmap. We also plan to propose and discuss potential use cases relevant to the company. This allows to better understand how the company would implement such a system, or why it would not consider it feasible. The following questions will be asked during the interviews (section 3.2.4):

- Do you think it is feasible to address these challenges, or they will persist also in the future?
- Given the theoretical feasibility, do you think it is practical and convenient to address them?
- Are you currently addressing these challenges within your company?
- Consider implementing a system that does X. Would it be feasible? Why?

#### 3.1.2 COMPANIES SELECTION

To ensure quality of the research through comparability of the results, we set some criteria to filter the companies we will reach out to. We decided to focus on mostly manufacturing companies, that adopted AI for decision making only in a later stage, for two reasons:

- To better understand the impact of the *introduction* of AI
- Companies with a data-intensive business model are already widely covered by the literature (Butters, 2023; Garg et al., 2017)


However, to enrich the results we also wanted to include at least a consulting and a company specialized in BI, as they can offer insights on the impact of AI in other companies. Since the company identity is not relevant for the analysis, we omit information that can identify the companies or their spokespeople.

To include the discussion of technical aspects, we opted to interview for technical executives, and senior managers in charge of AI and analytics operations. Representatives of the chosen companies were contacted via cold calls, cold emails and LinkedIn connection requests. We asked for interest in participating in our research on the impact of AI technologies in decision making. In case of a positive response, the four RQ were shared to them, and a video call was arranged in the following days. By sharing the questions in advance, we allowed the representatives to prepare for the interview, resulting in more insightful answers. In addition, they can name someone else from their company if they consider him better suited to answer the questions. However, the framework in table 2.8 was not shared in advance, to avoid altering the answers received. Interviews were carried out in English and Italian, according to the preferred language of the interviewee. The conversations were recorded, paraphrased and split in different sections to better answer the research questions one by one.

#### 3.1.3 LIST OF SELECTED COMPANIES

In total, we reached out to 94 companies. Of these, 78 did not answer, two refused to participate because they were not employing any kind of AI, and four refused for other reasons. Only ten of the original 94 companies accepted to schedule an interview. This is represented graphically in figure 3.1.

In the following section, we will briefly present the companies. Table 3.1 serves as a summary of various data collected from the companies that were interviewed. These data points are crucial in identifying and classifying the type of company interviewed, all while ensuring the confidentiality of the company identity. *ID* is an arbitrary identifier for each company, *Industry* represents the main business or industrial sector, *Employees* is a  $10^n$  approximation of the workforce size, and is used as a proxy for the company scale and complexity. *Spokesperson title* represents instead the job title and main responsibilities of



**Figure 3.1:** Sankey diagram of companies invited to the interview

the interviewee. Geographical data such as the headquarters country was not collected to increase anonymity, but all interviewed companies are either US or EU-based.

Seven of the ten interviewed companies manufacture goods, while three offer services to other companies (C2, C6, C7). C1 is a company producing coffee makers for bars, restaurants and hotels. It is undergoing a long-term digitalization process, which involves direct interaction with customers and collection of data from the sensors located inside their appliances. The Chief Operating Officer (COO) supervises this project and has a deep knowledge of the technologies employed in their products and within the company. C2 is a global consulting firm specializing in management and technology solutions, with a focus on improving business performance and driving innovation. The spokesperson is the Managing Director responsible for the data science and AI practice in Italy. C3 is widely recognized for producing vehicle components and offering travel related services to both institutions and end users. The Chief Innovation Officer established long-standing partnerships with leading technology universities and startups to research bleeding edge technologies. C4 is an European company producing semiconductor devices for both manufacturing companies and residential customers. The interviewee is responsible for the silicon foundry and strategic decisions such as make or buy, monitoring production performance, and allocating production budget. C5 is a multinational technology company that designs and manufactures hardware and software for both professional productivity and gaming purposes. We interviewed a general manager in charge of data-driven and AI-driven transformations of corporate clients. C6 is a leading provider of workplace financial wellness solutions, offering unbiased financial education and personalized guidance as employee benefit. The representative is an experienced manager in charge of developing

**Table 3.1:** Properties of interviewed companies

ID	INDUSTRY	EMPLOYEES	SPOKESPERSON TITLE
C1	Coffee makers	100	Chief Operating Officer
C2	Management Consulting	1 000 000	Country Director of AI
C3	Automotive components	100 000	Chief Innovation Officer
C4	Electronic components	10 000	Senior Operations Manager
C5	Hardware and software	100 000	Senior AI Manager
C6	Financial coaching	100	Director of AI
C7	Analytical software	1 000	Senior BI Manager
C8	Blast chillers	100	Senior Product Manager
C9	Alcoholic beverages	1 000	Director of Analytics
C10	Sports cars	1 000	Director of Analytics

AI products for the company and its customers. **C7** is a supply chain risk management company that provides its clients a platform to streamline supplier prequalification and contractor management. It also helps customers identifying gaps, inconsistencies, and areas for improvement in their risk management and regulatory compliance. The spokesperson is a manager with more than 25 years of experience in internal and external BI implementations. **C8** is a manufacturer of professional blast chillers, shock freezers, and holding cabinets for the food service industry. The company representative is a senior product manager supervising the R&D department. **C9** is a global alcoholic beverage company known for its iconic brands and diverse portfolio. The director of data and analytics we interviewed offered various insights on the undergoing implementation of the company long-term data strategy. **C10** is a renowned luxury sports car manufacturer, whose business also articulates in lifestyle fashion and racing divisions. The spokesperson is in charge of the data science and business analytics department, and supervises the internal application of AI solutions.

## 3.2 INTERVIEWS RESULTS

**E**VEN if companies were interviewed individually on a series of questions, we found more practical to group the answers based on a particular topic, rather than on a company basis. Interviews were structured to introduce the research goal, ask the companies for the decisions where they are currently employing AI, the reasons behind this adoption, the challenges in extending the technology to other type of decisions, and personal considerations on the mitigation of such challenges.

### 3.2.1 CURRENT AI EMPLOYMENT

In this section, we collect the list of decisions for which interviewed companies are already employing AI (RQ 1a). As part of our research, we focus only on current employment to have a more realistic representation of how these technologies are employed to take decisions.

At the end, we summarize these answers in table 3.2 as related to either more strategic or operational decisions. This separation is based on the properties of the decisions that emerge during the interview.

This list of decisions will be recalled in to discuss why AI is employed for each case.

**C1** clarified that their main digitalization goals are to uncover new market opportunities, enhance existing products, and improve their corporate communication on social media. They are collecting usage data through sensors in components like solenoid valves, and they analyze this data with ML algorithms to comprehend the factors influencing the values. Based on their analysis, they either introduce new product features such as predictive maintenance, or make improvements to existing products. This process also allows the company to align its value proposition with customer needs, preventing the inclusion of unnecessary features that would increase costs. In their social media communications, the company focuses on the most discussed product features and promotes proper usage practices to ensure optimal sustainability of the product. However, to understand the most interesting features for their clients and end user they rely on focus groups and workshops, and not on social media analysis.

**C2** is mostly using AI internally to automate back office tasks, rather than supporting decisions. According to the AI Director for Italy, these decisions are always taken by domain experts, after risk-base assumptions are made and analyzed. Consultants in the company

often advise clients on digital platforms that embed some sort of AI, for instance a Customer Relationship Management (CRM) solution that can infer customers' loyalty from their number of previous suppliers, quality-price positioning and sociodemographic data.

Within his digital service division, C3 primarily utilizes ML through two trained models. One of these models is a CV solution, capable of identifying objects, obstacles, and road cracks from videos taken with dashboard cameras. The second model analyzes telemetry data from vehicle sensors to identify areas of risk in roads. These findings aid road maintainers in making safety-related decisions.

C4 does not utilize AI for high level managing processes, but rather for conducting optical inspection and workforce scheduling. The purpose of optical inspection is to ensure quality control during the manufacturing process of semiconductor devices, by determining which products meet the quality threshold. This CV model was trained using labeled pictures of acceptable and faulty products manufactured in the past three years. On the other hand, shift scheduling is carried out on a daily basis to determine the most optimal allocation of available workers and machineries. The trained model was provided with data regarding the production output of a shift, the operators and equipment present, and each operator ability to use these tools. In the training phase, the model identified enhancements in output resulting from changes in the inputs. In the employment phase, it calculates the likelihood of producing more wafers with a specific allocation of workers and equipments. Initially, this model was implemented just for a limited amount of products, but its usage is been expanded due to positive outcomes.

During the interview, C5 spokesperson mentioned that their Financial Planning and Analysis division has been utilizing a proprietary ML model for already two years to predict revenues. Additionally, the company extensively develops and employs other AI products to enhance employee productivity, even before they are made available to the market. An example is an AI extension to their collaboration platform, which is able to provide an executive summary of hour-long video meetings, highlighting the identified channels and action items.

C6 has recently created an internal LLM in order to enhance the retrieval of information from their exclusive knowledge base. The company financial coaches experienced faster decision-making processes as a result of this productivity enhancement, as the model can summarize documents and answer their questions. They have intentions to make this tool



available to their Business to Business customers, who have expressed interest in it and would like to train the model on their own resources, such as HR information, R&D data, and industry-specific data that is not readily accessible.

C7 primarily leverages AI as tools to boost productivity for software developers and employees in charge of auditing customers' operations. On the customer side, AI is predominantly employed for performing natural language searches on data, rather than implementing predictive or prescriptive analytics solutions. Moreover, the company is arranging an internal hackathon to enable employees to suggest their preferred applications of AI.

C8 incorporates ML into their products by implementing an automated defrosting procedure. Every chiller assesses its usage over a span of one week to identify the period with the least amount of activity, and then performs in this time frame the periodic defrosting. By autonomously determining the most suitable timing for maintenance tasks, the device eliminates the need for user intervention and ensures uninterrupted service. While this process is achieved through predetermined rules, it is adjusted automatically to suit the specific usage patterns of each client. In addition, the company is offering its customers a subscription-based application to monitor and control their appliances. This service represents a complimentary source of revenue, and also doubles as a data source for telemetry. The products' sensors are continuously monitored through Internet, facilitating the detection and explanation of product malfunctions. Notably, the company acknowledges that analyzing customer feedback and reviews would be valuable for determining implementation approaches and methods during design discussions.

Even if C9 already reached a comprehensive level of business reporting, it is not using data extensively to make decisions. The director of analytics admitted that the company recognized the importance of data only in the last few years. This did not happen before because data is not part of the core business, which remains alcoholic beverages. On the other hand, industries such as banking, insurance, credit, pharmaceuticals, and retail have leveraged data the most, as they have been dealing with a large volume of information since their establishment. The company is undergoing the implementation of a supply chain planning solution to optimize the production and delivery of their goods. This solution utilizes a weighted moving average algorithm to predict the demand for each good in a given location and time bucket. However, all the steps required to convert this demand forecast into actual production and stock transfer orders involve rule-based calculations. Among other factors, these calculations consider and optimize the available production capacity

for each plant and the ideal stock levels for each location. Decisions that C9 considers strategic are taken through extensive ad-hoc analyses. Often the company relies on external agencies to conduct various analyses such as price elasticity, volume impact, long-term consumer impact, sustainability, and new market opportunities. In the future, AI will assist the company in connecting the dots, but it will not be sufficient to replace all these agencies.

Despite not focusing on autonomous drive vehicles, C10 is actively investing in a data-driven strategy to make its production processes more efficient. The company has a small team dedicated to analytics and data science and they have recently started to hire engineers to expand their know-how in ML. The company has conducted internal investigations to identify areas where AI can bring improvements by targeting specific issues. Their marketing and communication departments are employing CV models to automatically label photos and videos based on their content. After being labeled, these media are used to train new ML models. GenAI is being adopted by the finance, legal, and communications departments to retrieve, summarize and write documents. Text-to-image and text-to-3D models are widely employed by the design department to speed up creative workflows. On the other hand, the R&D and manufacturing departments make use of CV for chassis anomaly detection, and predictive maintenance to avoid equipment failures. Different business users are also starting to use the CRM through extensions that automatically interpret data and provide insights, for instance “This variable increased by X% over the last year”, or “Model X has the biggest market share”. Additionally, they are incorporating ML into their product features, including predictive braking and virtual sensors.

**Table 3.2:** Current AI employment in decision making

ID	AI EMPLOYMENT IN MORE STRATEGIC DECISIONS	AI EMPLOYMENT IN MORE OPERATIONAL DECISIONS
C1	N/A	Product failures analysis
C2	N/A	Robotic process automation
C3	N/A	Road risk analysis
C4	N/A	Quality control, shift scheduling
C5	N/A	Revenue forecasting, software development, content creation
C6	N/A	Financial advisory
C7	N/A	Documentation auditing, software development
C8	N/A	Product failures analysis
C9	N/A	Demand forecasting
C10	N/A	Media labeling, predictive maintenance, car performance analysis, quality control, content creation

### 3.2.2 DRIVERS OF AI EMPLOYMENT

In this section, we collect the main factors that led interviewed companies to the adoption of AI for the decisions listed in table 3.2. We expect to talk about extrinsic properties related to the situation, and properties intrinsic in the technology (RQ 1b). The advantages of AI to potential use cases were not included as not relevant to current usage of AI. At the end, we summarize these answers in table 3.3.

As stated by C1 spokesperson, the main benefit of utilizing ML algorithms to determine whether to enhance an existing product or introduce a new model, is the ability to analyze the big amount of data collected. With the capacity to make accurate predictions based on past data, it becomes possible to implement preventive maintenance features in their products, thereby preventing downtime in the core operations of their end-users.

Despite the lack of examples showcasing AI's influence on C2 decision-making, accord-

ing to the Italy Director of AI, modern AI excels in three domains:

- **Data Analysis:** ML algorithms can be employed to forecast the future and estimate the likelihood of our action outcomes
- **Natural Language Processing:** NLP has the capability to interpret unstructured data from various sources, including social networks. It can perform important tasks like analyzing sentiment, generating maps of common topics, provide insights from customer care inquiries, and identifying frequently mentioned products in a company or industry.
- **Feature extraction:** CV makes possible to identify various components within images and videos, which has proven beneficial in projects involving roads and buildings

C3 is employing ML in their service offer primarily due to the exceptional accuracy of their models in identifying the objects they were trained on, as well as their ability to learn new objects such as cones, barriers, and obstacles. Unlike rule-based systems, which struggle to handle exceptions caused by variations in camera angles or unexpected environmental conditions, ML models can maintain the same level of accuracy even when faced with exceptions like a moved or mud-covered cone. This has been achieved by employing the model on additional

C4 utilizes both rule-based algorithms and ML for pattern recognition in quality control. ML algorithms have an advantage over traditional ones because they can be trained using examples of known good or bad outputs. This eliminates the need to manually code every test for the optical inspection algorithm to identify failures, such as delta comparisons with a germanium wafer representing the golden standard. The DL model can be trained with labeled production examples from past years to recognize most of the situations that can happen. Instead, the low accuracy of traditional algorithms often required manual inspection to confirm the analysis results. Another factor to consider is that traditional algorithms cannot improve their accuracy based on further training or human feedback, as they require manual modifications of the rule set. The company aims to achieve such high accuracy in quality control that rules no longer need to be hard-coded, or people employed for the task.

Regarding shift scheduling, the company uses over 800 different pieces of equipment that can be combined for different production processes, and therefore are not organized in separate production lines. The scheduling of equipment is based on production demand

and available workers. The ANN model currently employed optimizes resource allocation based on available tools, workers, and production orders. This is achieved by training the model to identify patterns where using different equipment combinations improved throughput. Then, the model categorizes inputs and provides results, similar to a lookup table. Every day, the shift scheduler enters the resources available for production, and the model suggests the optimal configuration. This approach is faster and less resource-intensive compared to calculating everything on the spot.

AI adoption for C5 revenue forecast was driven by its superior performance (time and accuracy) compared to humans. This not only resulted in a significant time reduction from 3 weeks to just 30 minutes, but also automated a process that was less valuable for humans and difficult to analyze. This time saving allows the Chief Financial Officer (CFO) to effectively react to market changes and make important decisions. GenAI instead sees wider adoption as it can create a strong initial draft of emails or job descriptions, avoiding the need to start from scratch. People can therefore save time and focus on tailoring the document to their specific situation, a task that the model cannot handle.

C6 is employing a LLM because nowadays this technology can emulate human reasoning and understanding well enough to extract summaries and give specific answers to novel questions.

A BI manager at C7 said that even if AI products can suggest software developers ideas they did not consider, they cannot entirely replace the developer role. In a similar way, the company is employing tools that facilitate the review of clients' operating manuals, but human judgement is always needed to validate the veracity and completeness of the results. In the audit process a LLM may hallucinate, overlook certain errors or falsely flag non-errors. The spokesperson additionally stated that an internally created chatbot yielded unsatisfactory outcomes, as it failed to adequately answer the specific questions that were asked.

The early motivation behind C8's decision to incorporate the automatic defrosting feature was to ease the burden of maintenance for the users. The initial response to this feature has been favorable, prompting the company to explore the development of additional intelligent functionalities, such as automatically adjusting settings to enhance performances, and a cloud-based monitoring of appliances. When the products are connected to the internet, clients can remotely oversee their chillers and perform operations. The company

can also benefit from this situation by collecting a greater amount of data than they could from internal laboratory testing. This wealth of data will is then analyzed to more accurately identify the causes of early failures in real-world situations, such as customer misuse, power surges or faulty components.

Nowadays, C9 is performing demand forecasting through a time series ML algorithm embedded in their supply chain planning software. This solution offers superior performance and accuracy compared to the manual and time-consuming calculations that were previously performed on spreadsheets. Advancements in ML and NLP made possible to analyze qualitative data in a more structured way.

According to the Analytics Director at C10, GenAI is more helpful than traditional AI applications due to its versatility in different workflows and its ability to process unstructured input.

**Table 3.3:** Drivers of AI employment in decision making

ID	DRIVERS IN MORE STRATEGIC DECISIONS	DRIVERS IN MORE OPERATIONAL DECISIONS
C1	N/A	Availability of telemetry
C2	N/A	Accuracy in forecasting and pattern recognition, capacity to generate insights from natural language
C3	N/A	Availability of telemetry and dashcam videos, accuracy in identifying objects and areas of risk, capacity to improve from additional examples
C4	N/A	Availability of historical production data, accuracy in pattern recognition, capacity to improve from additional examples or human feedback, avoidance of coding rules, faster results
C5	N/A	Availability of financial data, accuracy in revenue forecasting, possibility to automate the forecast process, faster results, scalability to other data sources, capacity to draft documents
C6	N/A	Availability of an internal knowledge base, capacity to summarize and answer questions
C7	N/A	Possibility to automate the audit process, capacity to suggest methods to software developers
C8	N/A	Availability of telemetry
C9	N/A	Availability of historical demand of goods, accuracy in demand forecasting
C10	N/A	Availability of multimedia content, sensorial and production data, versatility of GenAI

### 3.2.3 BARRIERS TO AI EMPLOYMENT

In this section, we collect the main challenges faced by interviewed companies in employing AI in other kind of decisions (RQ 1c). We expect to talk about the intrinsic limitations of AI technologies and the implementation issues that arose in the companies. At the end, we summarize these answers in table 3.4.

According to the COO of C1, it is now simple to collect a large amount of data, but the real challenge lies in determining which information can be leveraged and which one is redundant or useless. To be useful, quantitative values coming from individual components must be combined with qualitative insights from markets and external environments that represent the bigger picture. While usage metrics can provide tangible means of monitoring performance or indicating areas for improvement, they alone cannot explain the cause of anomalies.

For example, they cannot determine whether an abnormal pressure level is the result of water impurities obstructing the pump, or a damaged internal pipe. The company is therefore conducting workshops and focus groups with coffee roasters and bartenders to collect these kind of information.

The AI lead at C2 claimed that currently only few AI systems are currently used by companies. This is due to three primary factors:

- **Cost:** these systems are difficult to build as they require a reliable data pipeline and experts to develop it
- **Specificity:** AI systems are specific and intricate, even in the case of more plug-and-play solutions like AutoML
- **Explainability:** even if this issue is not receiving central attention, it poses a major challenge for the organizational employment of these technologies: companies must understand and explain why the algorithm returned a specific output

For these reason, rule-based systems are still the most used within companies.

C3 is offering a pioneer AI solution able to identify areas of major risk in a road. The Chief Innovation Officer argues that this can be seen as a *descriptive analytics* system according to the Gartner Analytics Ascendancy Model (GAAM) framework (Maoz, 2013). After contextualizing the safety results from the model with factors like rain, night and rush hour, and assessing the impact of possible risk mitigations, an analyst provides actionable



prescriptions to the road authority, such as reducing the speed limit or deploying a police officer. Despite its effectiveness, there are several reasons why this solution cannot fully replace the human analyst, effectively achieving the *prescriptive analytics* stage:

- **Legal liability:** the client may be held accountable for accepting automated recommendations without performing appropriate due diligence
- **Human parity:** current technology has yet to reach the level of human reasoning due to the highly complex nature of our processing abilities
- **Customer trust:** clients would find it hard to believe that a software can advise in civil engineering and road design

Similarly, people still expect the doctor to make the final decision and the pilot to control the airplane, even if technology is offloading more and more tasks from their jobs.

Since accuracy alone is not enough to win client trust in feasibility, the company is partnering with leading universities that publish in peer-reviewed journals to demonstrate that automatic road risk identification is possible. Once the idea is considered feasible by major scientific publications, the model needs to reach a high accuracy level before it can be employed in real projects.

The interviewee also had an opportunity to test a ML-driven tool from a young startup, in order to identify disruptive companies in a certain industry. Despite the promising concept, the tool fell short in terms of data accuracy, as the results were not pertinent to his needs. The idea was proven to be feasible, but challenging to execute. He also explained that large companies like C3 would rather invest in someone with three decades of expertise over a tool that only recently left beta testing.

There are other areas where they can and should utilize ML, including compound development, machinery improvements, and expand their service portfolio. However, the challenge arises from having a restricted budget to implement all these solutions. Priority is given to the ones providing clear advantages, measurable ROI, and the possibility of replacing humans. It is not a matter of *if*, but rather *when* these initiatives will be pursued.

According to C4 spokesperson, the main challenge of applying AI to decisions about the future is that the technology cannot deal well with novel situations. Training a model on historical data featuring the same products or worker-skill combination can be useful, but will not lead to any innovation by itself. Even if a model is fed with data on thousands of innovative products launches, it will just learn the patterns on how these products have been launched, and not how the next product will look like, particularly in the case of a

disruptive technology. Similarly, GenAI can generate text in a well-established format, like soccer match commentary, but it is unlikely to excel in novel forms of communication.

One cannot just ask the model a simple question and expect it to handle all the complications of real life. This is why autonomous vehicles are challenging to develop: even if ideally they take people from point A to point B, they also have to deal with situations like cats crossing the street or vehicles coming from other directions. The ML model needs guidance on how elaborate data, and this is implemented through a learning function that either optimizes the gains or minimizes the losses. In case of quality control, C4 labeled the training data as either good or faulty products, and in shift scheduling they modeled the production output of each time period as a function of the combination of available resources in the given period.

Due to the large number of unique production machines, the company cannot collect enough usage data for each model, and thus does not consider feasible to implement advanced process control or fault tolerant control. Instead, they believe that manufacturers should incorporate these functionalities directly into the tools.

C5 spokesperson cited Mollick (2023) by defining GenAI as “the world greatest intern”. Even when relying on this technology for an initial draft, employees still need their human brain, creativity, and understanding to finalize the document. In addition, it is crucial to maintain a human touch, as people prefer personal interactions over robotic ones. Used in this way, AI serves as a secondary brain which offloads simpler tasks, but not the ones where human expertise is needed. The ultimate decision-making authority rests with the human mind, as AI cannot decide for us. Software developers inside the company can ask GenAI to write code snippets, but not to build the entire program they are looking for. Another important issue is creating a responsible AI, being careful when training the model to avoid overfitting, making sure that the system returns an unbiased result, and correcting it otherwise.

LLMs can be fine-tuned, but they still represent a danger as they can hallucinate or return inaccurate information contained in their training material. According to C6 spokesperson, training from scratch on a smaller quantity of high-quality data should work better than training on a large dataset and then performing fine-tuning. The main challenge in applying AI to other areas of the company is that the main business is still human-centered, and some of the questions that the financial coaches have to answer are too complicated and require too much reasoning to be handled by a machine. Sometimes it may happen

that the answer to a client request is not even contained in the internal knowledge base, if the request is too niche or complex. Therefore, future AI implementations in the company do not aim to replace coaches, but rather to make them more productive, delivering more value to the customers.

C7 is cautious around the immaturity of these technologies, mostly for two reasons:

- **Legal liability:** companies are legally liable in case they share confidential information with third parties without prior consent. Therefore, C7 is wary of taking risks such as processing customers' private information through external services, especially if they are using the received data to further improve their models (Simens, 2021).
- **Result accuracy:** often, the reliability of results produced by AI tools is questionable. The solutions proposed may not work, or it may be nonsense as in the case of LLM hallucinations.

According to a BI manager in the company, AI is not at the point of saying what you should do, as it takes too many inputs other than data to make a useful recommendation. This concept was understood when the company was testing a new productivity tool for its employees: despite all the software employed to develop and test the product, only the final users could provide concrete feedback on how to improve the user experience.

According to a senior product manager at C8, there are other initiatives on analytics that the company could have pursued before. However, the mentality of executives opposed for long time to investments in data-collection processes, as they would not have the same short-term returns of enhancements in the core business. It has been difficult to use AI for decisions related to the design, manufacturing, marketing, sales, and customer service of blast chillers. The main reason is that these decisions often require human expertise and judgement, and any software that aims to advise or replace humans would need to fully capture the nuances and complexities of human decision making in these areas.

Currently, many decisions in C9 are made based on gut feeling and experience, and despite the efforts to establish a data-driven culture, it remains difficult to change how people take decisions. Another challenge the company is facing is to collect real time data. Since there is no way to collect statistics at the bar level, such as how many drinks made with a certain liquor are served each night, many figures are just estimates. When relying on external agencies, C9 often received insights summarized in a slideshow, without the raw data for further analysis.

When asked how an intelligent system can propose a brand to acquire to expand the product portfolio, the managers shared the following considerations. An AI system cannot tell the company which brand to buy, but it can potentially describe where most of the consumers are located today, where there is a need to improve our distribution channels, and in which areas the company should expand its presence through an acquisition. For instance, if in Chile the typical consumer is similar to the Italian one in terms of drinking habits, food pairing and working hours, but C9 is not present there, it can make sense to export an Italian liquor in the country or acquire a local brand. However, to perform an acquisition there must be a company willing to sell, and C9 has limited control over this external situation.

Due to its limited size, C10 relies on external partners when it comes to developing, hosting and integrating analytical systems to its business needs. The carbon neutrality objective of the company also represents a challenge, as training in-house large ML models comes with a high computational cost. Being present in a variety of industries from motorsports to fashion, the company has access to a wealth of data on its cars, customers and markets. Rather than collection or modeling, the main challenges lie in the organization of this data, since it has to be stored and classified properly. In the case of automated CRM reports, these insights are limited to describing historical figures of the company, and no business recommendation is generated by the software. The management may take into account these reports in their decisional processes, but often the information they provide do not play a primary role. This mostly happens for two reasons:

- The belief that these results can be inaccurate
- Executives are not used to take decisions based on data

When it comes to create content with GenAI, the main problems are the quality of the results and the risk of hallucinations. This technology also introduced ethical and privacy concerns, as confidential data is transmitted to third parties for processing. It is difficult to evaluate AI investments that improve the efficiency of operations, as their benefits are expressed in time savings and not in euro. These AI solutions are more complex than rule-based counterparts, they require more testing, and their models must be periodically monitored to prevent deviations of the results.

**Table 3.4:** Barriers to AI employment in decision making

ID	BARRIERS IN MORE STRATEGIC DECISIONS	BARRIERS IN MORE OPERATIONAL DECISIONS
C1	Inability to explain causes or collect qualitative data	Need to determine what data is useful, need to link metrics with business usage
C2	Costs of data pipelines and employees, specificity of business cases, explainability of the results	Impossibility to automate other processes
C3	Unmatched human reasoning, customers' doubts about feasibility, customers' aversion to automation of impactful decisions, mandatory human evaluation of automated recommendations	Need for accurate results, opportunity costs of AI investments
C4	Inability to innovate or predict disruptions, vagueness of business decisions	Definition of a learning function, data gathering and labeling
C5	Unmatched human creativity and understanding of situations	Attention to training and evaluation, people preference for human interactions
C6	Human-centered nature of the business	Complexity or specificity of customer inquiries, risk of hallucinations
C7	Inability to provide feedback or to consider factors other than data	Cautiousness in handling confidential data, unsatisfactory results
C8	Unmatched human judgement and decision making	Short-term ROI mentality, opportunity costs of AI investments
C9	Technology can only perform targeted analysis	Lack of data-driven mentality and real-time data
C10	Lack of data-driven mentality	Dependency on external companies, energy consumption, unsatisfactory results quality, need for periodic re-evaluation of models

### 3.2.4 CONSIDERATIONS ON AI EMPLOYMENT

In this section, we collect the activities that interviewed companies are performing to address the adoption barriers listed in table 3.4. We also invite experts to express their opinions on these limiting factors, and if they anticipate any change in the foreseeable future (RQ 1d). At the end, we summarize these answers in table 3.5.

The Chief Operating Officer of C1 underlines the importance of interacting with customers and end users to get new data and point of view. ML can instead be used for data-intensive problems when the outcome of the data analysis is known, for instance understand the improvement potential of individual components of a coffee maker. Quantitative data needs to be paired with qualitative information to take more unstructured decisions. An example is improving a component of coffee makers: qualitative analysis can highlight what is important to focus on, and why it is important for the bartenders. Quantitative data can determine how much this aspect can be improved, and how to do it. People will remain fundamental, as they can reason and decide to pursue digitalization processes or workshops with clients. Technology only comes at a later stage, as a tool to analyze and implement what emerged from these meetings. He said the main challenge will be upskilling employees to understand how these technologies work and for what tasks of their jobs they can be used.

The Italian lead of AI at C2 argues that LLM frameworks like *LangChain* (Chase, 2022) that can bridge structured and unstructured data have adoption potential in business contexts. When asked to comment the Palantir AIP demo video (Palantir Technologies, 2023), in which an employee seeks suggestions from a chatbot on how to address an upcoming hurricane that could impact the distribution centers, he explained that the most interesting part of the demo is conversational reporting, a trend that is being adopted by major BI vendors. In conversational reporting, instead of accessing a dashboard, users can utilize a chat interface to input their requests, as if they were asking an analyst for the information. Nevertheless, the query itself does not generate the data or the inference model on the fly. A company would still need:

- Historical data on hurricanes
- Real-time data from a weather providers
- Real-time data from sensors in the distribution centers
- Data scientists and weather experts to build the model that calculates the risk

The chat interface can eventually replace the need for a dashboard designer, but none of the aforementioned resources or scientists can be replaced. Therefore, only companies at risk of frequent or impactful hurricanes may find it beneficial to establish such a system and hire the experts it needs.

According to the Chief Innovation Officer at C3, many new opportunities for ML employment will arise with time. For instance, simulations and virtual analysis on digital twins of production machines can avoid the need of physical sensors in those machines. By considering ML as a tool to process information, a company should define what it wants to accomplish with it, and then collect the data needed to train the model. He also highlighted that the speed of ML adoption varies between industries, with financial, pharmaceutical and biotech companies as the fastest adopters, and manufacturers as the slowest. He estimates that it will take at least ten years to see this technology adopted throughout C3, while for other manufacturing companies this can take multiple decades.

Digital literacy of employees is also necessary to achieve an extensive adoption: Two decades ago, people perceived AI as a sentient computer whose aim was conquering the world. Nowadays, people are understanding the basics of AI and the value that AI applications can bring. It is also important to recognize and navigate hype cycles, as a lot of buzz and excitement are actually caused by sales and marketing departments. This pattern happens again and again: 5G connectivity has not changed the world, and was extremely overhyped.

When asked about the feasibility of a ML system that suggest a product to launch, the interviewee expressed the following considerations. Deciding what to launch is a broad question with many variables. If a ML-based system able to *fully answer* this question existed, it would need to ingest and analyze a massively huge amount of data. Existing tools can help by answering *only part* of that question, and they do not necessarily involve ML. A company can rely on a specialized agency to source raw materials and negotiate a fair price for them. However, this entire agency only covers a single factor, among the ones that have to be accounted for to launch a new product. Therefore, the capabilities of such a tool would be so extended to replace entire companies by itself. When it comes to big

decisions, it is not only about the data, there is also a human side to consider. This involves understanding the competitor nature, the decision alignment with the company strategy and culture and the important role of intuition and instinct. In the future, we will see AI recommendations (rather than AI decisions), and the human will decide based on them. To conclude, he said that when it comes to driving a business, he does not think we will ever see an AI taking over managers and make business decisions.

According to C4 spokesperson, applying AI to strategic decisions remains challenging as it requires modeling numerous indicators, some of which are difficult to quantify. Important decisions may require timely reaction to external factors that are not under the control of the company, such as new export restrictions (gallium and germanium produced in China), unforeseen outbreaks (COVID-19) or geopolitical incidents (the Russian invasion of Ukraine). All these events will not appear in predictive statistics or simulations, and the lack of historical data on them hinders the feasibility of ML-driven decision making.

Business decisions may be too complex to be addressed just by relying on structured data. AI can be employed for simple or standard business questions for which lot of data can be collected. For instance, the company could set a price based on price elasticity, or predict the reaction to *changes in packaging* based on the customer preferences it collected. This is different from using AI to come up with a *new product* and predict how millions of people will react to its launch. The military use case seems less complex, since there are fewer variables to consider and a set of constraints that limit the universe of choices. In this situation technology can potentially be used to take more strategic decisions, by representing the world as a 2D map and considering streets and railways as the only way to move units.

AI can definitely help when a company has a lot of data that it can structure and turn into training material. But if the problem involves high-level reasoning, you do not need AI, you need human experts capable of reasoning. AI can be a powerful weapon, but it would be excessive to apply it to any decision.

According to an experienced manager at C5, when it comes to dealing with customers there are information that cannot be coded into data. To take a decision, one would need to understand the relationship with the client, the current context, and what you want to do with them. AI cannot help in this area for now, but it can take some easy decisions. The manager said that even having access to AI software before it is available on the market, he does not think AI will ever compete with human cognition, and that this technology is targeted to assist, rather than replace human decisions.



However, he said that we overestimate the impact of technology in short term, and we underestimate the impact of technology in the long term. We don't know the impact that AI can have in the long term, but we see more and more use cases every day in every area, from traveling to work to education.

The director of AI at C6 argued that even if modern AI pass the Turing test, they do not replace human intelligence because humans are still capable of orchestrating knowledge in a certain way that computers cannot. Chain of Thought (CoT) is focusing in this direction, trying to replicate how our brain works. This process can be employed to tame LLMs: even if these models are not good at reasoning or math, they can be used to take decisions, if the prompt is formulated as a series of question or a course of action. With CoT, one can break down a complex decision in multiple small prompts, and derive individual answers. However, humans still needs to define these intermediate steps, limit bias and hallucinations by employing the most suitable tool for each task, and check the result of each step.

According to a BI manager at C7, companies need to address their basic data issues, before jumping in ML and advanced analytics. All decisions need some data as input, but decision makers should not rely solely on BI. They also need to consider the risks and the dependencies associated to each outcome, and keep in mind the overall strategy of the business. Important decisions often require a collaborative back and forth, involving internal or external actors.

The interviewee also shared some thoughts about building a predictive analytics system. She said that even if some companies may consider hiring data scientists to implement a prescriptive analytics system, many others cannot make this leap yet. Based on her experience, most organizations did not succeed in implementing their BI platform, as they cannot even answer basic question on what already happened. Major reasons for these failures are the lack of data literacy among employees, and the complexity of data itself. Unlike a simple table in a spreadsheet, this data may need specific business logics to be interpreted correctly. This is particularly true if data coming from the reporting system does not match the one from the transactional systems, probably due to a different aggregation logic. For companies that started to collect and organize data many decades ago, now the issue could be finding the appropriate dashboard and navigate to the relevant content, in case of too many alternative dashboards.

The senior product manager at C8 argued that many of the decisions related to the design, manufacturing, marketing, sales, and customer service might not require ML. These

decisions could be based on market research, customer feedback, traditional engineering principles and domain knowledge.

According to the analytics director at C9, every organization should follow these steps to define their data strategy:

1. Understand the growth drivers of the company (e.g. increasing revenues)
2. Individuate the business needs linked to these drivers (e.g. raising the price of a product)
3. Prepare a data substrate that provides insights about these needs (e.g. price elasticity of customers, brand loyalty)

These steps must be approached in order, as the information system is just an enabler of the business strategy, and must be fed with the data needed to take decisions. It is crucial that executives know where their business is headed to, because no technology, no matter how advanced, can tell them what are their business needs. Finding these needs is not be trivial, and may require a structured analysis by itself. In parallel, the company must put effort in change management, as the information system can only deliver value when employees understand its link to the business needs and stop taking decisions based on opinions or gut feeling. What is happening on ML is a cultural change more than a technological one, and this change may take companies over ten years after they realize the importance of data.

The director of analytics at C10 is confident that over time the executive team will start relying more and more on data for their decisions, as they understand the benefits of this approach. C10 is constantly monitoring internal use cases where AI can bring benefits, and learning from the external environments about new developments and possible applications.

**Table 3.5:** Considerations on AI employment in decision making

ID	CONSIDERATIONS ON MORE STRATEGIC DECISIONS	CONSIDERATIONS ON MORE OPERATIONAL DECISIONS
C1	Interact with clients to collect data, consider ML as a tool for specific analyses	Educate employees on AI and its potential use cases
C2	Evaluate the costs and benefits of implementing AI solutions	Monitor developments in LLM frameworks
C3	Aim to obtain recommendations rather than automatic decisions, monitor and promote scientific research on novel use cases	Have a narrow objective for the model, rely on simulations to reduce costs
C4	Consider how possible external disruptions may impact ML forecasting, acknowledge AI shortcomings in high-level reasoning	Limit ML employment to straightforward and data-intensive decisions for which data is available
C5	Consider AI applications as complementary to human cognition, monitor novel use cases	Prefer direct interactions and personalized communications when dealing with clients
C6	Monitor developments in CoT	Have experts to address complex or niche requests, train models on unbiased data, employ LLMs where hallucinations are less problematic
C7	Acknowledge that data is just one of the input needed to take a decision and that information systems can only process data	Educate employees on how to get the best results from their tools
C8	Evaluate where human expertise is more effective than ML applications	Consider long-term impacts of digital investments other than short-term costs
C9	Understand what the company needs and only then implement a targeted insight system	Educate employees on how data can help addressing the company needs
C10	Educate executives on the benefits of an analytical decision making approach	Monitor technology developments and novel use cases

### 3.3 ANALYSIS AND DISCUSSION OF RESULTS

INTERVIEWING large companies on their current AI employment in decision making resulted in insightful conversations that offered interesting point of views and actionable advices for executives. By summarizing the answers received in four thematic tables, we offer an overview of the current employment of AI in decision making, and we facilitate the following discussion.

In this section, we analyze and discuss thematically the data we collected from the interviews. Section 3.3.1 will analyze the main AI use cases that emerged from the interviews, while section 3.3.2 will discuss the reasons behind this adoption. Section 3.3.3 will comment the main challenges that are preventing companies from applying AI to other kind of decisions. Section 3.3.4 will summarize the main considerations of interviewed managers on the AI employment.

#### 3.3.1 CURRENT AI EMPLOYMENT

The most common AI use cases among the interviewed companies are:

- Employee productivity (C2, C5, C6, C7, C10)
- Telemetry data analysis (C1, C3, C8, C10)
- Optical pattern recognition (C3, C4, C10)
- Time-series forecasting (C5, C9)

The following use cases are instead unique within the sample:

- Road risk identification (C3)
- Shift scheduling (C4)
- Documentation auditing (C7)
- Media labeling (C10)
- Predictive maintenance (C10)

Some of the interviewed companies are embedding AI in their client offering, rather than using it internally:

- C1 would like to embed predictive maintenance in their coffee makers
- C2 is implementing third-party ML software (e.c. CRM) for their clients
- C3 is providing two CV models to road authorities
- C5 is embedding LLMs in their commercial productivity software
- C6 is developing a different kind of language model that can be trained from scratch on proprietary data

C10 also appears as the only company that succeeded in implementing predictive maintenance. C4 said that this activity is not practical, due to the high number of unique machines employed in production. C1 instead plans to embed this feature in the coffee makers it produces. C6 is not providing automated financial advices to its customers, but the coaches in the company are leveraging AI tools to seek information faster.

The ML analysis of telemetry by C1 and C8 turned out to play only a marginal role in the entire analysis process. C1 said to use focus groups rather than NLP to determine the issues to address first and cover in their social media communication.

To the best of our knowledge, we could not identify an example of AI employment for a strategic business decision. The strategic-organizational classification of table 3.2 also considers the answers received in sections 3.2.3 and 3.2.4, as it emerged that companies find challenging to apply AI to more complex or important decisions. Surprisingly, not even the largest companies operating also in the software industry (C2, C3, C5) could describe an internal case of AI employment in decisions that they consider strategic. When asked about this lack of strategic use cases, they explained that they do not consider it feasible (sections 3.3.3 and 3.3.4).

#### 3.3.2 DRIVERS OF AI EMPLOYMENT

When interviewed companies explained their reasons behind these adoptions, they both talked about properties of the decisions and the technology (section 3.2.2). Eight companies out of ten mentioned the availability of data as a key factor:

- Structured real-time telemetry (C1, C3, C8, C10)
- Structured historical data (C4, C5, C9, C10)
- Unstructured images and videos (C4, C10)
- Unstructured text (C6)

Another common feature of these tasks is that they can be automated:

- C2 is automating repetitive back office processes
- C3 is automating the data analysis step of road risk assessments
- C4 is replacing existing automation systems (rule-based optical quality inspection, and shift scheduling optimization algorithms) with more accurate or faster ML alternatives
- C5 is automating the forecast of revenues, the preparation of document, the generation of code snippets and the summarization of long video meetings
- C6 is automating the summarization of articles in their internal knowledge base through a chat interface
- C7 is automating the preliminary audit of clients' manuals and the generation of code snippets
- C9 is automating the forecast of demand of goods
- C10 is automating the tagging of images and videos, optical quality inspection and the creation of creative content

The most mentioned benefits that arise from AI employment are:

- Enhancements in employees' productivity (C5, C6, C7, C10)
- Faster results from automated processes (C4, C5, C7, C8)
- Accuracy in forecasting (C5, C9)
- Accuracy in pattern recognition (C3, C4)
- Capacity to improve from new examples (C3, C4)

It emerges that the main features of the decisions where interviewed companies employ AI, are the potential for automation and the availability of data. The main benefits related to the introduction of AI are time savings in both employees' productivity and automated processes. Other significant factors are the accuracy of forecasting and pattern recognition tasks.

#### 3.3.3 BARRIERS TO AI EMPLOYMENT

The main reasons for which AI employment is challenging for more strategic decisions are:

- Limited cognitive capabilities of the technology (C1, C3, C4, C5, C6, C7, C8, C9)
- People mentality (C3, C10)

The main reasons regarding more operational decisions are instead:

- Need for acceptable results (C3, C6, C7, C10)
- People mentality (C5, C8, C9)
- Need to determine the data and model needed (C1, C4)
- Maintenance of models (C5, C10)
- Complexity of business decisions (C2, C6)
- Cost-benefit trade-off (C3, C8)

Among the reasons for strategic employment mentioned only once, C4 argues that business decisions are too complex to model, as they lack constraints and they need to account for many variables at the same time. C6 explain that the core business relies on human interactions, which cannot be replaced by software. C2 mentioned that explainability of results is not yet a central discussion topic, but it is crucial if decisions have influential consequences on a variety of stakeholders. In addition, companies need to evaluate the potential benefits against the costs of a complex AI solution

When C3 mentioned *legal liability* as a challenge to develop a fully automated prescriptive system, they meant that their client has the duty to perform an accurate risk evaluation before accepting an automatic prescription on the risk level. Instead, when C7 and C10 mentioned *legal liability*, they meant that they have the duty to not transmit confidential information to third parties, including AI providers for content generation purposes. Both C3 and C9 acknowledge that the adoption speed of analytics and AI changes between industries, with data-intensive companies having a hedge over manufacturing ones.

#### 3.3.4 CONSIDERATIONS ON AI EMPLOYMENT

The most common considerations on possible AI employments in strategic decision making are the following:

- Limit AI employment to narrow or data-intensive tasks (C1, C3, C4, C5, C8, C9)
- Be aware of the factors not considered by AI models (C1, C4, C7)
- Stay current with novel developments and use cases (C5, C6)

When it comes to more operational use cases instead, the most common considerations are:

- Upskill people and spread a data-driven mentality (C1, C7, C8, C9)
- Limit AI employment to narrow or data-intensive tasks (C3, C4, C6)
- Stay current with novel developments and use cases (C2, C10)
- Maintain human relationships with clients (C5, C6)

Other considerations on more strategic use cases is the need to start from a real business necessity before implementing an information system (C9). Companies should also evaluate the cost-benefit trade-off of a complex AI solutions (C2), as gathering the necessary experts and software is expensive. C10 stress the importance of having a data-driven mentality among executives, while C1, C7, C8 and C9 focus on employees.

In general, managers are well aware of both the limitations of current AI technologies and the complexities of their business cases. We found that most of them (C3, C4, C5, C7, C8, C9) resonated with Jordan (2019) in saying that current AI is not capable of human level cognition, and therefore cannot understand complex situations and take strategic decision. Some of these experts (C3, C4, C5, C9) take this opinion further, arguing that achieving this goal will not be possible.

#### 3.3.5 ANSWER TO RESEARCH QUESTION

Based on the interviewed sample and our analysis of the results, we conclude that AI is more suited for operational rather than strategic decision making. This is explained as follows.

From the interviews, it emerged that the decisions that can benefit the most from AI employment are present a series of common features. The data needed to address these issues is collectible or already available; the decision steps have potential for automation (section 3.3.2). These decisions do not require complex reasoning (section 3.3.3), have limited impact in case of a mistake and are narrow in their scope (section 3.3.4). We also argue that these decisions tend to happen routinely rather than one-off (table 2.8), based on all the use cases discussed (section 3.3.1).

Decisions of this kind benefit from the main strengths of AI technologies: accuracy in forecasting, accuracy in pattern recognition, faster results after the training, capacity to learn from examples (section 3.3.2).

Strategic decision makers of today remain executives that rely on their experience, reasoning, intuition and ad-hoc data analysis to make these decisions (section 3.3.4).



## 3.4 CONCLUSIONS

THIS master's thesis has presented an extensive exploration of AI in decision making. The research process involved a thorough literature review and a series of interviews to investigate the impact of AI technologies in decision making.

By starting with an exploratory bibliometric approach in the broad field of data-driven decision making, we identified AI as one of the main research themes. After a brief but necessary introduction on the historical meaning of AI, we clarify its research area and current relevance for organizational employment. Then, we narrowed down our research on the impact of AI in organizational decision making. We explained how current literature is organized (bibliographic coupling) and by what themes it was influenced (co-citation analysis). This approach allowed to individuate an important research gap.

We elaborated on well-known frameworks and theories to propose a framework to separate business decisions in more strategic and more operational, according to the properties of these decisions. The original contribution of this thesis lies in explaining what kind of business decisions can benefit the most from AI. To answer this question, we conduct empirical analysis on four main aspects: current AI employment, drivers of AI employment, barriers to AI employment, and experts' considerations on AI employment.

By collecting primary qualitative data on large manufacturing companies, we aim for a more accurate representation of the impact of AI technologies in decision making. Ten experienced professionals took part to the interview, explaining their current use cases for AI technologies, their reasons behind this adoption, the challenges that are stopping them from expanding their AI usage, and their point of view on possible future employments of AI.

The main results obtained from this research indicate that AI is more adapt to organizational rather than strategic decisions. According to the interview sample, the decisions that can benefit the most from AI employment present one or more of the following features: they can be answered from data that is easily collectible or already available, they can be structured and automated, they do not require human-level cognition, they are narrow in their scope, and they happen routinely. Companies value AI technologies for their accuracy in forecasting, accuracy in pattern recognition, fast results compared to large computations and their capacity to learn from examples.

As for any scientific study, it is important to acknowledge the limitations of this research. Firstly, the proposed framework to categorize decisions is not free from biases as must be

approached with critical judgement. There is no such thing as a dichotomy of business decisions, but the framework can help understanding why strategic decisions tend to differ from more operational ones. Secondly, we opted to have semi-structured interviews to better discuss the factors driving or hindering AI employments. This research methodology did not allow for a larger sample of companies or a more structured set of results to analyze. Thirdly, we did not restrict our limited sample to a specific industry, as we preferred to give voices to different point of views. Fourthly, what emerged in the interviews may not represent the entire situation of interviewed companies, especially for larger ones. To mitigate this issue, we reached out to the people that to the best of our knowledge were the most aware about internal AI applications.

As for future research directions, there are several avenues worth exploring to clarify the impact of AI in organizational decision making. Firstly, quantitative studies such as surveys can target larger samples and analyze them with statistical tools. Secondly, by restricting the analysis to a specific industry or geography, one can better study the differences and similarities within a more homogeneous sample. Thirdly, individual case studies are needed to cover in details all the existing AI applications to decision making inside a company. Fourthly, a periodic evaluation on the impact of AI in decision making can shed the light on AI adoption rates over time, and cover novel use cases of these technologies.

As a closing remark, we recommend people willing to expand their usage of technologies, to keep in mind the *Law of the instrument* (Maslow, 1966):

“If the only tool you have is a hammer, it is tempting to treat everything as if it were a nail.”

To avoid this cognitive bias, it is important to start from the problem, understand its properties, gather a list of alternative solutions, and only then evaluate if AI represents the best way to address the issue.

# BIBLIOGRAPHY

## ARTICLES

- Arnott, D., & Pervan, G. (2005). A critical analysis of decision support systems research. *Journal of Information Technology*, 20(2), 67–87. <https://doi.org/10.1057/palgrave.jit.2000035> (cit. on p. 19).
- Atanassov, K. T. (1986). Intuitionistic fuzzy sets. *Fuzzy Sets and Systems*, 20(1), 87–96. [https://doi.org/10.1016/S0165-0114\(86\)80034-3](https://doi.org/10.1016/S0165-0114(86)80034-3) (cit. on pp. 33, 47, 48).
- Baas, J., Schotten, M., Plume, A., Côté, G., & Karimi, R. (2020). Scopus as a curated, high-quality bibliometric data source for academic research in quantitative science studies. *Quantitative Science Studies*, 1(1), 377–386. [https://doi.org/10.1162/qss\\_a\\_00019](https://doi.org/10.1162/qss_a_00019) (cit. on p. 17).
- Bączkiewicz, A., Kizielewicz, B., Shekhovtsov, A., Wątróbski, J., & Sałabun, W. (2021). Methodical aspects of mcdm based e-commerce recommender system. *Journal of Theoretical and Applied Electronic Commerce Research*, 16(6), 2192–2229. <https://doi.org/10.3390/jtaer16060122> (cit. on pp. 46, 47, 56).
- Barredo Arrieta, A., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., Garcia, S., Gil-Lopez, S., Molina, D., Benjamins, R., Chatila, R., & Herrera, F. (2020). Explainable artificial intelligence (xai): Concepts, taxonomies, opportunities and challenges toward responsible ai. *Information Fusion*, 58, 82–115. <https://doi.org/10.1016/j.inffus.2019.12.012> (cit. on pp. 37, 38, 58).
- Belhadi, A., Kamble, S., Fosso Wamba, S., & Queiroz, M. M. (2022). Building supply-chain resilience: An artificial intelligence-based technique and decision-making framework. *International Journal of Production Research*, 60(14), 4487–4507 (cit. on p. 40).
- Benzidia, S., Makaoui, N., & Bentahar, O. (2021). The impact of big data analytics and artificial intelligence on green supply chain process integration and hospital environmental performance. *Technological Forecasting and Social Change*, 165, 120557. <https://doi.org/10.1016/j.techfore.2020.120557> (cit. on p. 46).
- Bonjour, T., Haliem, M., Alsalem, A., Thomas, S., Li, H., Aggarwal, V., Kejriwal, M., & Bhargava, B. (2022). Decision making in monopoly using a hybrid deep reinforce-

- ment learning approach. *IEEE Transactions on Emerging Topics in Computational Intelligence*, 6(6), 1335–1344 (cit. on pp. 56, 57, 59).
- Chen, H., Chiang, R. H., & Storey, V. C. (2012). Business intelligence and analytics: From big data to big impact. *MIS quarterly*, 1165–1188 (cit. on pp. 12, 41).
- Chen, L., Jiang, M., Jia, F., & Liu, G. (2022). Artificial intelligence adoption in business-to-business marketing: Toward a conceptual framework. *Journal of Business & Industrial Marketing*, 37(5), 1025–1044. <https://doi.org/10.1108/JBIM-09-2020-0448> (cit. on pp. 54, 59).
- Dandolo, D., Masiero, C., Carletti, M., Dalle Pezze, D., & Susto, G. A. (2023). Acme—accelerated model-agnostic explanations: Fast whitening of the machine-learning black box. *Expert Systems with Applications*, 214, 119115 (cit. on p. 56).
- Davis, F. D. (1989). Perceived usefulness, perceived ease of use, and user acceptance of information technology. *MIS quarterly*, 319–340 (cit. on p. 41).
- Dempster, A. P. (1967). Upper and lower probabilities induced by a multivalued mapping. *The Annals of Mathematical Statistics*, 38(2), 325–339. Retrieved May 6, 2023, from <http://doi.org/10.1214/aoms/1177698950> (cit. on p. 34).
- Deng, Y. (2020). Uncertainty measure in evidence theory. *Science China Information Sciences*, 63(11), 210201. <https://doi.org/10.1007/s11432-020-3006-9> (cit. on p. 35).
- Dhir, A., Khan, S. J., Islam, N., Ractham, P., & Meenakshi, N. (2023). Drivers of sustainable business model innovations. an upper echelon theory perspective. *Technological Forecasting and Social Change*, 191, 122409 (cit. on p. 21).
- Dietvorst, B., Simmons, J., & Massey, C. (2014). Algorithm aversion: People erroneously avoid algorithms after seeing them err. *Journal of experimental psychology. General*, 144. <https://doi.org/10.1037/xge0000033> (cit. on pp. 32, 43, 51, 56).
- Donthu, N., Kumar, S., Mukherjee, D., Pandey, N., & Lim, W. M. (2021). How to conduct a bibliometric analysis: An overview and guidelines. *Journal of Business Research*, 133, 285–296 (cit. on pp. 11, 12, 14, 16, 17).
- Duan, Y., Edwards, J. S., & Dwivedi, Y. K. (2019). Artificial intelligence for decision making in the era of big data - evolution, challenges and research agenda. *Int. J. Inf. Manag.*, 48, 63–71. <http://dblp.uni-trier.de/db/journals/ijinfoman/ijinfoman48.html#DuanED19> (cit. on pp. 19, 25, 32, 42, 56, 58).
- Dubey, R., Gunasekaran, A., Childe, S. J., Bryde, D. J., Giannakis, M., Foropon, C., Roubaud, D., & Hazen, B. T. (2020). Big data analytics and artificial intelligence pathway

- to operational performance under the effects of entrepreneurial orientation and environmental dynamism: A study of manufacturing organisations. *International Journal of Production Economics*, 226, 107599. <https://doi.org/10.1016/j.ijpe.2019.107599> (cit. on pp. 32, 39).
- Edwards, J. S., Duan, Y., & Robins, P. C. (2000). An analysis of expert systems for business decision making at different levels and in different roles. *European Journal of Information Systems*, 9(1), 36–46. <https://doi.org/10.1057/palgrave.ejis.3000344> (cit. on pp. 56, 57, 59).
- Fabianova, J., Janekova, J., & Horbulak, J. (2021). Solving the bottleneck problem in a warehouse using simulations. *Acta Logistica*, 8(2), 107–116. <https://doi.org/10.22306/al.v8i2.209> (cit. on p. 21).
- Floridi, L. (2023). AI as agency without intelligence: On ChatGPT, large language models, and other generative models. *Philosophy & Technology*, 36(1), 15. <https://doi.org/10.1007/s13347-023-00621-y> (cit. on p. 25).
- Glikson, E., & Woolley, A. W. (2020). Human trust in artificial intelligence: Review of empirical research. *Academy of Management Annals*, 14(2), 627–660 (cit. on p. 44).
- Helo, P., & Hao, Y. (2022). Artificial intelligence in operations management and supply chain management: An exploratory case study. *Production Planning & Control*, 33(16), 1573–1590 (cit. on pp. 40, 63).
- Herrera-Viedma, E., Palomares, I., Li, C.-C., Cabrerizo, F. J., Dong, Y., Chiclana, F., & Herrera, F. (2021). Revisiting fuzzy and linguistic decision making: Scenarios and challenges for making wiser decisions in a better way. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 51(1), 191–208 (cit. on p. 48).
- Highhouse, S. (2008). Stubborn reliance on intuition and subjectivity in employee selection. *Industrial and Organizational Psychology*, 1(3), 333–342 (cit. on p. 43).
- Hong, I. B., & Vogel, D. R. (1991). Data and model management in a generalized mcdm-dss. *Decision Sciences*, 22(1), 1–25 (cit. on p. 19).
- Huang, M.-H., & Rust, R. T. (2018). Artificial intelligence in service. *Journal of Service Research*, 21(2), 155–172. <https://doi.org/10.1177/1094670517752459> (cit. on pp. 41, 56, 58).
- Jarke, M. (1986). Knowledge sharing and negotiation support in multiperson decision support systems. *Decision Support Systems*, 2(1), 93–102. [https://doi.org/10.1016/0167-9236\(86\)90125-9](https://doi.org/10.1016/0167-9236(86)90125-9) (cit. on p. 19).

- Jarke, M., & Radermacher, F. J. (1988). The ai potential of model management and its central role in decision support. *Decision Support Systems*, 4(4), 387–404 (cit. on p. 19).
- Johnson, M., Albizri, A., & Simsek, S. (2022). Artificial intelligence in healthcare operations to enhance treatment outcomes: A framework to predict lung cancer prognosis. *Annals of Operations Research*, 1–31 (cit. on p. 50).
- Jordan, M. I., & Mitchell, T. M. (2015). Machine learning: Trends, perspectives, and prospects. *Science*, 349(6245), 255–260. <https://doi.org/10.1126/science.aaa8415> (cit. on pp. 32, 36).
- Jordan, M. I. (2019). Artificial intelligence - the revolution hasn't happened yet. *Harvard Data Science Review*, 1(1), 1–9 (cit. on pp. 25, 26, 58, 92).
- Ju, J., Liu, L., & Feng, Y. (2018). Citizen-centered big data analysis-driven governance intelligence framework for smart cities. *Telecommunications Policy*, 42(10), 881–896 (cit. on p. 21).
- Keding, C., & Meissner, P. (2021). Managerial overreliance on ai-augmented decision-making processes: How the use of ai-based advisory systems shapes choice behavior in r&d investment decisions. *Technological forecasting and social change*, 171, 120970 (cit. on pp. 46, 51, 56, 61).
- Krakowski, S., Luger, J., & Raisch, S. (2022). Artificial intelligence and the changing sources of competitive advantage. *Strategic Management Journal*. <https://doi.org/10.1002/smj.3387> (cit. on pp. 46, 55, 58).
- Lawler, J. J., & Elliot, R. (1996). Artificial intelligence in hrm: An experimental study of an expert system. *Journal of Management*, 22(1), 85–111. [https://doi.org/10.1016/S0149-2063\(96\)90013-6](https://doi.org/10.1016/S0149-2063(96)90013-6) (cit. on pp. 42, 56, 58).
- Lemos, S. I., Ferreira, F. A., Zopounidis, C., Galariotis, E., & Ferreira, N. C. (2022). Artificial intelligence and change management in small and medium-sized enterprises: An analysis of dynamics within adaptation initiatives. *Annals of Operations Research*, 1–27 (cit. on pp. 53, 58, 63).
- Leyer, M., & Schneider, S. (2021). Decision augmentation and automation with artificial intelligence: Threat or opportunity for managers? *Business Horizons*, 64(5), 711–724 (cit. on pp. 2, 51, 52, 58, 62).
- Logg, J. M., Minson, J. A., & Moore, D. A. (2019). Algorithm appreciation: People prefer algorithmic to human judgment. *Organizational Behavior and Human Decision Processes*, 151, 90–103 (cit. on pp. 43, 51, 56).

- Loureiro, S. M. C., Guerreiro, J., & Tussyadiah, I. (2021). Artificial intelligence in business: State of the art and future research agenda. *Journal of business research*, 129, 911–926 (cit. on p. 12).
- Martin, N. H. (1956). Differential decisions in the management of an industrial plant. *The Journal of Business*, 29(4), 249–260 (cit. on p. 4).
- Michalek, J. J., Feinberg, F. M., & Papalambros, P. Y. (2005). Linking marketing and engineering product design decisions via analytical target cascadin. *Journal of Product Innovation Management*, 22(1), 42–62. <https://doi.org/10.1111/j.0737-6782.2005.00102.x> (cit. on p. 21).
- Min, H. (2010). Artificial intelligence in supply chain management: Theory and applications. *International Journal of Logistics: Research and Applications*, 13(1), 13–39 (cit. on pp. 39, 40, 55–58).
- Mintzberg, H., Raisinghani, D., & Théorêt, A. (1976). The structure of "unstructured" decision processes. *Administrative Science Quarterly*, 21(2), 246–275. Retrieved September 18, 2023, from <https://doi.org/10.2307/2392045> (cit. on pp. 3, 4).
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., Petersen, S., Beattie, C., Sadik, A., Antonoglou, I., King, H., Kumaran, D., Wierstra, D., Legg, S., & Hassabis, D. (2015). Human-level control through deep reinforcement learning. *Nature*, 518(7540), 529–533. <https://doi.org/10.1038/nature14236> (cit. on p. 36).
- Murphy, C. K. (2000). Combining belief functions when evidence conflicts. *Decision Support Systems*, 29(1), 1–9. [https://doi.org/10.1016/S0167-9236\(99\)00084-6](https://doi.org/10.1016/S0167-9236(99)00084-6) (cit. on p. 35).
- Nguyen, H. D., Tran, K. P., Thomassey, S., & Hamad, M. (2021). Forecasting and anomaly detection approaches using lstm and lstm autoencoder techniques with the applications in supply chain management. *International Journal of Information Management*, 57, 102282 (cit. on p. 56).
- Pietronudo, M. C., Croidieu, G., & Schiavone, F. (2022). A solution looking for problems? a systematic literature review of the rationalizing influence of artificial intelligence on decision-making in innovation management. *Technological Forecasting and Social Change*, 182, 121828 (cit. on pp. 12, 59).
- Qamar, Y., Agrawal, R. K., Samad, T. A., & Chiappetta Jabbour, C. J. (2021). When technology meets people: The interplay of artificial intelligence and human resource management. *Journal of Enterprise Information Management*, 34(5), 1339–1370. <https://doi.org/10.1108/JEIM-11-2020-0436> (cit. on pp. 46, 53).

- Qiao, H., Chen, J., & Huang, X. (2022). A survey of Brain-Inspired intelligent robots: Integration of vision, decision, motion control, and musculoskeletal systems. *IEEE Trans Cybern*, 52(10), 11267–11280 (cit. on pp. 46, 49).
- Rajagopal, N. K., Qureshi, N. I., Durga, S., Ramirez Asis, E. H., Huerta Soto, R. M., Gupta, S. K., & Deepak, S. (2022). Future of business culture: An artificial intelligence-driven digital framework for organization decision-making process. *Complexity*, 2022 (cit. on pp. 3, 5, 6, 55, 56).
- Raza, S. A., Govindaluri, S. M., & Bhutta, M. K. (2023). Research themes in machine learning applications in supply chain management using bibliometric analysis tools. *Benchmarking: An International Journal*, 30(3), 834–867 (cit. on pp. 12, 55, 62).
- Riggs, R., Roldán, J. L., Real, J. C., & Felipe, C. M. (2023). Opening the black box of big data sustainable value creation: The mediating role of supply chain management capabilities and circular economy practices. *International Journal of Physical Distribution & Logistics Management* (cit. on p. 21).
- Rodríguez, R. M., Martínez-López, L., & Herrera, F. (2012). Hesitant fuzzy linguistic term sets for decision making. *IEEE Trans. Fuzzy Syst.*, 20(1), 109–119. <http://dblp.uni-trier.de/db/journals/tfs/tfs20.html#RodriguezMH12> (cit. on pp. 32, 33, 48).
- Schmitt, M. (2023). Deep learning in business analytics: A clash of expectations and reality. *International Journal of Information Management Data Insights*, 3(1), 100146 (cit. on p. 56).
- Shajalal, M., Boden, A., & Stevens, G. (2022). Explainable product backorder prediction exploiting cnn: Introducing explainable models in businesses. *Electronic Markets*, 32(4), 2107–2122. <https://doi.org/10.1007/s12525-022-00599-z> (cit. on pp. 46, 49, 50, 58, 59).
- Sharma, R., Shishodia, A., Gunasekaran, A., Min, H., & Haque Munim, Z. (2022). The role of artificial intelligence in supply chain management: Mapping the territory. *International Journal of Production Research*, 60(24), 7527–7550. <https://doi.org/10.1080/00207543.2022.2029611> (cit. on pp. 19, 21).
- Shrestha, Y. R., Krishna, V., & von Krogh, G. (2021). Augmenting organizational decision-making with deep learning algorithms: Principles, promises, and challenges. *Journal of Business Research*, 123, 588–603 (cit. on pp. 56, 58, 59).
- Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., van den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., Dieleman, S., Grewe, D., Nham, J., Kalchbrenner, N., Sutskever, I., Lillicrap, T., Leach, M., Kavukcuoglu, K., Graepel,



- T., & Hassabis, D. (2016). Mastering the game of go with deep neural networks and tree search. *Nature*, 529(7587), 484–489. <https://doi.org/10.1038/nature16961> (cit. on pp. 36, 56).
- Tang, M., & Liao, H. (2021). From conventional group decision making to large-scale group decision making: What are the challenges and how to meet them in big data era? a state-of-the-art survey. *Omega*, 100, 102141. <https://doi.org/10.1016/j.omega.2019.102141> (cit. on pp. 12, 33, 34).
- Toorajipour, R., Sohrabpour, V., Nazarpour, A., Oghazi, P., & Fischl, M. (2021). Artificial intelligence in supply chain management: A systematic literature review. *Journal of Business Research*, 122, 502–517 (cit. on pp. 16, 40, 55, 58).
- Valenzuela, M., Ha, V., & Etzioni, O. (2015). Identifying meaningful citations (C. Caragea, C. L. Giles, N. Bhamidipati, D. Caragea, S. Das Gollapalli, S. Kataria, H. Liu, & F. Xia, Eds.). (WS-15-13), 21–26. <http://www.aaai.org/Library/Workshops/ws15-13.php> (cit. on p. 14).
- Xiao, F. (2020). Efmcdm: Evidential fuzzy multicriteria decision making based on belief entropy. *IEEE Transactions on Fuzzy Systems*, 28(7), 1477–1491. <https://doi.org/10.1109/TFUZZ.2019.2936368> (cit. on pp. 32, 35).
- Xie, D., Xiao, F., & Pedrycz, W. (2022). Information quality for intuitionistic fuzzy values with its application in decision making. *Engineering Applications of Artificial Intelligence*, 109, 104568. <https://doi.org/10.1016/j.engappai.2021.104568> (cit. on pp. 46, 47).
- Yan, M.-R., Hong, L.-Y., & Warren, K. (2022). Integrated knowledge visualization and the enterprise digital twin system for supporting strategic management decision. *Management Decision*, 60(4), 1095–1115 (cit. on p. 58).
- Ylijoki, O., & Porras, J. (2018). A recipe for big data value creation. *Business Process Management Journal*, 25. <https://doi.org/10.1108/BPMJ-03-2018-0082> (cit. on pp. 21, 56).
- Zadeh, L. (1975). The concept of a linguistic variable and its application to approximate reasoning. *Information Sciences*, 8(3), 199–249. [https://doi.org/10.1016/0020-0255\(75\)90036-5](https://doi.org/10.1016/0020-0255(75)90036-5) (cit. on p. 33).
- Zadeh, L. A. (1965). Fuzzy sets. *Information and Control*, 8, 338–353. [https://doi.org/10.1016/s0019-9958\(65\)90241-x](https://doi.org/10.1016/s0019-9958(65)90241-x) (cit. on pp. 32, 48).
- Zhang, Z., Wang, H., Geng, J., Jiang, W., Deng, X., & Miao, W. (2022). An information fusion method based on deep learning and fuzzy discount-weighting for target

- intention recognition. *Engineering Applications of Artificial Intelligence*, 109, 104610. <https://doi.org/10.1016/j.engappai.2021.104610> (cit. on pp. 49, 56).
- Zhou, Z.-J., Hu, G.-Y., Hu, C.-H., Wen, C.-L., & Chang, L.-L. (2021). A survey of belief rule-base expert system. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 51(8), 4944–4958. <https://doi.org/10.1109/TSMC.2019.2944893> (cit. on p. 47).

## BOOKS

- Anthony, R. N. (1965). *Planning and control systems: A framework for analysis*. Division of Research, Graduate School of Business Administration, Harvard University. <https://books.google.it/books?id=OY4uAAAAMAAJ> (cit. on p. 4).
- Davenport, T., Brynjolfsson, E., McAfee, A., & Wilson, H. (2019). *Artificial intelligence: The insights you need from harvard business review*. Harvard Business Review Press. <https://store.hbr.org/product/artificial-intelligence-the-insights-you-need-from-harvard-business-review/10281> (cit. on pp. 56, 59).
- Kahneman, D. (2011). *Thinking, fast and slow*. Farrar, Straus; Giroux. <https://us.macmillan.com/books/9781429969352/thinkingfastandslow> (cit. on p. 58).
- Lamarre, E., Smaje, K., & Zimmel, R. (2023). *Rewired: The McKinsey guide to outcompeting in the age of digital and AI*. Wiley. <https://books.google.it/books?id=sRO5zwEACAAJ> (cit. on p. 61).
- March, J., & Simon, H. (1958). *Organizations*. Wiley. <https://books.google.it/books?id=9oxEAAAIAAJ> (cit. on pp. 35, 57).
- Maslow, A. (1966). *The psychology of science: A reconnaissance*. Harper & Row. <https://books.google.it/books?id=qitgAAAAMAAJ> (cit. on p. 94).
- Mintzberg, H. (1979). *The structuring of organizations: A synthesis of the research*. Prentice-Hall. <https://www.nrc.gov/docs/ML0907/ML090710600.pdf> (cit. on pp. 2–4, 56–59).
- Russell, S., & Norvig, P. (2020). *Artificial intelligence: A modern approach* (4th ed.). Pearson. <http://aima.cs.berkeley.edu> (cit. on pp. 25, 26, 56–59).
- Shafer, G. (1976). *A mathematical theory of evidence*. Princeton University Press. <https://books.google.it/books?id=wug9DwAAQBAJ> (cit. on p. 34).
- Siegel, E., Glaeser, E. L., Kozyrkov, C., & Davenport, T. H. (2020). *Strategic analytics: The insights you need from harvard business review*. Harvard Business Review Press.

- <https://store.hbr.org/product/strategic-analytics-the-insights-you-need-from-harvard-business-review/10358> (cit. on pp. 5, 57, 59).
- Simon, H. (1957). *Models of man: Social and rational; mathematical essays on rational human behavior in society setting*. Wiley. <https://books.google.it/books?id=XhBKyAEACAAJ> (cit. on pp. 2, 58).
- Simon, H. (1997). *Models of bounded rationality: Models of bounded rationality / simon, herbert alexander. - cambridge, mass. [u.a.] : Mit press, 1982- vol. 3*. MIT Press. <https://mitpress.mit.edu/9780262519434/models-of-bounded-rationality> (cit. on p. 19).
- Sutton, R., & Barto, A. (2018). *Reinforcement learning, second edition: An introduction*. MIT Press. <https://mitpress.mit.edu/9780262039246/reinforcement-learning/> (cit. on p. 36).
- Weber, F. (2023). *Artificial intelligence for business analytics*. Springer Fachmedien Wiesbaden. <https://doi.org/10.1007/978-3-658-37599-7> (cit. on p. 56).

## REPORTS

- Elsevier. (2018). *Artificial intelligence: How knowledge is created, transferred, and used*. <https://www.elsevier.com/research-intelligence/resource-library/ai-report> (cit. on pp. 27, 30).
- IoT Analytics. (2023). *Quarterly trend report: What CEOs talked about in Q1/2023*. <https://iot-analytics.com/product/trend-report-what-ceos-talked-about-in-q1-2023/> (cit. on p. 29).
- Maoz, M. (2013). *How IT should deepen big data analysis to support customer-centricity*. Garner. <https://www.gartner.com/en/documents/2531116> (cit. on pp. 6, 56, 76).
- Siebert, M., Kohler, C., Scerri, A., & Tsatsaronis, G. (2018). *Technical background and methodology for the Elsevier's s artificial intelligence report*. <https://www.elsevier.com/?a=829143> (cit. on p. 27).


## WEBSITES

- Bornet, P. (2022). *Is decision intelligence the new AI?* Retrieved July 4, 2023, from <https://www.forbes.com/sites/forbestechcouncil/2022/05/25/is-decision-intelligence-the-new-ai/> (cit. on p. 56).
- Butters, J. (2023). *Highest number of S&P 500 companies citing “AI” on Q1 earnings calls in over 10 years.* Retrieved June 2, 2023, from <https://insight.factset.com/highest-number-of-sp-500-companies-citing-ai-on-q1-earnings-calls-in-over-10-years> (cit. on pp. 27–29, 64).
- Gadd, L. (2020). *AI-based citation evaluation tools: Good, bad or ugly?* Retrieved April 14, 2023, from <https://thebibliomagician.wordpress.com/2020/07/23/ai-based-citation-evaluation-tools-good-bad-or-ugly/> (cit. on pp. 12, 14).
- International Organization for Standardization. (1987). *Iso 9075:1987 - information processing systems — database language — sql.* Retrieved April 14, 2023, from <https://www.iso.org/standard/16661.html> (cit. on p. 19).
- Marcus, G. (2023). *The sparks of AGI? Or the end of science?* Retrieved March 30, 2023, from <https://garymarcus.substack.com/p/the-sparks-of-agi-or-the-end-of-science> (cit. on p. 25).
- Matiisen, T. (2015). *Demystifying deep reinforcement learning.* Retrieved May 7, 2023, from <https://neuro.cs.ut.ee/demystifying-deep-reinforcement-learning/> (cit. on p. 36).
- Mazzanti, S. (2020). *SHAP explained the way i wish someone explained it to me — towards-datascience.com.* Retrieved June 1, 2023, from <https://towardsdatascience.com/shap-explained-the-way-i-wish-someone-explained-it-to-me-ab81cc69ef30> (cit. on p. 50).
- Merriam-Webster. (1831). *Definition of model.* Retrieved July 11, 2023, from <https://www.merriam-webster.com/dictionary/model> (cit. on p. 58).
- Mintzberg, H. (2015). *The soft underbelly of hard data.* Retrieved July 5, 2023, from <https://mintzberg.org/blog/hard-data> (cit. on pp. 58, 59).
- Mollick, E. (2023). *On-boarding your AI intern.* Retrieved August 8, 2023, from <https://www.oneusefulthing.org/p/on-boarding-your-ai-intern> (cit. on p. 78).
- Quinn, B. (2023). *Popular tools - artificial intelligence tools for research and writing - guides at texas tech university.* Retrieved April 15, 2023, from <https://guides.library.ttu.edu/artificialintelligencetools/popular> (cit. on p. 12).

- Sibony, O., Lovallo, D., Kahneman, D., & Klein, G. (2010). *Strategic decisions: When can you trust your gut?* Retrieved June 3, 2023, from <https://www.mckinsey.com/capabilities/strategy-and-corporate-finance/our-insights/strategic-decisions-when-can-you-trust-your-gut> (cit. on p. 58).
- Simens, M. (2021). *How your data is used to improve model performance*. Retrieved September 15, 2023, from <https://help.openai.com/en/articles/5722486-how-your-data-is-used-to-improve-model-performance> (cit. on p. 79).
- The Open University. (2023). *Pros and cons of bibliometrics*. Retrieved April 2, 2023, from <https://www.open.ac.uk/library-research-support/bibliometrics/pros-and-cons-bibliometrics> (cit. on p. 13).
- Tiku, N. (2022). *Google engineer Blake Lemoine thinks its LaMDA AI has come to life*. Retrieved April 20, 2023, from <https://www.washingtonpost.com/technology/2022/06/11/google-ai-lamda-blake-lemoine/> (cit. on p. 25).

## SOFTWARE

- Aria, M., & Cuccurullo, C. (2017). *Bibliometrix: An R-tool for comprehensive science mapping analysis* (Version 4.1.2). <https://www.bibliometrix.org/> (cit. on p. 12).
- Chase, H. (2022). *LangChain: building applications with LLMs through composability*. <https://www.langchain.com/> (cit. on p. 82).
- Elsevier. (2004). *Scopus: Expertly curated abstract & citation database*. <https://www.scopus.com> (cit. on p. 17).
- OpenAI. (2021). *ChatGPT: A large language model for natural language processing*. <https://chat.openai.com> (cit. on p. 25).
- Palantir Technologies. (2023). *Palantir artificial intelligence platform*. <https://www.palantir.com/platforms/aip> (cit. on p. 82).
- Scite. (2018). *Scite: See how research has been cited*. <https://scite.ai/> (cit. on p. 13).
- Van Eck, N., & Waltman, L. (2010). *VOSviewer: A computer program for bibliometric mapping* (Version 1.6.19). <https://www.vosviewer.com> (cit. on p. 12).



**T**HIS thesis was typeset using  $\LaTeX$ . The main layout is inspired by the beauty, quality and variety of the dissertations shared by the community over the years. The colors and cover layout follow the visual identity manual of the University of Padova. The body is set in 12 point Minion 3. Other fonts used are  $\TeX$  Gyre Pagella for chapter numbers, Fira Code for monospaced text, and Libertinus Math for mathematical sections. The template used for this thesis will be released under a permissive license at <https://github.com/AlphaJack/masterthesis>. The three illustrations above represent Padova, Clermont-Ferrand and Milano, the main cities where I lived, studied and worked to achieve my Master of Business Administration.



